

## The Zone System and Digital Photography

The Zone System is a precise and powerful method for controlling not only exposure, but also to control the final appearance of a photographic print.

The Zone System was introduced to a larger audience in 1948, when Ansel Adams (1902-1984) published the second volume of his *Photography Series*, titled the *The Negative*. (The first was titled *The Camera [and Lens]*, and the third, *The Print*).

In this ground-breaking trilogy, Adams assumes that the photographer uses a large format view camera and single sheet film and takes control over the entire workflow right down to mixing his or her own chemicals. Newer books on the system often focus on the use of roll film and standard chemicals, and how to deal with the problems that arise from not being able to develop each negative individually.

The Zone System has been in use for roll film photography for many years. However, it was not immediately obvious how it could be used in conjunction with digital photography.

After a lot of trial and (mostly) error, it was found that the Zone System is just as useful with digital as it is with film. However, digital photography is different from using film, so the Zone System must be *adapted* to be useful to digital. This note contains some findings and thoughts about how the Zone system can be adapted.

This treatment owes a lot to Norman Koren's (2005a) excellent article on the simplified Zone System. However, while Koren's adaptation deliberately limits itself to controlling exposure, this method is to emulate Adams' original scheme where the system encompasses the entire workflow, up to and including the final print.

### Prerequisites

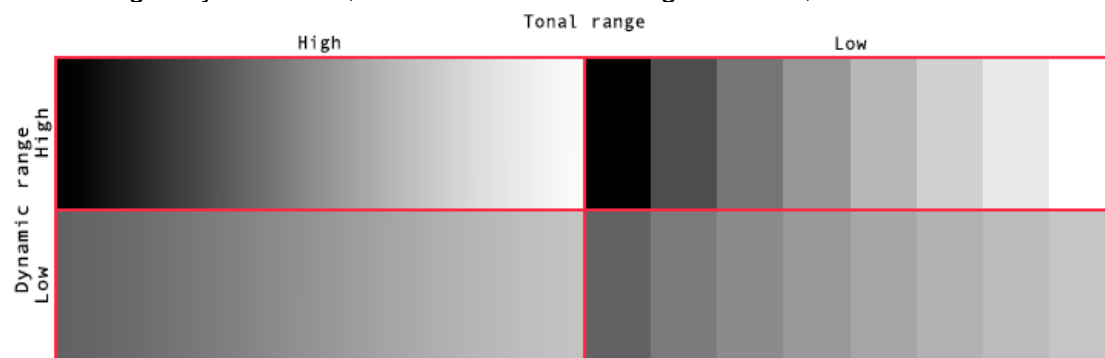
If this is the first time you have heard about it, what follows will probably not make much sense. For an introduction to the Zone System, please read a reprint edition of *The Negative* (Adams 2002) or pick up a modern textbook about it, such as *The Practical Zone System* by Chris Johnson (2006).

This also assumes that the reader is familiar with such terms as aperture (f-number), shutter speed, film/sensor speed (ISO) and Exposure Value (EV) and their relationship, and have a basic knowledge about digital processing, including monitor and printer calibration, RAW-conversion, reading histograms, and gamma. To have read the book by Chavez and Blatner (2009) suggested under further reading is strongly recommended, whether one uses Photoshop for processing or not.

### Motivation

The Zone System is about capturing and printing photographic images in such a way that the original scene is reproduced as well as it is possible within the limitations of the photographic medium. To achieve this, the photographer must adopt a workflow that loses as little as possible of the scene's dynamic and tonal range.

Dynamic and tonal ranges are related, and are sometimes referred to interchangeably. However, as is illustrated in the figure below,



they are independent. We can have image data that have a high or low tonal range,

and also image data with a high or low dynamic range, and any combination of the two.

The *dynamic* range is a measure of how much the darkest bits in a recorded scene differ from the lightest. It is usually expressed in EV, where an increase in luminance equal to 1 EV representing a doubling of the light.

The *tonal* range is a measure of the granularity we use when real world tones are mapped onto an recording medium. With a high tonal range, gradients are smooth. With a low tonal range, the gradients are abrupt, and we see an image defect usually called banding.

The photon wells in the sensor in a digital camera count photons. Within the dynamic range of the sensor, each photon well will produce a charge that is directly proportional to the amount of light that strikes it. This means that within this range, digital sensors are strictly linear devices.

The lowest possible output from the sensor is produced when the sensor is not exposed to any light. This is called the sensor's "noise floor" and is greater than zero. From the noise floor, the charge will continue to accumulate until the photon well reaches its capacity.

In digital photography, the dynamic range is the difference between the noise floor and the full well charge capacity.

Likewise, tonal range is a function of the number of bits available to represent tones. With too few bits available, tonal range will suffer.

