Extensions to the PNG 1.1 Specification:
Special-Purpose Public Chunks
Version 1.1.0
October 26, 1998

The latest versions of this document and related information can always be found at the PNG FTP archive site, <ftp://ftp.uu.net/graphics/png/>. The maintainers of the PNG specification can be contacted by e-mail at png-info@uunet.uu.net.

Abstract

This document is an extension to the Portable Network Graphics (PNG) specification, version 1.1. It describes additional public chunk types and contains additional information for use in PNG images.

This document, together with the PNG specification, contains the entire list of registered "public" PNG chunks. The additional registered chunks appearing in this document are the oFFs, pCAL, sCAL, gIFg, gIFs, and fRAc chunks, plus the deprecated gIFt chunk.

Chunks described here are expected to be less widely supported than those defined in the basic specification. However, application authors are encouraged to use these chunk types whenever appropriate for their applications.

This document also describes data representations that do not occur in the core PNG format, but are used in one or more special-purpose chunks. New chunks should use these representations whenever applicable, in order to maximize portability and simplify decoders.

Contents

1 Introduction **3**

2 Data Representation **3**
2.1 Integer values . . . 3
2.2 Floating-point values . . . 3

3 Summary of Special-Purpose Chunks **4**

4 Chunk Descriptions **4**
4.1 oFFs Image offset . . . 4
4.2 pCAL Calibration of pixel values . . . 5
4.3 sCAL Physical scale of image subject . . . 8
4.4 gIFg GIF Graphic Control Extension . . . 9
4.5 gIFx GIF Application Extension . . . 9

5 Chunks Not Described Here **10**
5.1 fRAc Fractal image parameters . . . 10

6 Text Chunk Keywords **10**

7 Deprecated Chunks **10**
7.1 gIFt GIF Plain Text Extension . . . 10

8 Security Considerations **11**

9 Appendix: Sample code **11**
9.1 pCAL . . . 12
9.2 Fixed-point gamma correction . . . 15

10 Appendix: Rationale **19**
10.1 pCAL . . . 19

11 Appendix: Revision History **20**

12 References **20**

13 Credits **21**

1 Introduction

This document is an extension to the Portable Network Graphics (PNG) specification, version 1.1. It describes additional public chunk types and contains additional information for use in PNG images. The basic PNG specification is available from the PNG FTP archive mentioned above.

Chunks described here are expected to be less widely supported than those defined in the basic specification. However, application authors are encouraged to use these chunk types whenever appropriate for their applications. Additional chunk types may be proposed for inclusion in this list by contacting the PNG specification maintainers at png-info@uunet.uu.net.

This document also describes data representations that do not occur in the core PNG format, but are used in one or more special-purpose chunks. New chunks should use these representations whenever applicable, in order to maximize portability and simplify decoders.

2 Data Representation

2.1 Integer values

Refer to Section 2.1 of the PNG specification for the format and range of integer values.

2.2 Floating-point values

The core of PNG does not use floating-point numbers anywhere; it uses integers or, where applicable, fixed-point fractional values. However, special-purpose chunks may need to represent values that do not fit comfortably in fixed-point notation. The textual floating-point notation defined here is recommended for use in all such cases. This representation is simple, has no a priori limits on range or precision, and is portable across all machines.

A floating-point value in this notation is represented by an ASCII text string in a standardized decimal floating-point format. The string is variable-length and must be terminated by a null (zero) character unless it is the last item in its chunk. The string consists of an optional sign ("+" or "-"), an integer part, a fraction part beginning with a decimal point ("."), and an exponent part beginning with an "E" or "e" and optional sign. The integer, fraction, and exponent parts each contain one or more digits (ASCII "0" to "9"). Either the integer part or the fraction part, but not both, may be omitted. A decimal point is allowed, but not required, if there is no fraction part. The exponent part may be omitted. No spaces or any other character besides those specified may appear.

