

METHOD AND APPARATUS FOR IMAGE MANIPULATION VIA A DITHER MATRIX

5 FIELD OF THE INVENTION

The present invention relates to digital image processing. In particular, the invention relates to enhancing digital colour images while processing them for printing to a media substrate.

10 CROSS REFERENCES TO RELATED APPLICATIONS

Various methods, systems and apparatus relating to the present invention are disclosed in the following US Patents/ Patent Applications filed by the applicant or assignee of the present invention:

09/517539	6566858	6331946	6246970	6442525	09/517384	09/505951
6374354	09/517608	6816968	6757832	6334190	6745331	09/517541
10/203559	10/203560	10/203564	10/636263	10/636283	10/866608	10/902889
10/902833	10/940653	10/942858	10/727181	10/727162	10/727163	10/727245
10/727204	10/727233	10/727280	10/727157	10/727178	10/727210	10/727257
10/727238	10/727251	10/727159	10/727180	10/727179	10/727192	10/727274
10/727164	10/727161	10/727198	10/727158	10/754536	10/754938	10/727227
10/727160	10/934720	11/212702	11/272491	10/296522	6795215	10/296535
09/575109	6805419	6859289	6977751	6398332	6394573	6622923
6747760	6921144	10/884881	10/943941	10/949294	11/039866	11/123011
6986560	7008033	11/148237	11/248435	11/248426	10/922846	10/922845
10/854521	10/854522	10/854488	10/854487	10/854503	10/854504	10/854509
10/854510	10/854496	10/854497	10/854495	10/854498	10/854511	10/854512
10/854525	10/854526	10/854516	10/854508	10/854507	10/854515	10/854506
10/854505	10/854493	10/854494	10/854489	10/854490	10/854492	10/854491
10/854528	10/854523	10/854527	10/854524	10/854520	10/854514	10/854519
10/854513	10/854499	10/854501	10/854500	10/854502	10/854518	10/854517
10/934628	11/212823	10/728804	10/728952	10/728806	6991322	10/728790

10/728884	10/728970	10/728784	10/728783	10/728925	6962402	10/728803
10/728780	10/728779	10/773189	10/773204	10/773198	10/773199	6830318
10/773201	10/773191	10/773183	10/773195	10/773196	10/773186	10/773200
10/773185	10/773192	10/773197	10/773203	10/773187	10/773202	10/773188
10/773194	10/773193	10/773184	11/008118	11/060751	11/060805	11/188017
11/298773	11/298774	11/329157	6623101	6406129	6505916	6457809
6550895	6457812	10/296434	6428133	6746105	10/407212	10/407207
10/683064	10/683041	6750901	6476863	6788336	11/097308	11/097309
11/097335	11/097299	11/097310	11/097213	11/210687	11/097212	11/212637
11/246687	11/246718	11/246685	11/246686	11/246703	11/246691	11/246711
11/246690	11/246712	11/246717	11/246709	11/246700	11/246701	11/246702
11/246668	11/246697	11/246698	11/246699	11/246675	11/246674	11/246667
11/246684	11/246672	11/246673	11/246683	11/246682	10/760272	10/760273
10/760187	10/760182	10/760188	10/760218	10/760217	10/760216	10/760233
10/760246	10/760212	10/760243	10/760201	10/760185	10/760253	10/760255
10/760209	10/760208	10/760194	10/760238	10/760234	10/760235	10/760183
10/760189	10/760262	10/760232	10/760231	10/760200	10/760190	10/760191
10/760227	10/760207	10/760181	10/815625	10/815624	10/815628	10/913375
10/913373	10/913374	10/913372	10/913377	10/913378	10/913380	10/913379
10/913376	10/913381	10/986402	11/172816	11/172815	11/172814	11/003786
11/003616	11/003418	11/003334	11/003600	11/003404	11/003419	11/003700
11/003601	11/003618	11/003615	11/003337	11/003698	11/003420	6984017
11/003699	11/071473	11/003463	11/003701	11/003683	11/003614	11/003702
11/003684	11/003619	11/003617	11/293800	11/293802	11/293801	11/293808
11/293809	11/246676	11/246677	11/246678	11/246679	11/246680	11/246681
11/246714	11/246713	11/246689	11/246671	11/246670	11/246669	11/246704
11/246710	11/246688	11/246716	11/246715	11/246707	11/246706	11/246705
11/246708	11/246693	11/246692	11/246696	11/246695	11/246694	11/293832
11/293838	11/293825	11/293841	11/293799	11/293796	11/293797	11/293798
10/760254	10/760210	10/760202	10/760197	10/760198	10/760249	10/760263
10/760196	10/760247	10/760223	10/760264	10/760244	10/760245	10/760222
10/760248	10/760236	10/760192	10/760203	10/760204	10/760205	10/760206
10/760267	10/760270	10/760259	10/760271	10/760275	10/760274	10/760268