The camera records RAW sensor data. When RAW sensor data are adjusted for contrast and the resulting tone values are (usually) compressed and redistributed to fit into the dynamic range and tonal response of the rendering medium. At this point, one may run into problems if the data is not suited for these transformations. For example, if our exposure has clipped the highlights or put the shadows into the noise floor, it will lose detail in those areas. If we have wasted tone levels by underexposing there may be too few bits containing actual data available for a high tonal range, and image quality will suffer.

The motivation behind a digital Zone System is to have a workflow that creates image data that preserves both highlight and shadow detail, and at the same time makes optimal use of the bits available for recording.

## Limitations

According to Fred Parker's *Ultimate Exposure Computer*, real world luminance goes from EV -6 (night, away from city lights, subject under starlight only) to EV +16 (subjects in bright daylight on sand or snow). This is a span of 23 EV. The dynamic range of the human eye, from the darkest we can perceive, to the brightest light we can tolerate, is about one to one million. Expressed in EVs, this is equivalent to around 20 EV ( $\log_2 1000000 = 19.9$ ).

The dynamic range of recording and rendering mediums used for photography is much smaller than this.

Unfortunately, it has been difficult to find good data about the dynamic range of film. Kreunen (2003) says that *Fuji Realia 100* colour negative film has a dynamic range of 15 EV. However, his method involves making a set of exposures, and using normalised data where he includes the entire non-linear region in the result. Doubts have been raised as to whether this method produces valid results.

According to Koren (2005a), *Kodak Supra 100* colour negative film has a dynamic range of 10 EV, and *Kodak Ektachrome 100VS* colour slide film has a dynamic range of about 5.6 EV. These figures are based upon interpreting the tone curve for the films in Kodak data sheets, not from actual measurements.

There are reports about famous photographer Bruce Barnbaum capturing an dynamic range of 18 EV on black and white negative film, but he is a master of development and does not use standard techniques.

In a discussion on the Photo.net bboard in 2001, several different figures are given for black and white negative film, but nobody reveals how the figures are measured. The most believable figure comes from John Hicks, who says that Ilford HP5+ has a dynamic range of at least 14 EV.

DPreview measures dynamic range by taking a *single* shot of a backlit calibrated Stouffer 39 step wedge that span 13 stops. Dynamic range is then measured by counting the steps that is reproduced as a shade of grey. It is thought that this method is more appropriate than the method suggested by Kreunen (2003).

Using this method, DPreview finds that the Nikon D3 has an in-camera JPEG dynamic range of 8.6 EV (1:390) at ISO 200, and RAW dynamic range of 12 EV (1:4000).

Digital techniques such as HDR has dramatically expanded the photographers ability to capture a large dynamic range, by combining a number of images with different dynamic ranges, an arbitrarily large dynamic range can be recorded.

Photographs may be rendered on two different devices: screens, and paper. Monitor screens come in different grades, and while an inexpensive consumer monitor may have a dynamic range as low as 1:100, a professional grade monitor may be capable of reproducing a ten times larger dynamic range. Paper and inks also comes in different grades, but some print papers and quality inks may be capable of producing a dynamic range of 1:250. [Source: Martin (2006)]

These observations are summarised in the table below:

Dynamic Range			
Medium		Range	EV
	Real world	1:8400000	23.0
	Human eye	1:1000000	19.9
Recording	Colour slide film	1:50	5.6
	JPEG image data	1:400	8.6
	Colour negative film	1:1000	10
	RAW image data	1:4000	12
	B&W negative film	1:16000	14
	HDR image data	?	?
Rendering	Monitor (consumer)	1:100	6.6
	Print paper	1:250	8.0
	Monitor (pro. grade)	1:1000	10.0

Because the mediums and materials available to the photographer today is much more limited in the range they can record and render than the range of values present in real world, the photographers must bridge this gap. The Zone System does this by providing the photographer with a systematic way to determine correct exposure and further processing of images.

### 3. The Basics

The Zone System was invented by Ansel Adams and Fred Archer around 1940 as a simple and straightforward method for controlling exposure and producing fine prints. Before Adams and Archer, the photographic industry had more or less standardised on aperture (f-number), shutter speed (1/sec.), and film sensitivity (ISO) as the three main parameters for exposure control, but there existed no systematic approach for

determining how these should be set to produce the best possible print. The Zone System remedied this.

The original Zone System takes the tones that appear in a black & white photographic print, and divides this into eleven discrete “zones”, from Zone 0 (total black) to Zone X (pure white). However, of the 11 zones, only 9 can hold information. Zone 0 and Zone X are “off the scale”. Zone 0 represents unexposed silver halide (dark current noise if we are talking about digital images). Zone X represents specular highlights that completely fog the negative (or makes photon wells overflow). In principle, Zone 0 and Zone X can span an infinite number of levels, while Zone I through IX is evenly spaced throughout the dynamic range captured.