Note in particular that C-language "F" and "L" suffixes are not allowed, the string "." is not allowed as a shorthand for 0 as in some other programming languages, and no commas or underscores are allowed. This format ought to be easily readable in all programming environments.

3 Summary of Special-Purpose Chunks

This table summarizes some properties of the chunks described in this document.

Name	Multiple OK?	Ordering constraints
oFFs	No	Before IDAT
pCAL	No	Before IDAT
sCAL	No	Before IDAT
gIFg	Yes	None
gIFt	Yes	None (this chunk is deprecated)
gIFx	Yes	None
fRAc	Yes	None

4 Chunk Descriptions

4.1 oFFs Image offset

The oFFs chunk gives the position on a printed page at which the image should be output when printed alone. It can also be used to define the image's location with respect to a larger screen or other application-specific coordinate system.

The oFFs chunk contains:

```
X position:     4 bytes (signed integer)
Y position:     4 bytes (signed integer)
Unit specifier: 1 byte
```

Both position values are signed. The following values are legal for the unit specifier:

```
0: unit is the pixel (true dimensions unspecified)
1: unit is the micrometer
```

Conversion note: one inch is equal to exactly 25,400 micrometers. A micrometer (also called a micron) is 1/1,000,000th of a meter.

The X position is measured rightwards from the left edge of the page to the left edge of the image; the Y position is measured downwards from the top edge of the page to the top edge of the image. Note that negative values are permitted, and denote displacement in the opposite directions. Although oFFs can specify an image placement that is partially or wholly outside the page boundaries, the result of such placement is application-dependent.

If present, this chunk must precede the first IDAT chunk.

4.2 pCAL Calibration of pixel values

When a PNG file is being used to store physical data other than color values, such as a two-dimensional temperature field, the pCAL chunk can be used to record the relationship (mapping) between stored pixel samples, original samples, and actual physical values. The pCAL data might be used to construct a reference color bar beside the image, or to extract the original physical data values from the file. It is not expected to affect the way the pixels are displayed. Another method should be used if the encoder wants the decoder to modify the sample values for display purposes.

The pCAL chunk contains:

```
Calibration name:      1-79 bytes (character string)
Null separator:        1 byte
Original zero (x0):    4 bytes (signed integer)
Original max  (x1):    4 bytes (signed integer)
Equation type:         1 byte
Number of parameters: 1 byte
Unit name:             0 or more bytes (character string)
Null separator:        1 byte
Parameter 0 (p0):      1 or more bytes (ASCII floating-point)
Null separator:        1 byte
Parameter 1 (p1):      1 or more bytes (ASCII floating-point)
... etc ...
```

There is no null separator after the final parameter (or after the unit name, if there are zero parameters). The number of parameters field must agree with the actual number of parameters present in the chunk, and must be correct for the specified equation type (see below).

The calibration name can be any convenient name for referring to the mapping, and is subject to the same restrictions as a tEXt keyword: it must contain only printable Latin-1 [ISO/IEC-8859-1] characters (33-126 and 161-255) and spaces (32), but no leading, trailing, or consecutive spaces. The calibration name can permit applications or people to choose the appropriate pCAL chunk when more than one is present (this could occur in a multiple-image file, but not in a PNG file). For example, a calibration name of “SI” or “English” could be used to identify the system of units in the pCAL chunk as well as in other chunk types, to permit a decoder to select an appropriate set of chunks based on their names.

The pCAL chunk defines two mappings:

- A mapping from the stored samples, which are unsigned integers in the range 0..max, where max=(2^bitdepth)-1, to the original samples, which are signed integers. The x0 and x1 fields, together with the bit depth for the image, define this mapping.
- A mapping from the original samples to the physical values, which are usually real numbers with units. This mapping is defined by x0, x1, the equation type, parameters, and unit name.

The mapping between the stored samples and the original samples is given by the following equations:

```
original_sample =
   (stored_sample * (x1-x0) + max/2) / max + x0

stored_sample =
    ((original_sample - x0) * max + (x1-x0)/2) / (x1-x0)
    clipped to the range 0..max
```

In these equations, "/" means integer division that rounds toward negative infinity, so n/d = integer(floor(real(a)/real(b)))). Note that this is the same as the "/" operator in the C programming language when n and d are nonnegative, but not necessarily when n or d is negative.