10/760184	10/760195	10/760186	10/760261	10/760258	11/293804	11/293840
11/293803	11/293833	11/293834	11/293835	11/293836	11/293837	11/293792
11/293794	11/293839	11/293826	11/293829	11/293830	11/293827	11/293828
11/293795	11/293823	11/293824	11/293831	11/293815	11/293819	11/293818
11/293817	11/293816	11/014764	11/014763	11/014748	11/014747	11/014761
11/014760	11/014757	11/014714	11/014713	11/014762	11/014724	11/014723
11/014756	11/014736	11/014759	11/014758	11/014725	11/014739	11/014738
11/014737	11/014726	11/014745	11/014712	11/014715	11/014751	11/014735
11/014734	11/014719	11/014750	11/014749	11/014746	11/014769	11/014729
11/014743	11/014733	11/014754	11/014755	11/014765	11/014766	11/014740
11/014720	11/014753	11/014752	11/014744	11/014741	11/014768	11/014767
11/014718	11/014717	11/014716	11/014732	11/014742	11/097268	11/097185
11/097184	11/293820	11/293813	11/293822	11/293812	11/293821	11/293814
11/293793	11/293842	11/293811	11/293807	11/293806	11/293805	11/293810
09/575197	09/575195	09/575159	09/575123	6825945	09/575165	6813039
6987506	09/575131	6980318	6816274	09/575139	09/575186	6681045
6728000	09/575145	09/575192	09/575181	09/575193	09/575183	6789194
6789191	6644642	6502614	6622999	6669385	6549935	09/575187
6727996	6591884	6439706	6760119	09/575198	6290349	6428155
6785016	09/575174	09/575163	6737591	09/575154	09/575129	6830196
6832717	6957768	09/575162	09/575172	09/575170	09/575171	09/575161

The disclosures of these applications and patents are incorporated herein by reference.

BACKGROUND OF THE INVENTION

- 5 A market has developed for inkjet printers that can dock directly with digital cameras.

These printers are expected to download and print the photographs quickly and with a minimum of user input. A significant part of the delay between docking the camera and printing the photographs, is the processing time needed to prepare the image data for the printhead.

The major processing tasks are JPEG decompression, colour space conversion, image rotation and halftoning to convert the colour values into dots. These tasks are briefly explained below.

- 5 Colour can be specified using three independent variables. The variables are essentially coordinates in a colour space. The same colour can be specified in different colour spaces using different variables. Each colour space has a particular use or application. RGB (red, green, blue) is the natural colour space for devices that emit light such as television screens or computer monitors. CMY (cyan, magenta, yellow) is the natural colour space for
- 10 devices that display images with reflected light such as printed images from a printer. $Y C_R C_B$ (luminance, chrominance red, chrominance blue) separates the luminance from the chrominance (usually abbreviated to 'chroma') channels for more convenient data compression. Humans are more sensitive to luminance than chroma so any change in chrominance resulting from compression and subsequent decompression will be less
- 15 noticeable than an equivalent change in luminance. This means both chroma channels can be heavily compressed as long as the luminance is lightly compressed. With two of the three channels heavily compressed, YCC image data is more efficiently handled by the processor.
- 20 Digital cameras capture images natively in RGB. For efficient storage, the images are converted to YCC and compressed. Image data downloaded from a camera is typically in sYCC which is a widely recognized standard form of YCC. This must then be colour space converted when it is output to a screen or printer.

If the image is downloaded to a printer, the data is converted to the printer's colour space and the separate colour channels are *halftoned* with a *dither matrix*. Halftoning exploits the eye's perception of a spatial average of printed dots to reproduce *contone* (continuous tone) images. Inkjet printers can either print a dot at any one of the addressable locations on the media, or not. However, dots dispersed over area of white (say) paper, will appear to eye as a contone shade somewhere between white and the dot colour, depending on the number of dots.