### **Sensor Speed: ISO**

To control the exposure of an image on to a light sensitive imager (i.e. film or sensor), we need to know how sensitive it is to light. Light sensitivity of the imager is measured on a scale defined by the *International Organisation for Standardisation* (ISO). This measurement is referred to as the “speed” of the film or sensor, and the unit of measurement is known as “ISO”.

How to measure this value is defined by a series of ISO standards. For instance, the standard known as *ISO 6:1993* (ISO 1993) tells how to measure ISO speed for ordinary black-and-white negative film (as a function of density over base fog), and *ISO 12232:2006* (ISO 2006) does the same for the electronic sensors used in digital cameras. For digital cameras, the standard defines sensor speed in such a way that a digital sensor and a film that both are rated with the same ISO value will behave in a similar way when exposed to the same amount of light. This means that light meters and exposure techniques that works well with film, also will work when the photographer is using a digital camera. (Note: *ISO 12232:2006* is not available for free on-line, but *Kodak application note MTD/PS-0234* (2009) gives a good summary.)

ISO actually defines two parallel scales, one linear (arithmetic) scale and one logarithmic scale. This is because the ISO scale was created in 1987 by merging two older scales known as “ASA” and “DIN”. The ISO linear scale corresponds to the “ASA” scale, and the ISO logarithmic scale corresponds to “DIN” scale.

In the ISO linear scale, doubling the value indicates a doubling in speed. In the ISO logarithmic scale, adding 3 to the numeric value indicates a doubling in speed. ISO suggests that both values are written, separated by a slash (/) and that the logarithmic value is marked with a degree (°) symbol. Given this notation, a film or digital sensor with a speed equal to ISO 200/24° is twice as fast (sensitive) as a film or sensor with a speed equal to ISO 100/21°.

In practice, the ISO logarithmic scale is rarely used, and you will instead see the film or sensor speed measured in the linear values, e.g. “ISO 100”, “ISO 200”, and so on. In (most) digital cameras, ISO works like this: The electronic sensor collects photo-electrons (produced by photons colliding with the image sensor and knocking electrons loose from the valence band) into photo wells. This process accumulates charge in the photo wells. The amount of charge that has accumulated in a photo well indicates the amount of light that has illuminated the photo well. Photo wells have limited capacity, when a photo well has accumulated the maximum charge it can hold, supplying more light will only result in non-linearity (burnout), and maybe also spillage into adjacent pixels (blooming). The native quantum efficiency (the percentage of time an incident photon actually results in a photo-electron) and the charge well capacity determine the so-called *base ISO* of the sensor.

The accumulated charge from each photo well can be converted into a voltage. This voltage is fed into an Analog-to-Digital Converter (ADC) to produce the RAW digital data that (eventually) will result in a digital image.

Lowering the sensor's ISO *below* the base ISO is not really practical. (Some cameras let the photographer do this, for example calling it “Lo-1” to indicate that it is off the scale. However, this is not recommended, and may result in loss of tonal range.)

All digital cameras let the photographer *increase* the ISO above the base ISO. One way to implement this is to amplify the voltage read from each photo well before it is fed into the ADC. E.g.: If at ISO 100 the signal is 100 mV, we can get ISO 200 by using an amplifier to boost it to 200 mV. For ISO 1600, we can five-double it to 1600 mV, and so on.

The analogue amplification ensures that the full input voltage range of the ADC is utilised. This means that the full bit depth of the ADC converter is used, so we get the maximum tonal range from the ADC. However, when we amplify the signal, we also amplify any noise, so we will lose some image quality.

Instead of analogue amplification, we can use an ADC converter with higher than required bit depth and just rescale the output data digitally afterwards. You lose one bit for each multiplication by two (1 f-stop). If we have a sensor with ISO 100 as base ISO, using a 16-bit ADC would allow us to crank the speed up to ISO 1600 and still output the standard 12 bits per photo well, without the need for any analogue amplification. However, digital multiplication also magnifies noise, just like analogue magnification.

Signal processing can remove noise from a digital file. Unfortunately this may also remove some fine detail. However, when used in moderation, the noise removal may be acceptable.

The upper limit for high ISO in a digital cameras are set by the manufacturer to a value that gives the user some flexibility when taking photographs in poor light, but not so high that the resulting images (due to noise and/or noise removal) is of very inferior quality.

Exposure Value: EV

Exposure Value (EV) is defined by ISO (1974) through the following equation:

$$2^{EV} = N^2/t = (L \times S)/K$$

Where *EV* is the exposure value, *N* is aperture expressed as a f-number, *t* is the shutter time in seconds, *L* is luminance expressed in candelas per square meter, *S* is linear ISO speed, and *K* is the reflective meter calibration constant.

Inverting the formula, we get:

$$EV = \log_2(N^2/t) = \log_2((L \times S)/K)$$

The EV-scale is logarithmic. Going from EV to the next (e.g. from EV 10 to EV 11) represents a doubling in luminance or 1 f-stop.

