

Quantizing Noise in Digital Photography

Douglas A. Kerr

Issue 2
April 28, 2025

FOREWORD

In a digital camera, when we take the voltage output of an individual sensor element and convert it to digital form, a phenomenon called *quantizing* (often called *quantization*) takes place. As a result, the digital representation does not exactly represent the voltage. At the end of the digital image chain, this results in a discrepancy between the reconstructed image and the original image.

In the area of waveform-based digital representation of speech waveforms, we sometimes characterize this discrepancy between the original data and the reconstructed data as a special kind of pseudo-noise, *quantizing noise*.

Some workers suggest that this concept is pertinent to the impact of quantizing error on digital images, and that *quantizing noise* should be reckoned among the ingredients of noise in a digital imaging system, a notion with which the author disagrees. In any case, the process of quantizing does have an effect on how noise already present in the sensor voltage is seen in the digital representation. In this article, both these matters are examined.

An appendix discusses a similar concept followed in the field of digital video engineering.

BACKGROUND

Quantizing

Analog information is characterized by its “continuous” nature. That is, in general, the variable in which we are interested can take on any value within certain limits. The actual temperature of the air outside can be (approximately) 28.17652° C, or 28.17653° C.

When we digitize an analog variable, we take each occurrence of its value and assign it a number taken from a discrete repertoire of values. In effect (in the “classical” arrangement), we “round” the actual value to the nearest value in that repertoire. This process is called *quantizing*. The term comes from the notion that we treat all the values to be digitally represented as if they were made up of an integral number of increments (*quanta*), usually of some fixed size.

“Quantization” and “quantizing”

We often encounter the words *quantization* and *quantizing*, which seem to be almost synonyms. Usage in this regard is not too well regulated. The trail starts with the verb, which is (to) *quantize*. From the standpoint of a linguistic purist, the process would be called *quantizing* (the gerund form of the noun), but *quantization* is often used for the result (and sometimes for the process as well). The related adjective is often seen as either *quantizing* (as in *quantizing error*) or *quantization* (as in *quantization error*).

In the face of all that editorial sophistry and ambiguity, I will use here the term *quantizing* in all situations.

Quantizing error

The result of this process is that the digital representation is never exactly equal to the actual value being represented. The discrepancy is called *quantizing error*.

For example, suppose that our digital structure uses an 8-bit number, interpreted as providing integer values from (in decimal) -128 through +127. If we encode the temperature 28.17652° C in a straightforward way, it will be given the “data number” (DN) 28 (in decimal). In this case, the temperature indicated by that DN will differ from the actual temperature by -0.17652° C. This is the *quantizing error* in this specific case.

If a suite of data numbers represents the “samples” that directly describe, for example, an audio waveform, or a photographic image, then when we reconstruct the waveform or image from that suite of numbers the result will not exactly match the original waveform or image.

In the digital representation of a waveform, we can aptly speak of the end result as *quantizing distortion*. There is no equally tidy description of the overall effect on a photographic image.

The quantizing step size

In a common quantizing situation, the values in the discrete repertoire represent evenly-spaced values of the quantity being represented. In fact the increment of the data number (DN) itself is usually 1, but often there is a scaling factor involved. Thus, in a temperature system, a change in 1 in the DN might represent a change of 0.2° C in the temperature being represented.

The spacing between available values, in the scale of the actual quantity, is called the *quantizing step*. In the example just above, that is (exactly) 0.2° C.

In this article, when we speak in general, abstract terms, we assume that the scale factor involved is such that the quantizing step is 1 “unit”.

The nature of quantizing error

We may ask whether quantizing error is random or not. Well, yes and no.

If the actual value being encoded is known, then the quantizing error is determined. If the value is constant, so will be the quantizing error. Assuming that the quantizing step is 1 (and that the boundaries between steps fall halfway between integers, the “classical rounding” rule), then a source value of 7.028 will always receive a DN of 7, and the quantizing error will always be -0.028. There is nothing random about this.

But in the overall operation of a system, the source values vary (else there would be no need to encode and transmit them—the recipient would already know!). Thus, the quantizing error will typically be different for each value processed.

To analytically evaluate the quantizing error behavior of a system on an “overall” basis, we often assume that the values we encounter will be **uniformly** distributed over a range that is an integral number of quantizing steps in width, or else over a large (but not necessarily integral) number of quantizing steps. That is, we assume it would be equally likely that an individual value would lie in the range 15-15.2 units as in the range 22.3-22.5 units.