The dither matrix covers a small area of the image at a time. The matrix has a range of threshold values dispersed throughout its sites. The contone colour levels for each pixel are compared to the spatially corresponding threshold values within the matrix. If the contone level exceeds the threshold value, a dot of that colour is printed (or equivalently, a dot is printed if the contone level is greater than or equal to the threshold, or less than, or less than or equal to the threshold value). This will produce many micro-differences between the contone and halftone image, but the eye is largely insensitive to these high frequency differences.

To produce a colour image, the separate halftoned images for each of the three colour channels are superimposed by the printer on the media. Printers typically have cyan, magenta, yellow and sometimes black (to conserve the other inks and provide a 'truer' black). This is abbreviated to CMYK (Cyan, Magenta, Yellow and blackK). If the print resolution, or dots per inch (dpi) is high enough, halftoning can reproduce any colour in the printer's *gamut* (palette of printable colours). Accordingly, the individual dots in CMY(K) space are colour averaged by the eye to reproduce the colours of the original image.

Downloaded images can be manipulated and enhanced on a computer prior to printing.

However, if downloaded directly from a camera to a photo printer, the user does not have an opportunity to manually enhance and view the images prior to printing. Even so, it is possible to incorporate some relatively basic image enhancement into the printer that can

5 be automatically and uniformly applied to the images prior to printing.

One common and relatively basic image enhancement technique is *histogram expansion*.

It improves the colour contrast by expanding the range of colours present in the raw image data so that it is more evenly spread across the entire range of available colours. To do

10 this, it is necessary to collect image statistics and build histograms for each colour channel.

This involves the collection of the three colour levels for each pixel and recording the number of pixels that fall into a range of discrete colour level intervals to build the

histograms. Usually the original image will have histograms with at least one sparsely populated region. By re-assigning all the pixels in the sparsely populated region into one

15 of the colour levels, the rest of the histogram can expand into the vacated region.

Spreading the pixels more evenly across the histogram improves the colour contrast. As the number of pixels in the sparsely populated regions is insignificant, reassigning them to a single colour level has little, if any, detrimental effect on the image. So in the vast majority of cases, the net effect of histogram expansion is an enhancement of the image.

20

Unfortunately, the processing tasks associated with histogram expansion can delay printing. In some printing applications, there is an expectation that the printer will start printing the downloaded images almost instantly. Photograph printers that dock directly with the camera are one such example. The captured images will typically download from

25 the camera upon docking and automatically print the images to 6 inch by 4 inch photo

grade paper. Users prefer, if not expect, to see their photos being printed within a few seconds. More importantly, users expect good quality prints, but, as discussed above, computationally intensive image enhancement is counter to quickly initiating the printing of downloaded photos.

5

SUMMARY OF THE INVENTION

According to a first aspect, the present invention provides a method of manipulating contone image data to be halftoned with a dither matrix, the method comprising:

- determining at least one characteristic of the contone image data;
- 10 using the at least one characteristic to derive a secondary dither matrix from a predetermined primary dither matrix; and,
- halftoning the contone image data with the secondary dither matrix.

- Manipulating the dither matrix can be equivalent to performing image enhancement
- 15 manipulations to the input image data. However, it far less computationally intensive to manipulate the dither matrix instead of all the raw image data. For example, if the dither matrix is (say) 64 x 64, with each element being an 8 bit value, it has roughly 4 kilobytes of data. By comparison, a digital photograph (6 inches by 4 inches) at 3 mega-pixel native resolution, is about 10 megabytes of data. So, in this example, the number of
- 20 manipulations to the dither matrix is about three orders of magnitude less than that needed for the image data to achieve the same net effect. As discussed above, histogram expansions are a very common image enhancement technique and simply compressing the range of threshold values in the dither matrix as an inverse to the desired expansion of the histogram provides an equivalent result with far less processing. Furthermore, as the data
- 25 in the dither matrix is a tiny fraction of the input data, the dither matrix can be given added

complexity, or granularity, in order to achieve a better result than a normal histogram expansion, while still providing greater computation efficiency.