Notice that x0 and x1 are the original samples that correspond to the stored samples 0 and max, respectively. Encoders will usually set x0=0 and x1=max to indicate that the stored samples are equal to the original samples. Note that x0 is not constrained to be less than x1, and neither is constrained to be positive, but they must be different from each other.

This mapping is lossless and reversible when abs(x1-x0) <= max and the original sample is in the range x0..x1. If abs(x1-x0) > max then there can be no lossless reversible mapping, but the functions provide the best integer approximations to floating-point affine transformations.

The mapping between the original samples and the physical values is given by one of several equations, depending on the equation type, which may have the following values:

```
0: Linear mapping
1: Base-e exponential mapping
2: Arbitrary-base exponential mapping
3: Hypberbolic mapping
```

For equation type 0:

```
physical_value = p0 + p1 * original_sample / (x1-x0)
```

For equation type 1:

```
physical_value =
   p0 + p1 * exp(p2 * original_sample / (x1-x0))
```

For equation type 2:

```
physical_value =
   p0 + p1 * pow(p2, (original_sample / (x1-x0)))
```

For equation type 3:

```
physical_value =
   p0 + p1 * sinh(p2 * (original_sample - p3) / (x1-x0))
```

For these physical value equations, "/" means floating-point division.

The function exp(x) is e raised to the power of x, where e is the base of the natural logarithms, approximately 2.71828182846. The exponential function exp() is the inverse the natural logarithm function ln().

The function pow(x,y) is x raised to the power of y.

```
pow(x,y) = exp(y * ln(x))
```

The function sinh(x) is the hyperbolic sine of x.

```
sinh(x) = 0.5 * (exp(x) - exp(-x))
```

The units for the physical values are given by the unit name, which may contain any number of printable Latin-1 characters, with no limitation on the number and position of blanks. For example, "K", "population density", "MPa". A zero-length string can be used for dimensionless data.

For color types 0 (Gray) and 4 (Gray-alpha), the mappings apply to the gray sample values (but not to th3e alpha sample). For color types 2 (RGB), 3 (indexed RGB), and 6 (RGBA), the mappings apply independently to each of the red, green, and blue sample values (but not the alpha sample). In the case of color type 3 (indexed RGB), the mapping refers to the RGB samples and not to the index values.

Linear data can be expressed with equation type 0.

Pure logarithmic data can be expressed with either equation type 1 or 2:

```
Equation type 1        Equation type 2
x0 = 0                 x0 = 0
x1 = max               x1 = max
p0 = 0                 p0 = 0
p1 = bottom            p1 = bottom
p2 = ln(top/bottom)    p2 = top/bottom
```

Equation types 1 and 2 are functionally equivalent; both are defined because authors may find one or the other more convenient.

Using equation type 3, floating-point data can be reduced (with loss) to a set of integer samples such that the resolution of the stored data is roughly proportional to its magnitude. For example, floating-point data ranging from -10^31 to 10^31 (the usual range of 32-bit floating-point numbers) can be represented with:

```
Equation type 3
x0 = 0
x1 = 65535
p0 = 0.0
p1 = 1.0e-30
p2 = 280.0
p3 = 32767.0
```

The resolution near zero is about 10^-33, while the resolution near 10^31 or -10^31 is about 10^28. Everywhere the resolution is about 0.4 percent of the magnitude.

Note that those floating-point paramaters could be stored in the chunk more compactly as follows:

```
p0 = 0
p1 = 1e-30
p2 = 280
p3 = 32767
```

Applications should use double precision arithmetic (or take other precautions) while performing the mappings for equation types 1, 2, and 3, to prevent overflow of intermediate results when p1 is small and the exp(), pow(), or sinh() function is large.