How do we describe the magnitude of the quantizing error?

It is often attractive to have a single “metric” that describes the magnitude of the quantizing error phenomenon in a certain system. We might consider using the average quantizing error, but in fact, under our assumptions, that would be zero (since errors would occur that were both positive and negative).

The quantizing error is limited to $\pm b/2$, where b is the quantizing band size. Thus we might want to use the maximum absolute value of the error (which is always $b/2$).

However, especially when we look into the relationship between quantizing error and other discrepancies in the reconstructed data (such as noise), we often draw upon to a common measure from the world of statistics, the *standard deviation* of the error.

The standard deviation of a variable’s value is obtained by taking the various instances of the value (the error quantities, in this case),

squaring them, averaging these squared values, and taking the square root of the average.¹ This process is identical to the determination of the root-mean-square (RMS) value of a varying electrical voltage or current (such as an AC waveform). And in fact there is an electrical rationale for the choice of this particular metric of quantizing error (which we will hear of shortly).

In statistical work, the symbol for standard deviation is σ (lower-case Greek *sigma*).

In line with the principle described earlier, to evaluate a particular scheme, we look at the quantizing error that would occur for a large number of possible variable values (usually evenly distributed over one quantizing band, or over an arbitrary large range), and then determine the σ for the entire set of errors attending those different variable values.²

Note that I have not suggested that this metric for the effect of quantizing error is useful in the area of digital imaging. It is considered useful in a number of other situations

Controlling quantizing error

We haven't yet discussed the actual adverse impact of quantizing error, but it seems obvious at this point that it is not a good thing, so having less would be better.

We mentioned above that the maximum magnitude of the quantizing error that can occur is $b/2$, where b is the size of the quantizing step. Clearly, if we reduce b , then the maximum quantizing error decreases, and not surprisingly the measure we mentioned of overall quantizing error (σ) decreases as well.

Given that we must, in any particular system, accommodate a certain range of the variable of interest, decreasing the size of the quantizing step means that we must have more steps overall (more possible DN values). In a binary context, this means an increase in the number of the bits used to represent the DN. Essentially, each bit we add cuts

¹ To be precise, we take the differences of the individual values **from the average** of the suite, square those differences, and so forth. Here, where we often assume the average quantizing error to be zero, we can use the simple description above.

² Note that, even though we assume each of those different variable values is equally-probable, the overall σ is not the average of the σ s for the different values. Rather, it is the square root of the sum of the squares of the individual σ s. In other words, we just extend the averaging of the squares of the individual errors, a step in the determination of σ , over the entire suite of test data.

the size of the quantizing step in half, and the measure of the quantizing error in half.

Note that, in general, we cannot do this “cost free”, since there is a certain cost associated with storing or transmitting a bit.

IMPACT OF QUANTIZING ERROR IN DIGITAL AUDIO

Principle of digital representation of an audio waveform

When we represent an audio waveform in digital form (in the most straightforward way), we first “sample” the waveform. That is, we capture the instantaneous value of the signal voltage at repetitive intervals. If the rate at which we take these samples is greater (even by a little) than twice the highest frequency contained in the waveform, then, according to the Nyquist-Shannon sampling theorem, from the suite of samples values alone we can exactly reconstruct the entire original waveform. Note that this doesn’t say “a good approximation of the original waveform”—it says “the original waveform” (implying a precise reconstruction).

However, perfectly achieving this promise requires that when we capture the instantaneous voltage, we do so “exactly”, to unlimited precision and with zero error. Of course, this is impossible in practice. But to the degree we approach that, we can actually approach the promise of the theorem.

We then take each sample voltage and digitize it—we assign it a data number (DN). Of course, quantizing is involved here, and thus quantizing error emerges.

When we want to reconstruct the original waveform, we reconstruct each sample voltage based on the DN that describes it, and then through a filter process, from that recreated train of samples we reconstruct the waveform.

But the DNs do not exactly represent the samples (owing to quantizing error), and as a consequence, the recreated samples are not identical to the original sample. Thus the waveform reconstructed from them is not identical to the original waveform.

This impairment is called *quantizing distortion*.

The practical perceptual impact of quantizing distortion is that, if we have enough of it, the recovered audio “sounds funny” (“buzzy”, in severe cases).