Many texts on the zone system claim that the difference between adjacent zones is 1 EV (1 f-stop). This is not true. EVs and f-stops express relative difference in levels of light present in a scene. Zones express relative difference in levels of density present in a photographic print, which may or may not reproduce exactly the relative levels of the original scene.

Zones are *not* f-stops or EVs.

The original Zone System was slanted towards printing black and white positives from negative film. The core idea behind the original system was to be able to produce negatives with a full tonal range could be printed on grade 2 (soft/normal) paper, so having eleven zones with ten discrete steps between them mapped onto a dynamic range equal to  $2^{10} = 1024$ , or two more zones than the 1:250 ratio that characterised B&W print paper.

### **The Zone Scale**

The table below outlines the core of the Zone System, the Zone Scale of tone values. Each zone in the Zone Scale corresponds to a specific visual representation of tonal values in a photographic *print*, going from total black to pure white with middle grey (18 % reflectance) in the middle.

Zone descriptions. (Adapted from Adams 2002, p. 52-60)

DR	TR	Spot	Zone	Description
			0	Total black. Complete lack of density, other than dark current noise (or film base density + fog in the case of a film negative). Should appear as total black in the print.
Low			I	Near black, no detail. Effective threshold. First step above complete black in the print. Slight tonality, but no texture.
			II	Dark grey-black. First suggestion of texture. Very dark details in shadows. Deep tonalities, representing the darkest part of the image in which some slight detail is required.
		Shadow	III	Very dark grey. Dark textured bark on shadow side of tree. Average dark materials. Good texture and detail can be seen. This is where you will want to place shadow details.
			IV	Medium-dark grey. Average dark green foliage, shadow side of skin, dark stone, landscape shadow. Details plainly visible. This where you want to place the shadow side of Caucasian portraits in sunlight.
Mid	Text- ural range	Average	V	Middle grey. Standard Kodak 18 % grey reflectance card. Clear northern sky (panchromatic rendering), dark skin, grey stone, average weathered wood. Excellent detail visible.
			VI	Rich mid-tone grey. Caucasian skin in sunlight, light stone and sand, shadows in snow in brightly sunlit snowscapes. Sharp fine detail visible.
			VII	Off white or bright light grey. White with texture, very light skin, silver hair, weathered white paint, snow with acute side lighting. Highest Zone that will still hold good details.
High			VIII	Almost white (not blank whites). Textured snow in sun, reflected highlights on Caucasian skin. Delicate texture and some gradation exist, but no detail.
			IX	Nearly pure white without texture (must be compared to pure white to tell difference). Glaring white surfaces, snow in flat sunlight. No detail or significant texture visible.
			X	Pure white. Specular highlights, glares or light sources in the picture area. Danger of photon well overflow. Rendered as the maximum white value of the paper surface.

The range of zones which convey definite qualities of texture and the recognition of of substance is known as the *textural range* (TR) and encompasses Zone II through VIII.

Three zones are particularly important: The zone used to measure for the shadows for film (Zone III), the zone used for 18 % grey exposure (Zone V), and the zone used to measure the highlights for digital (Zone VII). These zones are labelled “shadow”, “average”, and “highlight” in the zone description table.

Many users of the Zone System find the system easier to work when they disregard the Zones 0 and X, which mean only working with 9 zones. These can be nicely divided into three tone ranges (Low, Mid and High), with Zone V in the middle of the scale. Adams refers to this (Zone I to Zone IX) as the *dynamic range* (DR).

## Tones, Zones, Gamma and the Histogram

Every digital image rendering process consists of two steps. The first is the computing of the linear luminance values from the charge accumulated by the sensor. These are the luminance values recorded in the RAW data file produced by the camera. The second is the mapping of the computed values to the values appropriate for being viewed by humans. This process is known as *tone mapping*.

The Zone System wedge below and its normalised pixel values are lifted from Norman Koren's *A simplified Zone System for making good exposures*. It shows normalised pixel values as a decimal fraction of 1, and in decimal and hexadecimal absolute values, for the nine zones of his simplified system. If you are using a monitor set up for gamma = 2.2 (the default for Microsoft Windows and most newer Apple systems), and your browser follows W3C recommendations for colour rendering, the tones should be fairly accurate.

To be able to use the Zone System on a digital camera, you need to familiarise yourself with how your meter corresponds to your histogram. As a start, make a careful reading of some uniform surface, photograph it, and look at its histogram. You will see a narrow column. With most meters, this column will appear to the left of the centre of the histogram.

The two images below show how the Zone Scale applies to photographs. Below each image is its histogram. The histogram shows how the tones in the image map onto the Zone Scale.



The image on the left, showing jazz drummer Billy Cobham in action, is a photograph where the tonal range is spread fairly evenly across all nine zones.



There is just a slight peak in Zone I (dark background, black hair and beard). The drummers' forehead lies within Zone V, which is the Zone most photographers would want to use for dark skin, and the rest of the tones in the photograph spread out around this.