If present, the pCAL chunk must appear before the first IDAT chunk. Only one instance of the pCAL chunk is permitted in a PNG datastream.

4.3 sCAL Physical scale of image subject

While the pHYs chunk is used to record the physical size of the image itself as it was scanned or as it should be printed, certain images (such as maps, photomicrographs, astronomical surveys, floor plans, and others) may benefit from knowing the actual physical dimensions of the image's subject for remote measurement and other purposes. The sCAL chunk serves this need. It contains:

```
Unit specifier: 1 byte
Pixel width:    1 or more bytes (ASCII floating-point)
Null separator: 1 byte
Pixel height:   1 or more bytes (ASCII floating-point)
```

The following values are legal for the unit specifier:

```
1: unit is the meter
2: unit is the radian
```

Following the unit specifier are two ASCII strings. The first string defines the physical width represented by one image pixel; the second string defines the physical height represented by one pixel. The two strings are separated by a zero byte (null character). As in the tEXt chunk, there is no trailing zero byte for the final string. Each of these strings contains a floating-point constant in the format specified above (Floating-point values, Section 2.2). Both values are required to be greater than zero.

If present, this chunk must precede the first IDAT chunk.

4.4 gIFg GIF Graphic Control Extension

The gIFg chunk is provided for backwards compatibility with the GIF89a Graphic Control Extension. It contains:

```
Disposal Method: 1 byte
User Input Flag: 1 byte
Delay Time:      2 bytes (byte order converted from GIF)
```

The Disposal Method indicates the way in which the graphic is to be treated after being displayed. The User Input Flag indicates whether user input is required before continuing. The Delay Time specifies the number of hundredths (1/100) of a second to delay before continuing with the processing of the datastream. Note that this field is to be byte-order-converted.

The "Transparent Color Flag" and "Transparent Color Index" fields found in the GIF89a Graphic Control Extension are omitted from gIFg. These fields should be converted using the transparency features of basic PNG.

The GIF specification allows at most one Graphic Control Extension to preceed each graphic rendering block. Because each PNG file holds only one image, it is expected that gIFg will appear at most once, before IDAT, but there is no strict requirement.

4.5 gIFx GIF Application Extension

The gIFx chunk is provided for backwards compatibility with the GIF89a Application Extension. The Application Extension contains application-specific information. This chunk contains:

```
Application Identifier: 8 bytes
Authentication Code:    3 bytes
Application Data:       n bytes
```

The Application Identifier is a sequence of eight printable ASCII characters used to identify the application creating the Application Extension. The Authentication Code is three additional bytes that the application may use to further validate the Application Extension. The remainder of the chunk is application-specific data whose content is not defined by the GIF specification.

Note that GIF-to-PNG converters should not attempt to perform byte reordering on the contents of the Application Extension. The data is simply transcribed without any processing except for de-blocking GIF sub-blocks.

Applications that formerly used GIF Application Extensions may define special-purpose PNG chunks to replace their application extensions. If a GIF-to-PNG converter recognizes the Application Identifier and is aware of a corresponding PNG chunk, it may choose to convert the Application Extension into that PNG chunk type rather than using gIFx.

5 Chunks Not Described Here

The definitions of some public chunks are being maintained by groups other than the core PNG group. In general, these are chunks that are useful to more than one application (and thus are not private chunks), but are considered too specialized to list in the core PNG documentation.

5.1 fRAc Fractal image parameters

The fRAc chunk will describe the parameters used to generate a fractal image. The specification for the contents of the fRAc chunk is being developed by Tim Wegner, twegner@phoenix.net.

In the future, chunks will be fully specified before they are registered.

6 Text Chunk Keywords

It is expected that special-purpose keywords for tEXt and zTXt chunks will be registered and will appear in this document. However, no such keywords have yet been assigned.

7 Deprecated Chunks

The chunks listed in this section are registered, but deprecated. Encoders are discouraged from using them, and decoders are not encouraged to support them.