“Quantizing noise”

As we make detailed plans for a system for digital representation of audio waveforms, or test the behavior of an actual system, we need some *metric*, a “score” by which we can describe numerically “how bad is the quantizing distortion”.

In a telecommunication system, we can artificially choose to think of “noise” as being *anything that arrives in the delivered waveform that was not present in the original waveform*. Having done so, we can then think of the discrepancy between the delivered waveform and the original waveform, due to quantizing distortion, as a special kind of noise (perhaps a “pseudo-noise”).

The object of this outlook is that we can judge the extent of quantizing distortion using the same definition (and even the same measuring instruments) we use to judge the extent of the more familiar kinds of noise.

In particular, this involves, in principle:

- isolating the noise component of the delivered waveform (perhaps by “subtracting out” a replica of the original waveform)
- dissecting the noise component into sub-components, each spanning only a narrow range of audible frequencies
- multiplying the power³ in each such sub-component by a factor that represents the relative sensitivity of the human ear to sound at that frequency (called “weighting” the noise by frequency)
- adding all the results for all the sub-components

The sum (in units of power) is used as the metric describing the amount of noise.

The RMS measure

The power represented by an electrical waveform (under some assumption about circuit impedance) is proportional to the square of the voltage. But, in the case of a constantly-varying voltage (that is, a waveform), what voltage is that? It is the RMS measure of the voltage (which is its standard deviation). (The voltage of your household electrical supply is stated in terms of its RMS value.)

³ In general, the human response to an acoustical waveform is related to the power contained in it. Thus, when we are working with electrical audio waveforms or their components, we are interested in the power that they represent.

Now the loop has been closed—we see why, at least in the digital audio context, the assessment of quantizing error in terms of its *standard deviation* makes perfect sense: the standard deviation is a direct indicator of the *power* contained in the “quantizing noise”, which in turn is an indication of its impact on the listener.

In the more general case of quantizing error, where the values involved are not voltages along a waveform but perhaps thermometer readings, there is a statistical justification for the use of the standard deviation as the “metric” of quantizing error (in fact, as a metric for error of any type). I will not discuss that here. And in fact, its validity is not nearly so obvious in the general case as in the electrical waveform case; there is no physical property, equivalent to “power”, that comes from the square of the standard deviation of a suite of temperature readings.

Is it really “noise”?

Returning now to the concept of “quantizing noise”, our adoption of this notion does not mean that quantizing distortion in any way seems to the listener like other familiar kinds of noise. But practical tests show that this metric does correlate fairly well with the listener’s reaction to the impairment resulting from the quantizing distortion. In any case the adoption of the approach was heavily influenced by the pragmatic fact, when digital transmission of speech was introduced into the telephone network, apparatus for physically performing the steps listed above was widely available (it having been used for many years to measure noise of the conventional type in analog transmission channels). When in doubt, choose a metric that we can measure—especially with instruments we already have!

The nature of the perceived impairment

Why does an audio waveform afflicted by quantizing distortion sound “funny” to the listener? Well, if we consider the Fourier transforms of the original waveform and of the “distorted” waveform, we will find a difference in the frequency composition, and it is basically through the frequency composition that the human ear recognizes different waveforms. (The cochlea of the human ear essentially performs Fourier analysis on the “waveform” of the motion of the eardrum, transmitting to the brain its frequency distribution.)

The human perceptual system is “familiar” with a wide range of sounds, and does not expect to hear coming out of the mouth of a human a waveform whose frequency composition is different from any familiar model.

In popular music (especially guitar music), we often intentionally generate “distorted” waveforms that are not found in nature. It is

unlikely that such a waveform, if afflicted in digital transmission by substantial quantizing distortion, would be recognized by even a highly trained listener as “distorted” from the waveform as transmitted.

A complication in digital audio is that generally the quantizing interval is not constant. Rather, it uniformly increases as we go farther along the range of the variable. This arrangement gives better overall perceived performance for any number of bits in the encoding. It is a little like a constant percentage accuracy in a measuring instrument, in which the maximum error is a certain fraction of the voltage being measured.

In particular, it means that the “amount” of quantizing “noise” (its standard deviation, or RMS measure) varies with the signal amplitude, leading to sort of a uniform “signal-to-(quantizing) noise ratio”.

As a result, to get a single value of our “metric” that describes the quantizing distortion of the system, we need to not just test for different instantaneous values of the signal over some range, but must test with a sine-wave signal whose amplitude (full range) corresponds to the standard test level signal.