In the image on the right most of the data is in the low (Zone I, II, III) and high (Zone VII, VIII, IX) tones. Almost no area of the subject lies in the mid tones. This is a typical high contrast image, which is characterised by an U-shaped histogram.

In low contrast images, the peak in the histogram will be in the mid-tones. There are also tonal styles called "high key" and "low key", where the histograms peak will be in the high or low tone area of the histogram, respectively.

#### **4. Single Object Placement**

Most photographers are familiar with the use of the built-in light meter in their camera. The camera's meter measure the light that is reflected from the scene, and use this information to compute the exposure for the scene. The actual process may entail making more than one measurement and combining them (matrix/evaluative), weighing them (e.g. centre-weighted, partial or spot), but the measurements will always be used to compute a combination of aperture, shutter speed and ISO that presumably will give the "best" exposure for the metered scene.

With this type of metering, in a "normal" scene the meter should ideally be used to meter whatever reflectance the meter is calibrated for. Textbooks will usually tell you that the meter is supposed to meter "middle grey", or the tones in the scene should average out to middle grey or "18 % reflectance". This is not the way reflected light meters really work (see my note on exposure meters for details), but given the exposure latitude of modern films and sensors, it is close enough for jazz.

Of course, some scenes may not be "normal". An experienced photographer will recognise this, and know how to adjust for metering error by dialling in exposure compensation (EC). In digital photography, the photographer may even refer to the camera's histogram, or flashing clipping warning, to determine the amount of EC to use.

This usually works well, and for a lot of photography, this is all you need. However, it does not give the photographer much control over how the various tones and colours in the scene will be rendered in the final print. It is for this type of control the Zone System was created.

For the Zone System, light measurements are always done with a spot meter, preferably one with a one-degree coverage. A spot meter is essential because you will be measuring specific portions of the scene, and then "placing them" in a specific zone. It is this placement, and not the meter's reading, that determines exposure. When metering a scene, you should start by metering off a standard 18 % grey card. This gives you a reference point for Zone V.

With the Zone System, you will deliberately measure different parts of the scene and noting how they differ. The spot meter reading will always report the exposure that will render that part of the scene as middle grey (Zone V). However, we do not always want things to appear as middle grey. Therefore, we need to determine how we want the subject to appear in the final print, i.e. to decide what zone the object should ideally appear in. Ansel Adams called this process "visualisation". After doing a visualisation of the appropriate zone, we "place" the subject in the desired zone by modifying exposure up or down the scale to move the object from the measured to the desired zone. For instance, to place an object metered in Zone V in Zone VI, we use a +1 EV exposure adjustment.

When adjusting exposure, you are determining the tone values an object will have in photographic print independent of what tone values it has in real life. An experienced Zone System practitioner is capable of mentally visualise the change in an object's tone values as he or she moves it up and down the zone scale at the time of exposure.



## Some Examples

It is simpler to do this than to explain it, and some examples will make this clear.

Let us say we want to do a studio portrait of someone with light skin. We shall place the lit side of the subject's face (the most important portion of a portrait) in Zone VI. First, meter the face. What the meter gives you is the setting for a Zone V face. You will have to give the face *more* exposure than indicated by the meter (more light on the sensor) to place it in Zone VI. Opening the lens one f-stop (+1 EV) from the metered exposure will have this effect.

Now, let us move on to photographing an aubergine (egg plant). It's not black, but it's dark. Maybe you would like it to appear in Zone III in the final print. Again, your spot meter indicates exposure for Zone V. By closing down two f-stops (-2 EV), the aubergine will be placed in Zone III.

To recapitulate – the three things you need to know to use the Zone System to place single objects are:

The Zone scale is a progressive series of tone values. Each value is the equivalent of one full f-stop or one EV step.

The spot meter provides exposure readings for Zone V, giving you a correct exposure for a known Zone.

By adjusting exposure you can place the object in any Zone. On a calibrated monitor, and in the final print, the object will assume the tone value of the Zone in which it is placed.

When coming to grips with the Zone System, it's a good idea to practise metering all sorts of objects, deciding what Zone the object should be placed in, and then doing the necessary mental adjustment to effect this placement, i.e. a *visualisation* of the tone shift as it will appear on a calibrated monitor and in the final print.

With a digital camera, experimenting is much cheaper than using film, you can practice this and by looking at the results on your computer's monitor until using metered values to affect Zone placement becomes second nature for you.

However, some scenes may be more complex and contain multiple important objects.

## 5. Complex Scenes

In the preceding examples, the the Zone System was used to place a single object. In the real world, our scenes often contain many objects. Often there is not a single adjustment value that will place *all* the parts of the scene where we want them.