7.1 gIFt GIF Plain Text Extension

The gIFt chunk was originally provided for backwards compatibility with the GIF89a Plain Text Extension, but gIFt is now deprecated because it suffers from some fundamental design flaws.

- GIF considers a Plain Text Extension to be a Graphic Rendering Block, just like an image, so a GIF datastream containing an image and a Plain Text Extension is really a multi-image datastream with ordering issues (like associating each Graphic Control Extension with the proper Graphic Rendering Block). PNG, being a single-image format with no provisions for handling these ordering issues, is not equipped to contain both IDAT and gIFt simultaneously. Since IDAT is required, gIFt must be discouraged.
- Another design flaw is that the Text Foreground Color and Text Background Color fields of the Plain Text Extension are converted to RGB, rather than being converted to RGBA or left as palette indexes. Therefore, transparency information can be lost.

The gIFt chunk contains:

```
Text Grid Left Position: 4 bytes (signed integer,
                            byte order and size converted)
Text Grid Top Position:  4 bytes (signed integer,
                            byte order and size converted)
Text Grid Width:         4 bytes (unsigned integer,
                            byte order and size converted)
Text Grid Height:        4 bytes (unsigned integer,
                            byte order and size converted)
Character Cell Width:    1 byte
Character Cell Height:   1 byte
Text Foreground Color:   3 bytes (R,G,B samples)
Text Background Color:   3 bytes (R,G,B samples)
Plain Text Data:         n bytes
```

Text Grid Left Position, Top Position, Width, and Height specify the text area position and size in pixels. The converter must reformat these fields from 2 bytes LSB-first unsigned integers to 4 bytes MSB-first signed or unsigned integers. Note that GIF defines the position to be relative to the upper left corner of the logical screen. If an oFFs chunk is also present, a decoder should assume that the oFFs chunk defines the offset of the image relative to the GIF logical screen; hence subtracting the oFFs values (converted from micrometers to pixels if necessary) from the Text Grid Left and Top Positions gives the text area position relative to the main PNG image.

Character Cell Width and Height give the dimensions of each character in pixels.

Text Foreground and Background Color give the colors to be used to render text foreground and background. Note that the GIF-to-PNG converter must replace the palette index values found in the GIF Plain Text Extension block with the corresponding palette entry.

The remainder of the chunk is the text to be displayed. Note that this data is not in GIF sub-block format, but is a continuous datastream.

8 Security Considerations

The normal precautions (see the Security considerations section of the PNG specification) should be taken when displaying text contained in the calibration name, unit name, and parameters of the pCAL chunk.

Applications must take care to avoid underflow and overflow of intermediate results when converting data from one form to another according to the pCAL mappings.

9 Appendix: Sample code

This appendix provides some sample code that can be used in encoding and decoding PNG chunks. It does not form a part of the specification. In the event of a discrepancy between the sample code in this appendix

and the chunk definition, the chunk definition prevails.