Only through all this subtlety does the metaphor of “quantizing noise” as a model for quantifying quantizing distortion in digital audio waveforms take life.

Summary

Before we move to the real topic of this article, let’s circle back to the beginning of this last section. We need to realize that characterizing quantizing distortion as “noise” is only a conceit that lets us quantify the phenomenon using the same metrical concept, and physically measure it with the same instruments, as for actual noise. Hold that thought.

QUANTIZING ERROR IN DIGITAL PHOTOGRAPHY

Introduction

Now that we have some background, partly historical, in hand from another field, let’s look at the implications of quantizing error in digital photography.

Pertinent principles of digital imaging

In this situation, the “continuous variable” to be encoded describes the color of the scene (*color* including both the attributes of *luminance* and *chromaticity*). This variable is two-dimensional in its *domain* (since the scene, or its image, is two dimensional, in the familiar, geometric

sense), and is three-dimensional in its *range* (in the mathematical sense, since color actually requires three variables to describe it).

To simplify our discussion, we will (unless stated otherwise) think of this situation:

- The domain of our variable is actually just a one dimensional track across the image (like one scan line in a video camera).
- The system is monochrome, so we can replace *color* with just *luminance* (which is only a single-dimensional quantity, involving only a single variable).

Here again, the variable of interest (the luminance of the scene along a track) is sampled by the use of an array of discrete sensor elements (*sensels*). They can be thought of as each capturing the value of the variable at a particular spot in the image on the sensor array.⁴

The “determination” reported by each sensel is then turned into a data number (DN), which describes the value of our luminance variable at that point in the image.

Quantizing and quantizing error

Of course, quantizing is involved here. If the luminance at a certain sensel location (sampling point) is 156.23 units, it will receive a data number of 156 (decimal). Insofar as the data number representing the value of the variable of that point, it has an error (quantizing error) of -0.23 units.

These data numbers go through a very long process before we are able to again reconstruct the image, but for now let’s just suppose that we take the suite of data numbers representing the luminance of the scene along a track and from them immediately reconstruct the luminance variable. (And we will actually do this for every “track” across the image, so a complete two-dimensional image is reconstructed.)

This reconstructed image will vary from the actual image (and thus from the scene) scene because of *quantizing error*.

⁴ In reality, they pick up sort of an average luminance over a small area centered on a certain point, a distinction that has important technical consequences but which we can ignore for the moment.

Difference from the audio waveform context

Here we begin to understand the difference between the audio and photographic contexts. How do we generally become aware of the effect of quantizing error in this case? Not because it causes a more-or-less random discrepancy between the reconstructed luminance of every point in the image and the actual luminance of the scene point (although it does). In general, we would not be able to discern that as “anomalous” (the way a listener to a reconstructed audio signal affected by quantizing distortion would be able to, because of familiarity with typical speech waveforms).

There are two principal ways in which the impact of quantizing error are “recognizable” to the viewer of the image.

One is the effect on areas of the scene that have a fairly uniform and gradually-varying luminance. Since the digital representation can only convey luminances that differ by integral multiples of one digital unit, we see “bands” across such an area, each band having a constant luminance across it, with the luminance changing slightly, but probably visibly, from band to band (the difference corresponding to a change of one unit in the digital number).

Even when this banding is not prominent, it can negatively impact the image. But this phenomenon is greatly distinct from what we normally think of as “noise” in an image: a random discrepancy in the reconstructed luminance that appears in an area of constant or slowly varying luminance as a “mottled” or “granular” effect.

The second situation is on boundaries between well-defined scene areas (the edge of a door frame, for example). These typically involve a “luminance slope” *across* the boundary. Quantizing error can cause “banding” across that slope, which of itself may hardly be visible. But if the luminance is also slowly changing *along* the boundary, the result may be a subtle “zig-zag” effect on the perceived boundary (spoken of in video engineering as a “contour effect”).

What about other scenes, with mostly “finer detail” (my favorite example is a shot of the surface of a gravel road)? The quantizing distortion makes the reconstructed image different from the actual image on the sensor array, but not in a way that the viewer can usually notice.

A metric for quantizing error in this context

Now, would the standard deviation of the quantizing error give us a useful “score” for the particular quantizing scheme used in our camera? It doesn’t.

In fact, probably the most important single “metric” for quantizing behavior in the context of digital imaging is just the size of the quantizing step (perhaps expressed as a fraction of the overall luminance range).