For instance, let us assume we are metering a landscape. We see some good shadow details we want to record, and place them in Zone III (i.e. -2 EV). Now we read our highlight value and find that after making an -2 EV adjustment, it will be placed in Zone V. But we want it to appear in Zone VII *without* affecting the placement of shadow detail in Zone III. How can we do this?

That answer is that we can also control Zone placement through processing.

## Background

Ansel Adams discovered that he could influence negative contrast, and therefore the tonal range in a scene, by adjusting development. He referred to increasing contrast as *expansion*, and reducing contrast as *contraction*.

Adjusting development times works because development affect highlights more than shadows. By “pushing” development, Adams could move his highlights up one zone or two, while keeping the shadows in the zone he exposed for. By “halting” development, the reverse is possible. Highlights can be moved down one zone or two, while shadows are much less affected. This observation is the basis for the negative film photographer's adage: “Expose for the shadows, develop for the highlights.”

If we return to the landscape remember that the highlight value, after placing the shadows in Zone III, would move to Zone V. What if one wants to place it in Zone VII? Adams would accomplish this by adding a specific amount of development time that could be found by looking up the appropriate table. In this case, he would use what he called *N+2* development, which meant normal plus additional development to move the highlights up two zones without affecting the shadows.

The precise control over the dynamic range of the negative that controlled exposure combined with adjusted development times afforded, was what made the Zone System so powerful, compared to all other methods for controlling the final appearance of a photographic print.

### **Doing it Digitally – First Approach**

With digital workflow, we can no longer play around with development times. But there are two things that make digital photography very well suited for the Zone System. One of them is that unlike roll film, we can process each digital shot individually (so in that sense, it is just like sheet film). The other is our ability to move specific tones between zones through digital image processing.

Assume that we have photographed the landscape with a digital camera. As a first approach, let us also assume that we have determined exposure in the same way as explained in the background paragraph above. I.e., we have exposed to place shadow detail in Zone III (which is where we want them), but for that reason ended up with highlights in Zone V (and we want to place highlights in Zone VII). We can do this by using the *exposure* control in Photoshop ACR to *expand* the tonal range by moving the white point up two zones while at the same time keeping the black point fixed.

This means that by doing a simple exposure adjustment in ACR, it is possible to duplicate what Ansel Adams accomplished by adjusting development times.

However, this is not the whole answer. Under-exposing the highlights is necessary if we want to expand the dynamic range of the negative through a *N+2* development time or digital processing, but doing this type of adjustment in digital photography has unfortunate side-effects.

### **The Troublesome Highlights**

Unlike film, which has great exposure latitude at the highlight end, digital processing is very unforgiving in the case of over-exposure. Detail that is lost through over-exposure is clipped and lost forever. For this reason, we never want to *overexpose* highlights to the point off photon well overflow – not even in a single colour channel. This introduces clipped areas with absolutely no detail in the channel.

So why not simply place shadow detail in the lowest possible Zone? That should at least do the most to contain our highlights within safe limits.

The problem is that if you do this, you are wasting a lot of the bits the camera can capture. This problem has been noted by a number of experts on digital imaging, a good explanation of this can be found in a short Adobe white paper by Bruce Fraser (2004) at [http://www.adobe.com/digitalimag/pdfs/linear\\_gamma.pdf](http://www.adobe.com/digitalimag/pdfs/linear_gamma.pdf).

Unlike the eye and film, digital sensors measure light linearly. If the RAW file has a bit depth equal to 12 bit, a maximum of  $2^{12} = 4096$  different levels are possible. If those 4096 levels could be portioned equally over a 9 EV range, each EV should have  $4096/9=455$  levels to itself.

Unfortunately, this is not how things work out in practice. The linear capture of the camera's sensors means that if we try to capture a 9 EV range, corresponding to 9 zones, half of the 4096 levels (2048 levels) are devoted to Zone IX, half of the remainder (1024 levels) are devoted to Zone VIII, half of the remainder (512 levels) are devoted to the Zone VII, and so on. Zone V is represented by 128 levels, Zone III

by 32 levels, and the extreme shadows in Zone I is represented by only 8 different levels.

This means to underexpose to avoid clipping the highlights, you are running a significant risk of introducing noise and banding in the mid-tones and shadows. When you, as a result of underexposure, try to open up the shadows in the RAW conversion, you have to spread those 8 levels in the darkest stop over a wider tonal range, which exaggerates dark current noise and increases banding (quantification noise).

A much-quoted article by Michael Reichmann (2003), titled *Expose (to the) Right* argue that you should use your camera's histogram to evaluate the light in the scene, and push exposure towards over-exposure so that the histogram moves as far as possible to the *right edge* (without moving so far that highlights are blown as indicated by your camera's clipping warning) – hence the article's title.

***With digital: Expose for the highlights, process for the shadows.***

There are some problems with Reichmann's histogram recommendation.