9.1 pCAL

This section provides some sample code for the PNG pCAL chunk, written in the C Programming Language. The `pCAL_encode()` function takes an array of floating-point numbers and produces an array of 16-bit samples that can be stored as grayscale PNG pixels. The `pCAL_make_lut()` function produces a lookup table that can be used to extract the original physical values from PNG pixels.

```
/* Sample code for the PNG pCAL chunk */
#include <math.h>
/* Math.h supplies double precision exp(), pow(), and
 * sinh() functions.  If your math.h doesn't supply sinh(),
 * use sinh(x)=(exp(x)+exp(-x))/2.
 */
unsigned short limit(long low, double x, long high)
{
    if (low < high){
       if(x < low)return low;
       else if (x > high) return high;
       else return x;
    }
    else {
       if(x < high)return high;
       else if( x > low) return low;
       else return x;
    }
}
int pCAL_encode (unsigned short *stored_sample, long n,
 float *physical_value, unsigned int m, int equation_type,
 long x0, long x1, float *p)

/* returns 0 (success)
 *        -1 (error, x0==x1)
 *        -2 (unknown equation type)
 * input:
 *         n: number of samples
 *         physical_value[0..n-1]
 *         m: PNG sample depth
 *         equation_type: from pCAL chunk
 *         x0, x1: stored sample to original sample mapping
 *         p[]: equation parameters, from pCAL chunk
 * output:
 *         stored_samples[0..n-1] (caller must allocate space)
 */
```

```
{
    double d,dm;  /* force double precision arithmetic */
    long isample, osample, k;
    d=x1-x0;
    dm=m;

    if (x1 != x0) {
        if(equation_type == 0){
            for (k=0; k<n; k++) {
               isample=.5+d*(physical_value[k] - p[0])/p[1];
               osample=limit(x0, isample, x1);
               stored_sample[k]= floor(((osample-x0)*dm
                                +floor(d/2))/d);
            }
        }

        else if(equation_type == 1){
            for (k=0; k<n; k++) {
               isample= .5+d*(log(physical_value[k]
                         - p[0])/p[1])/p[2];
               osample=limit(x0, isample, x1);
               stored_sample[k]= floor(((osample-x0)*dm
                                +floor(d/2))/d);
            }
        }

        else if(equation_type == 2){
            double factor;
            factor=d/log(p[2]);
            for (k=0; k<n; k++) {
               isample=.5+log((physical_value[k]
                        -p[0])/p[1])*factor;
               osample=limit(x0, isample, x1);
               stored_sample[k]= floor(((osample-x0)*dm
                                +floor(d/2))/d);
            }
        }
```

```
        else if(equation_type == 3){
            for (k=0; k<n; k++) {
               isample= .5+p[3]+d*asinh((physical_value[k]
                        -p[0])/p[1])/p[2];
               osample=limit(x0, isample, x1);
               stored_sample[k]= floor(((osample-x0)*dm
                                +floor(d/2))/d);
            }
        }
        else return (-2); /* ERROR, unknown equation type */
    }
    else return (-1); /* ERROR, x0 == x1 */
    return (0);
}

int pCAL_make_lut (float *physical_value, unsigned int m,
             int equation_type, long x0, long x1, float *p)

/* returns 0 (success)
 *        -1 (error, x0==x1)
 *        -2 (unknown equation type)
 * input:
 *         m: PNG sample depth
 *         equation_type: from pCAL chunk
 *         x0, x1: stored sample to original sample mapping
 *         p[]: equation parameters, from pCAL chunk
 * output:
 *         physical_value[0..m] (caller must allocate space)
 */

{
    double d, dm; /* force double precision arithmetic */
    long sample, osample;
    d=x1-x0;
    dm=m;

    if (x1 != x0) {
        if(equation_type == 0){
            for (sample=0; sample<=m; sample++){
               osample=floor((sample*d+floor(dm/2))/dm) + x0;
               physical_value[sample] = p[0] + p[1]*osample/d;
            }
        }
```

```
        else if(equation_type == 1){
            for (sample=0; sample<=m; sample++){
               osample=floor((sample*d+floor(dm/2))/dm) + x0;
               physical_value[sample] = p[0]
                        + p[1]*exp(p[2]*osample/d);
            }
        }
        else if(equation_type == 2){
            for (sample=0; sample<=m; sample++){
               osample=floor((sample*d+floor(dm/2))/dm) + x0;
               physical_value[sample] = p[0]
                        + p[1]*pow(p[2],osample/d);
            }
        }

        else if(equation_type == 3){
            for (sample=0; sample<=m; sample++){
               osample=floor((sample*d+floor(dm/2))/dm) + x0;
               physical_value[sample] =
                  p[0] + p[1]*sinh(p[2]*(osample-p[3])/d);
            }
        }
        else return (-2); /* ERROR, unknown equation type */
    }
    else return (-1); /* ERROR, x0 == x1 */
    return (0);
}
```