And that having been said, it seems that the characterization of quantizing error as “noise” here will not, as it did in the digital audio context, lead us to a useful metric of quantizing error.

A conclusion

Accordingly, I discourage workers in this field from thinking in terms of “quantizing noise” when discussing the impact of quantizing error on the “faithful and accurate” reproduction of the original photographic image, or for including an item for “quantizing noise” in a workup of various ingredients of overall noise in a digital imaging system..

QUANTIZING AND “REAL” NOISE

Introduction

Let’s next accept the generally-recognized concept of noise in digital imaging: a random phenomenon that manifests as a mottled or granular appearance in an image area of uniform or slowly-varying luminance. Doesn’t quantizing error contribute to this?

No, it doesn’t. But it does have an influence on it.

Noise sources

In our digital camera image chain, there are various phenomena that result in the electrical signal from a sensel having a random departure from the voltage that would properly reflect the photometric exposure on the sensel. They include:

- Shot noise. This comes from the statistical nature of the arrival of photons.
- Thermal noise. This comes from random electron activity in the sensel detector and associated circuitry owing to their not being at a temperature of absolute zero.
- Reset noise. Each sensel detector has to be reset at the beginning of every exposure (initially charged, not discharged, as is commonly believed; the incidence of each photon discharges the detector by a certain increment). This may not occur consistently, for various reasons. Since the final charge of the detector is compared to the assumed initial charge to discern the detector’s

“report”, this phenomenon causes a random discrepancy in these reports.

These are what produce the “noise” we see in the image.

The statistical nature of the noise

Let’s take a moment to examine the nature of these random variations. They are generally considered to be (with a certain limitation) “Gaussian” in their nature. That means that:

- the probability that the voltage of the noise component for any given “sample” falls in a certain voltage range

follows essentially the “normal” statistical distribution. The probability density curve that presents this distribution is the famous “bell curve”.

Under this distribution, the probability that the noise voltage for a particular “sample” lies in a certain narrow band about zero has a certain value. The probability that it lies in an adjacent band of the same width is less, and in bands of the same width further and further from zero is successively less.⁵

In a true standard distribution, there would be no limit as to how large (positive or negative) the error voltage could be. There might be one chance in a billion that the voltage would lie in the range from +1000.0 to +1000.1 volts! Fortunately, physical and electrical realities mean that this would never actually happen. In a sense, the distribution is “truncated” by these realities.

Since the range of the standard distribution is (conceptually) unlimited (every bell curve has infinite width, even though we can’t draw “all of it”), how can we describe the degree of variation of a variable having a standard distribution (or approximately so)? The statistical measure we use is the *standard deviation* of the distribution, σ , which was discussed earlier in connection with evaluating quantizing “noise” in an audio transmission system. A distribution curve with a greater value of σ is “fatter”, and implies greater variation in the values.

⁵ Why do we not say “the probability that the voltage is a certain value”, but instead speak of it as “within a certain small range”? This is because the probability of a voltage having a certain value is zero, just as is the probability that a machined roller will have a diameter of exactly 2 inches. But we can speak of the probability that the diameter is between 1.99999999 inches and 2.00000001 inches.

One way to interpret the standard deviation of a “normal” distribution is that about 68.2% of the occurrences fall within \pm one standard deviation from zero (the range from -1σ through $+1\sigma$).⁶

Where we observe it

Of course, we do not see the effect of the various noise components while they are voltages (components of the overall sensel “output voltage”). We see their effect in the digital representation of the sensel voltage.

Now we know that quantizing has an effect on the “signal” component of that voltage—the part that actually, ideally represents the photometric exposure on the sensel. We have spoken of the “banding” that is the most common perceptible result of this.

What effect does it have on the random component of the sensel voltage—the “noise”? It is not magically exempt from quantizing error.

Let’s do a little thought game to get a hint at how this plays. Assume that the quantizing “step” in our system is 1 mV, and that quantizing follows the “classical rounding” rule (quantizing to the nearest allowable value). Imagine that, for all the sensels in a block receiving a uniform photometric exposure, the part of the sensel output voltage that reflects that photometric exposure (the “signal” component of the voltage) is 12.00 mV. Assume for simplicity’s sake that the noise has a “truncated uniform” distribution (not Gaussian), and so the noise voltage is strictly confined within a certain range, which in this case is from -0.04 mV to +0.04 mV, a “range” of 0.08 mV.