According to a second aspect, the present invention provides a print engine controller for
5 an inkjet printer, the print engine controller comprising:

a processor for receiving contone image data;
memory storing a predetermined primary dither matrix; wherein,
the processor is configured to determine at least one characteristic of the contone
image data and derive a secondary dither matrix from the primary dither matrix using the
10 at least one characteristic of the contone image data; such that,
the contone image data is halftoned with the secondary dither matrix prior to
printing.

Preferably, the contone image data has colour level values for pixels in the image, the
15 colour level values having a certain distribution within a predetermined range of discrete
colour levels, and, the at least one characteristic of the contone image data relates to said
certain distribution.

Preferably, the primary dither matrix has a range of threshold values and the secondary
20 dither matrix has a compressed range of threshold values for comparison to the colour level
values of the contone image data during halftoning.

In a further preferred form, the at least one characteristic related to the certain distribution
is the minimum number of contiguous discrete colour levels containing a predetermined
portion of the pixels, divided by the total number of levels in the predetermined range of

discrete colour levels. In some embodiments, the predetermined portion of the pixels is greater than 90%.

Optionally, the at least one characteristic related to the certain distribution is:

$$5 \quad (L_{\max} - L_{\min})/L_{\text{total}}$$

Where:

$L_{\max}$ is the maximum number of discrete colour levels if a top-most portion of the colour level values of the contone image data is disregarded;

$L_{\min}$ is the minimum number of discrete colour levels that contain a bottom-most
10 number of the colour level values of the contone image data; and,

L_{total} is the total number of levels in the range of discrete colour levels.

In these embodiments, the top-most portion may be the highest 5% of the colour level values of the contone image data. Likewise, the bottom-most portion may be the lowest
15 5% of the colour level values of the contone image data. In a less aggressive enhancement, the top and bottom-most portions may be 1%.

In preferred embodiments, the threshold values in the compressed range in the secondary matrix are determined in accordance with the following algorithm:

$$20 \quad T_{\text{new}} = L_{\min} + T_{\text{old}} \cdot (L_{\max} - L_{\min})/L_{\text{total}}$$

where:

T_{new} is the compressed threshold values in the secondary dither matrix; and,

T_{old} is the threshold value in the primary dither matrix.

- Optionally, at least some of the thresholds in the primary dither matrix are not whole numbers and the compressed threshold values in the secondary matrix are rounded or truncated to the nearest whole numbers. In other options, the threshold values occur in the primary dither matrix a predetermined number of times and the compressed threshold values occur in the secondary dither matrix a greater number of times, the greater number being approximately equal to $L_{total}/(L_{max}-L_{min})$ multiplied by the predetermined number associated with the threshold values of the primary matrix, or only one of two colliding threshold values from the primary matrix, that correspond to the compressed threshold.
- 10 In some preferred embodiments the processor only samples a portion of the pixels of the contone image data to determine L_{min} and L_{max} .

- Optionally, the colour levels values are an eight bit binary number such that there are 256 (2^8) levels in the range of discrete colour levels. Optionally, the dither matrix is 64 x 64 and the threshold levels range from 1 to 255 prior to compression.
- 15

BRIEF DESCRIPTION OF THE DRAWINGS

- 20 The invention will now be described, by way of example only, with reference to the preferred embodiments shown in the accompanying drawings in which:

Figure 1 shows a print engine pipeline in accordance with the present invention;

Figure 2 shows a dither matrix partially completed with threshold values;

- 25 Figure 3 shows a histogram of the image data for one of the colour channels;

Figure 4 shows the histogram of Figure 3 expanded to enhance colour contrast;

Figure 5 shows the dither matrix with compressed threshold values calculated to two decimal places; and,

Figure 6 shows the dither matrix with compressed threshold values rounded to the nearest

5 whole number.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

As discussed in the Background to the Invention, photograph printers have recently been
10 developed that dock directly with a digital camera and automatically printed the captured
images. They are expected to print the images quickly and with photographic quality.
Furthermore, these printers will only offer rudimentary image enhancement options, if any
at all. Users that want more complex image enhancement of their photos will download
the images to desktops or laptops and manipulate them with PhotoShop™ or similar
15 software (note that PhotoShop is a trademark of Adobe Systems Inc).

As the present invention offers basic image enhancement with computational efficiency, it
is well suited to photo printers. In light of this, it will be described with specific reference
to this application. However, skilled workers will readily appreciate that the invention is
20 not limited to photo printers and is suitable for a wide range of applications.