9.2 Fixed-point gamma correction

```
/*****************************************/
/* gamma-lookup.c 0.1.1 (Sat  5 Sep 1998)    */
/* by Adam M. Costello <amc @ cs.berkeley.edu> */
/*****************************************/

/* This is public domain ANSI C code.  It is example */
/* code for computing gamma correction lookup tables */
/* using integer arithmetic.                         */


#include <math.h>
#include <limits.h>


typedef unsigned char uint8;
```

```
#if UCHAR_MAX >= 0xffff
typedef unsigned char uint16;
#else
typedef unsigned short uint16;
#endif

#if UCHAR_MAX >= 0xffffffff
typedef unsigned char uint32;
#elif USHRT_MAX >= 0xffffffff
typedef unsigned short uint32;
#elif UINT_MAX >= 0xffffffff
typedef unsigned int uint32;
#else
typedef unsigned long uint32;
#endif


/*********************/
/* 16-bit arithmetic */

/* Precompute the log table (this requires floating point): */

void precompute16(uint16 L[511])
{
  int j;
  double f;

  f = 0xfe00 / log(1 / 510.0);  /* We want L[1] = 0xfe00. */

  for (j = 1;  j <= 510;  ++j)
    L[j] = log(j / 510.0) * f + 0.5;
}

/* Make a 256-entry gamma correction lookup table G[] with */
/* exponent g/pow(2,14), where g must not exceed 0xffff:   */

void gamma16(uint16 L[511], uint8 G[256], uint16 g)
{
  int i, j;
  uint16 x, y, xhi, ghi, xlo, glo;

  j = 1;
  G[0] = 0;

  for (i = 1;  i <= 255;  ++i) {
    x = L[i << 1];
    xhi = x >> 8;
    ghi = g >> 8;
    y = xhi * ghi;
```

```
    if (y > 0x3f80) {
      /* We could have overflowed later. */
      /* But now we know y << 2 > L[1].  */
      G[i] = 0;
      continue;
    }

    xlo = x & 0xff;
    glo = g & 0xff;
    y = (y << 2) + ((xhi * glo) >> 6) + ((xlo * ghi) >> 6);
    while (L[j] > y) ++j;
    G[i] = j >> 1;
  }
}


/*********************/
/* 32-bit arithmetic */

/* Precompute the log table (this requires floating point): */

void precompute32(uint32 L[511])
{
  int j;
  double f;

  f = 0x3fffffff / log(1 / 510.0);
  /* We want L[1] = 0x3fffffff. */

  for (j = 1;  j <= 510;  ++j)
    L[j] = log(j / 510.0) * f + 0.5;
}

/* Make a 256-entry gamma correction lookup table G[] with */
/* exponent g/pow(2,14), where g must not exceed 0xffff:   */

void gamma32(uint32 L[511], uint8 G[256], uint16 g)
{
  int i, j;
  uint32 x, y;

  j = 1;
  G[0] = 0;
```

```
  for (i = 1;  i <= 255;  ++i) {
    x = L[i << 1];
    y = (x >> 14) * g;
    while (L[j] > y) ++j;
    G[i] = j >> 1;
  }
}


/*************************************************/
/* floating-point arithmetic (for comparison) */

/* Make a 256-entry gamma correction lookup table G[i] with
 *  exponent g. */

void gamma_fp(uint8 G[256], double g)
{
  int i;

  G[0] = 0;

  for (i = 1;  i <= 255;  ++i)
    G[i] = pow(i/255.0, g) * 255 + 0.5;
}


/******************/
/* Test accuracy: */

#include <stdio.h>

/* Compare a gamma table against the gamma_fp version: */

void compare(uint8 G[256], double g)
{
  uint8 Gfloat[256];
  int i;

  printf("exponent = %f\n", g);
  gamma_fp(Gfloat,g);

  for (i = 0;  i <= 255;  ++i) {
    if (Gfloat[i] != G[i])
      printf("i = %3d, Gfloat = %3d, G = %3d\n", i, Gfloat[i],
          G[i]);
  }
}

/* Run comparisons for several values of g: */
```

```
main()
{
  double g[7] = { 2.2, 1.7, 1.45, 1.0, 1/1.45, 1/1.7, 1/2.2 };
  int k;
  uint16 L16[511];
  uint32 L32[511];
  uint8 G[256];

  precompute16(L16);
  precompute32(L32);

  printf("16-bit:\n");

  for (k = 0;  k < 7;  ++k) {
    gamma16(L16, G, (uint16) (g[k] * (1 << 14) + 0.5));
    compare(G, g[k]);
  }

  printf("\n32-bit:\n");

  for (k = 0;  k < 7;  ++k) {
    gamma32(L32, G, (uint16) (g[k] * (1 << 14) + 0.5));
    compare(G, g[k]);
  }
}
```