The camera shows a gamma-adjusted histogram. You would not like the look of a linear histogram, because that would show almost all the data clumped towards the darker (left) end. Most digital cameras apply a fairly strong S-curve to the RAW data so that the gamma and tone adjusted data have a somewhat film-like response. The result is that the on-camera histogram and flashing clipping warning will tell you that the highlights are blown when, in fact, they are not.

But if you don't have a spot-meter, using the histogram and clipping warning in the way Reichmann recommends will often help you get less noise and more bits for your highlights than just using the combination of ISO, aperture and shutter speed suggested by your camera's metering.

However the histogram won't replace a spot meter if you want to make use of the fullest possible dynamic range from a modern digital camera.

**Exposing for the Highlights**

If we sample various areas spread around the scene with the meter, we will probably discover that the dynamic range of the scene exceeds the 8-9 zones that is our digital camera's maximum dynamic range. For instance, the difference between some bright white clouds in the sky and the deep shadows under a bush may be as much as 11 or 12 EV. There is no way to record such a dynamic range with a single exposure. If you want to capture it, you need to make several exposures, as described in the short section on high dynamic range (HDR).

But assuming that we can live with only capturing some portion of the dynamic range of the scene, how should we determine exposure?

Firstly, we survey the scene, and identify the area containing highlight detail that we would like to reproduce in the final print.

Spot meter that area. The meter will record a value that translates into an aperture, shutter time and ISO combination that would place this part of the scene in Zone V.

Visualise the Zone we want highlight detail to appear in, in our print. Let's say we want it to appear in Zone VII. Moving the area from Zone V to Zone VII requires an exposure adjustment equal to +2 EV.

Make the adjustment to the camera's setting, and expose the image.

Given the same scene as in our previous example, moving highlight detail from Zone V to Zone VII will have the side-effect of moving shadow detail from Zone III to Zone V.

After doing all this, we now have recorded a file where the highlights should be where we want them. However, our shadows have now become too light.

## Processing

There now exist RAW converters that support the Zone System directly. The workflow description detailed uses Adobe Photoshop CS and ACR on Windows/XP. These are the tools used. With other tools, or on another operating system, the workflow may be slightly different.

This article was written using first version of Photoshop CS (aka PS 8). Some of controls have changed names and others have been added in ACR distributed with later versions of Photoshop. An attempt has been made to put the new names in square brackets for readers that use later versions.

To start by opening the RAW image file into ACR. The most important controls here are the *exposure* and *shadows [black]* sliders under the *Adjust [Basic]* tab. These two sliders give us the same kind of control over tone mapping as Ansel Adams' achieved when he adjusted the development times for his films. I.e. these sliders allow us to expand or contract the tonal range of the image.

After adjusting the dynamic range of the image in ACR, I continue in PS CS where sometimes it's necessary to apply a tone curve using the *Curves [Tone Curve]* dialogue, and concludes with sharpening just prior to printing. In general, Try to do as many adjustments as possible (except final sharpening) within the RAW converter, before going to PS CS.

After saving the adjusted file, it can be printed. Only use inks and papers made to match the printer along with the appropriate printer profiles for best results.

Here is more detailed description of how I work with the image in ACR. If you are using ACR (Adobe Bridge )with Photoshop CS4, you may toggle the clipping warning by pressing the letters «U» (shadows) and «O» (highlights) on the keyboard. In older versions of Photoshop CS, holding down the *alt*-key while moving the slider will bring up a clipping warning.

First pull the RAW image file into ACR, and set the white balance (if necessary). The *temperature* slider is used for major adjustments of colour temperature along the blue-yellow axis, and the *tint* slider is used for minor adjustments along the magenta-green axis.

Then use the *exposure* slider to work on the highlights. (With a perfect exposure for the highlights, there shouldn't anything to adjust, but even with very careful metering; it's often found to need for a little white point adjustment at this stage.) Moving the slider to the right *expand* the tonal range of the image, moving it to the left, *contract* the tonal range. Moving the white point down two zones with ACR exposure slider exactly duplicates what Adams referred to as *N-2* development. A few specular highlights may be desirable and is often unavoidable, but if some large area of the image is clipping in one or more channels, uniform tone will take the place of a gradient and the effect will usually be ugly.

With a properly exposed digital image, it is usually shadows placement we want to adjust. We do this in ACR with the *shadows [black]* slider. Moving the slider to the right *expand* the tonal range of the image, moving it to the left, *contract* the tonal range.

It's seldom necessary to use the other ACR controls. As for the controls under the other tabs (*Detail, Lens, Calibrate*), they are beyond the scope of this article. As for the remaining controls under the *Adjust* tab, they more or less duplicate functions that is also available in Photoshop CS. The ACR version is usually gentler then the PS CS versions, but otherwise, they work the same. Below is a short description of each:

The *fill light* and *recovery* sliders let you attempt to recover some details from dense shadows and blown highlights. They work like the Shadow/Highlight filter in Photoshop CS.