Thus, over our block, the sensel voltages will vary, randomly, over the range 11.96 mV through 12.04 mV.

But these voltages will all be given the digital number (DN) “12”. In an image reconstructed from those digital numbers, all the pixels will be identical. We will see no noise in numerical analysis of the DNs, and we will see no noise on the reconstructed image. The (admittedly small) noise in the sensel outputs has been discarded by the quantizing process. (How lovely!)

But before we get too excited, let’s consider another (and equally likely) case. Here, the actual signal component of the voltage of all the sensels is 12.50 mV. The noise is the same (a range of ± 0.04 mV). The actual sensel voltages will vary randomly from 12.46 mV through

⁶ In the general case, we would say “within ± 1 standard deviation from the average of the values”; since here we are dealing with error values, in a symmetrical system, the average is zero, and so the simpler interpretation (“about zero”) is meaningful.

12.64 mV. Since 12.5 mV is the boundary between two quantizing levels, about half the sensel voltages will be encoded with the DN 12, and about half with the DN 13. The average absolute error (that is, regardless of sign) will be about 0.5 mV (quite substantial compared to the range of the noise).

So how can we generalize what quantizing does to the actual sensel noise on its way to the digital world? Probably the best thing to do is to run a large suite of tests with the “signal” component of the “noisy” sensel voltages varying in small steps from perhaps 12.00 mV to 13.00 mV. For each trial, we will find a different amount of “noise” in the digital representation (we saw just recently the best and worst cases). Then perhaps an average of all these will tell us what quantizing does to the noise on its way to the digital world.

If we presume some specific distribution of noise voltages (such as the Gaussian distribution ordinarily assumed, with reservations), “we” can analytically predict that result. But “I” am not ready to take that on this week.

Rather, I set up a simulation (using an Excel spreadsheet) to give some insight into the result. The model uses a Gaussian distribution for the noise, and allows us to choose its standard deviation (how “great” was the noise). The model also allows us to use different sized quantizing intervals.

The model “ran trials” with the “signal” voltage at ten different values over an entire quantizing step (the phenomenon repeats at that interval). For each “trial”, the scheme determined the standard deviation of the data numbers of the digitized sensel voltages. Those results were combined for all ten trials to get an overall assessment of:

- the noise as seen in the digital representation
compared to
- the noise in the original sensel voltages.

The results are very interesting. Not surprisingly, they differed significantly depending on the size of the noise (as compared to the quantizing step). In each case, when we speak of the amount of noise, it is in terms of standard deviation. Because of practical shortcomings in the model (it used a discrete table of the standard deviation with only 100 steps, for example), these results are approximate, but clearly reveal the general pattern.

- For noise at the sensel output whose standard deviation was $1/2$ the quantizing step, the noise in the digital output was about 1.25 times the noise at the sensel output.

- For sensel noise whose standard deviation equal to the quantizing step, the digital noise was 1.1 times as big as the noise at the sensel output.
- For sensel noise whose standard deviation was 2 times the quantizing step, the digital noise was essentially the same size as the noise at the sensel output.
- For sensel noise whose standard deviation was 5 times the quantizing step, the digital noise was only 0.9 times as big as the noise at the sensel output.
- For sensel noise whose standard deviation was 10 times the quantizing step, the digital noise was only about 0.85 times as big as the noise at the sensel output.

CONCLUSIONS

Quantizing noise in digital imaging

- It is not meaningful, nor even pragmatically useful, to think of the impact of quantizing error in the reconstructed image as a form of noise ("quantizing noise")
- It is not meaningful to include "quantizing noise" in the "roundup" of noise ingredients in a digital imaging system

The effect of quantizing error on noise

- Quantizing error affects the transport of noise from the input to the digitization process to its digital output.
- The degree of noise perceived in the digital representation can be slightly larger, or slightly smaller, than the noise actually existing at the input to the digitizing process as a result of quantizing

ACKNOWLEDGEMENT

Thanks to Carla Kerr for her insightful copy editing of this difficult manuscript.

#

For more articles by this author, visit The Pumpkin at:
<http://doug.kerr.home.att.net/pumpkin>

APPENDIX A

“Quantizing noise” in digital video engineering

In digital video engineering, there is a fairly-widely recognized practice of characterizing the impact of quantizing error under a metaphor of “quantizing noise”, in particular the use of the metric *signal to quantizing noise ratio*. We must note that here, the specific metric for *signal* is the peak-to-peak value of the signal, and the specific metric for *noise* is the RMS value (standard deviation) of the quantizing error.