10 Appendix: Rationale

This appendix gives the reasoning behind some of the design decisions in the PNG extension chunks. It does not form a part of the specification.

10.1 pCAL

This section gives the reasoning behind some of the design decisions in the pCAL chunk. It does not form a part of the specification.

Redundant equation types

Equation types 1 and 2 seem to be equivalent. Why have both?

- We don't want to force people to do the exponentiation using ln() and exp(), since pow() may provide better accuracy in some floating-point math libraries. We also don't want to force people using base-10 logs to store a sufficiently accurate value of ln(10) in the pCAL chunk.

- When the base is e, we don't want to force people to encode a sufficiently accurate value of e in the pCAL chunk, or to use pow() when exp() is sufficient.

What are x0 and x1 for?

- First, x0 and x1 provide a way to recover the original data, losslessly, when the original range is not a power of two. Sometimes the digitized values do not have a range that fills the full depth of a PNG. For example, if the original samples range from 0 (corresponding to black) to 800 (corresponding to white), PNG requires that these samples be scaled to the range 0 to 65535. By recording x0=0 and x1=800 we can recover the original samples, and we indicate the precision of the data.
- Even if the original data had a range identical to a valid PNG image sample, like 0 (black) to 65535 (white), one might want to create a derived image by stretching the contrast in a limited intensity range containing the important details. For example, we might want to scale the samples so that 46000 becomes 0 (black) and 47000 becomes 65535 (white). As in the previous case, by recording x0=46000 and x1=47000, we can recover the original data samples that fell between 46000 and 47000.

Integer division

Why define integer divison to round toward negative infinity? This is different from many C implementations and from all Fortran implementations, which round toward zero.

We cannot leave the choice unspecified. If we were to specify rounding toward zero, we'd have to account for a discontinuity at zero. A division by positive d would map the 2d-1 values from -(d-1) through d-1 to zero, but would map only d values to any other value; for example, 3d through 4d-1 would be mapped to 3. Achieving lossless mappings in spite of this anomaly would be difficult.

11 Appendix: Revision History

- March 11, 1997: First public release.
- October 26, 1998:
  - Added pCAL chunk.
  - Deprecated the gIFT chunk.
  - Added sample gamma-correction code that uses integer arithmetic.

12 References

[ISO/IEC-8859]

International Organization for Standardization and International Electrotechnical Commission, "Information Technology—8-bit Single-Byte Coded Graphic Character Sets—Part 1: Latin Alphabet No. 1", IS 8859-1, 1998.
Also see sample files at
<ftp://ftp.uu.net/graphics/png/documents/iso_8859-1.*>

13 Credits

Editors

- Glenn Randers-Pehrson, randeg @ alumni.rpi.edu
- Tom Lane, tgl @ sss.pgh.pa.us (edited the first release of this document)

Contributors

Names of contributors not already listed in the PNG specification are presented in alphabetical order:

- John Bowler, jbowler @ acm.org
- Todd French, tfrench @ sandia.gov
- Alaric B. Williams, png @ abwillms.demon.co.uk

Trademarks

GIF is a service mark of CompuServe Incorporated.

Macintosh is a trademark of Apple Computer, Inc.

Windows is a trademark of Microsoft Corporation.

End of Extensions to the PNG Specification