The *brightness* slider is a non-linear adjustment that let you redistribute the mid-tone values without moving the black or white point. It duplicates the grey input slider in PS CS Levels dialogue.

The *contrast* slider expands and contracts the tonal range by moving the black and white points, keeping the mid-tone fixed. It works like the contrast slider in Photoshop's Brightness/Contrast dialogue.

The *clarity* slider increases local contrast, with greatest effect on the mid-tones. It works a like a large-radius low-amount unsharp mask.

The *vibrance* slider increases the saturation of all lower-saturated colours with less effect on the higher-saturated colours.

The *saturation* slider can be used to decrease or increase colour saturation. It affects all colours equally from -100 (monochrome) +100 (double the saturation). It works like the saturation slider in Photoshop's Hue/Saturation dialogue.

### **High Dynamic Range**

If we want to record a larger dynamic range than the dynamic range of the camera is capable of recording, there it is not possible to do this with a single exposure.

In such a situation, the solution is to make multiple exposures, each capturing a different dynamic range (e.g. one for the highlights and another for the shadows), and then blend them together.

We do this blending using special software designed for high dynamic range imaging (HDR).

### **Evaluating the Digital Zone System**

If you are used to working with film, the digital Zone System may appear strange at first. But when you get used to it, you'll probably find using digital means to make this type of adjustments much simpler than their chemical predecessors.

With digital tools we get immediate and visual feedback when we apply adjustments, and we can even "undo" adjustments and go back if we don't like the result. Using the ACR for expanding and contracting tonal range is also easier than adjusting development times.

In fact, many digital photographers argue that because digital image editing tools are so powerful and versatile, it is no longer necessary to use the precise spot measurements of the Zone System to determine exposure. They argue that the automatic settings determined by their camera's meter takes well enough care of exposure, and if this causes tones to not fall exactly where they should, this can be "fixed" in processing.

While digital image editing tools are indeed very powerful, you will always get better results if you get it right at the time of exposure, instead of having to "fix" things in processing. If you care about the tones in your prints, the Zone System is just as relevant for digital as it was with film. Digital editing tools just make some adjustments simpler to carry out. Just make sure that you use a calibrated monitor when you adjust the levels, and that you take equal care to make sure that the printer you use to print the result is calibrated, or that you use the correct profile for the ink and paper you use.

## **7. Sources, Resources and Further Reading**

Adams, Ansel (2002): *The Negative* (reprint of final edition from 1981); Little Brown and Company; Bulfinch.

Chavez, Conrad and David Blatner (2009): *Real World Adobe Photoshop CS4*; Peachpit Press.

Clark, R.N. (2005): *Dynamic Range and Transfer Functions of Digital Images and Comparison to Film*.

Clark, R.N. (2006): *Procedures for Evaluating Digital Camera Sensor Noise, Dynamic Range, and Full Well Capacities*.

Fraser, Bruce (2004): *Raw Capture, Linear Gamma, and Exposure*; (pdf) Adobe.

Gardner, Chuck (2006): *Digital and the Zone System*.

ISO (1974): *ISO 2720:1974. General Purpose Photographic Exposure Meters (Photoelectric Type) – Guide to Product Specification*; International Organisation for Standardisation.

ISO (1993): *ISO 6:1993. Photography – Black-and-white pictorial still camera negative film/process systems – Determination of ISO speed*; International Organisation for Standardisation.

ISO (2006): *ISO 12232:2006: Photography – Digital still cameras – Determination of exposure index, ISO speed ratings, standard output sensitivity, and recommended exposure index*; International Organisation for Standardisation.

Johnson, Chris (2006): *The Practical Zone System: For Film and Digital Photography* (4th ed.); Focal Press.

Kodak (2009): *Kodak Image Sensors ISO Measurement*; (pdf) application note MTD/PS-0234, rev. 5, September 28, Eastman Kodak Company.

Koren, Norman (2005a): *A simplified Zone system for making good exposures*.

Koren, Norman (2005b): *Making fine prints in your digital darkroom: Tonal quality and dynamic range in digital cameras*.

Kreunen, Ben (2003): *Dynamic Range*.

Martin, Richard (2006): *What You Need to Know About Dynamic Range*; New York Institute of Photography, Nov. 9.

Reichmann, Michael (2003): *Expose (to the) Right: Maximising S/N Ratio in Digital Photography*; The Luminous Landscape.

Wikipedia (undated): *Light Meter*.

#### Software

This section lists some Zone System software tools that I am aware of. Most of the programs have not been tested by me and I cannot guarantee that they will work as advertised.

Digital Film Tools: *Ozone*; (Photoshop plugin for tone mapping.)

Light Crafts: *LightZone*; (Zone oriented RAW converter and photo editor.)

On-line reviews and user reports about the software listed in this section:

Steinmueller, Uwe (2006): *Light Crafts LightZone™ Diary*; Digital Outback Photo.