I noted in the body of this article that, in a digital photography context, the amount of quantizing error depends on the actual value of the “signal” for any given pixel in a particular exposure. Thus, to get a single figure characterizing the degree of quantizing error, we must (actually or on paper) expose the system to a range of signal values, and determine the (RMS) value of the quantizing error over that entire exercise.

The very same situation, hardly surprisingly, is encountered in digital video. Again, there we expose the system to a variation in the basic signal value (representing, again, luminance) extending across the entire “tonal scale” of the system (from black to white), with a “uniform distribution”. (That is, the signal will be in any given width small region of the range for the same fraction of the time, regardless of where the range is; it will be in the range 0.15-0.16 V the same fraction of the time as in the range 0.53-0.54 V.)

We determine the RMS value of the suite of individual quantizing errors, and use this as the denominator of our *signal-to-quantizing-noise ratio*. The numerator (signal) is just the peak to peak value of the largest legitimate signal (recognizing that, as we might expect, there is a standardized “headroom” built into the scaling of the analog to digital conversion.)

We express this ratio in decibels (dB), as is the custom for (real) “signal to noise ratios”.

If the analog to digital conversion proceeds “ideally”, the result follows directly from the “bit depth” of the digital encoding (with no allowance for “headroom”:

$$R = 6.02n + 10.8 \quad (1)$$

where R is the signal to quantizing noise ratio and n is the number of bits in the digital representation.

(Those who have done theoretical derivations in the area of quantizing error in telecommunication systems may guess, correctly, that the constant “10.8” grows out of the appearance of the factor $\sqrt{12}$ at one stage of the process. It in fact is $20\log\sqrt{12}$).

And in fact the actual measured values for typical systems fall very close to this.

Now, back to the theme of this paper. Is this quantizing error really *noise*? Let’s see if it “looks like a duck” or “quacks like a duck”.

What are the perceptual manifestations of quantizing error in a digital image? Principally:

- Banding on areas of slowly-changing luminance (discussed in the body of the article).
- Contour effects (discussed in the body of the article).

The characterization of these, perceptually, as “noise” is questionable (especially for the former).

Now, let’s examine the role of the metric, *signal-to-quantizing-noise ratio*.

Does it tell us the amount of quantizing error. Yes, in this sense: if we have twice the quantizing error (owing to, say, a one bit smaller encoding), we have a 6.02 dB less value of the *signal-to-quantizing-noise ratio*. (For comparison, if we have twice the real noise, then the conventional *signal-to-noise ratio* is 6.02 dB less.)

But what about the absolute scale? Is the visual impairment from quantizing error, when the *signal-to-quantizing-noise ratio* is 46 dB (which would occur in a 6-bit system), perceptually comparable to the visual impairment caused by conventional noise giving a *signal-to-noise ratio* of 46 dB? No. In fact, the nature of their impact is so different that it is very difficult to get observers to even make a comparison. (“Which of these images do you like the best . . .”)

And probably, if you ask a video engineer familiar with this matter to answer, “how bad is a *signal-to-quantizing-noise ratio* of 46 dB, anyway”, she well might say, “well, you know—like what we get in that 6-bit system for our security cameras”.

Thankfully, nobody (well, almost) takes the use of the term “noise” in this definition to suggest that quantizing noise can be added to other types of noise (thermal noise in the analog amplifier chain, for example) to get an overall noise value, from which an overall *signal-to-noise ratio* can be reckoned.

Of course, intellectually, it would be better had the video engineering word decided to call their metric “signal to quantizing error ratio”. But they didn’t. And we can imagine why. Many workers in that field were aware of the clever usage of “quantizing noise” in the field of audio waveform encoding, and were anxious to carry that knowledge forward when digital video technique later emerged.

Why the attraction for the “dB” form? Video engineers like to think of every ratio in an “amplitude” domain in terms of dB. But, by rights, this is only applicable for power, or for quantities having a direct relationship to power (including when their squares correspond to power. The peak-to-peak range of a video waveform has no mathematical relationship to power.

Nevertheless, the use that is made of the concept here—to give an arbitrary metric for the “relative” degree of quantizing error—is fairly benign. And of course, the practice is so entrenched that there would be little point in my railing out against it (especially in an arena where I have no credentials).

#