Figure 1 shows a print engine pipeline for the image data from the camera 1 to the
printhead 13. When the camera docks with the printer, the images are downloaded to the
print engine controller (PEC) 2 as EXIF (exchangeable image file data) JPEG (joint
25 photographic expert group) files in sYCC colour space (or standard YCC colour space).

The PEC 2 decompresses the images with a contone decoder unit (CDU) 3. If the image is too large 4, it is downsampled 5 as it is decoded.

As the pixel data from each JPEG MCU (minimum coding unit) becomes available, it is
5 rotated and converted to the specific CMY colour space 7 of the printhead 13.

Once the data is in the CMY space, the PEC 2 can collect image statistics and build a histogram 8. The collection of image statistics involves building a histogram of the number of occurrences of each color level. Once the histograms for the image are known,
10 the degree of histogram expansion can be determined. This can be done in any number of ways and one particular method will be discussed below with reference to Figures 3 and 4.

Expanding the histogram requires new maximum and minimum colour levels to be determined 9. That is, a minimum level, L_{min} , is determined and mapped to 0. All levels
15 between 0 and L_{min} are also mapped to 0. Likewise, L_{max} is determined and mapped to the highest colour level value. For example, if the colour levels are 8 bit numbers, the highest is 255. Every level between L_{max} and 255 is also mapped to 255.

However, as the present invention manipulates the dither matrix to enhance the image
20 rather than the colour levels in the image data, the only characteristic of the histogram that the PEC 2 needs is the degree of histogram expansion that would be caused by the L_{min} and L_{max} determined by the chosen method. There is no need to map any of the colour levels in the image data to new levels.

13

The histogram is expanded by a factor of $255/(L_{\max}-L_{\min})$. Hence, the corresponding compression of the range of threshold values in the dither matrix is given by:

$$T_{\text{new}} = L_{\min} + T_{\text{old}} \cdot (L_{\max} - L_{\min}) / 255 \quad \text{EQ(1)}$$

Where:

- 5 T_{new} is the compressed threshold; and,
 T_{old} is the original threshold value.

If the dither matrix size is 64 x 64, then compressing the threshold values involves the manipulation of about 4kB of data, whereas the equivalent expansion of the input colour
10 levels is a manipulation of about 10MB of data, or possibly more depending on image resolution. Compressing the dither matrix is more computationally efficient than expanding the image data by several orders of magnitude. This can dramatically reduce any delay between docking the camera and printing the downloading images, and it also allows more complex image enhance techniques via the dither matrix while still remaining
15 far less computationally intensive than performing equivalent techniques on the input data.

Once the compressed dither matrix values have been calculated 10, the CMY colour values determined at stage 7 of the print engine pipeline, are compared directly with the compressed threshold values of the dither matrix to produce a halftone image of each
20 colour channel 11. It will be appreciated that the same matrix can be used for each channel, or separate dither matrices derived from respective histograms could be used for each colour channel.

The halftoned images are sent to the printhead 13 for printing 12 to complete the pipeline.

Figure 2 shows an example of a 64 by 64 dither matrix. For simplicity, it is only partially completed with the range of threshold values. If the threshold values are 8 bit (corresponding to 8 bit colour levels), there are 255 thresholds. All 255 threshold levels appear in the matrix a number of times and the number of times a particular threshold occurs depends on the characteristics of the printhead and the aim to achieve a perceptually uniform colour space.

As discussed above, the dither matrix is repeatedly tiled across the contone image and the separate colour levels for pixel are compared to a corresponding threshold in the dither matrix. If the colour level exceeds the threshold, the printhead will eject a drop (of that particular colour) at that location, and if it is less than the threshold, no drop is ejected. Because the eye spatially averages the colour, it does not see the high frequency differences between the contone image and the halftoned image.

Referring now to Figure 4, the histogram for input image data is shown. $L_{\min}$ and $L_{\max}$ can be derived in any number of ways. For example, many images will not have a pixel at level 0, or a pixel at level 256. In this case, $L_{\min}$ and $L_{\max}$ can simply be the highest and lowest colour levels sampled. However, this does not take into account the tendency for the highest and lowest sampled colour levels to be outliers from the histogram distribution. Hence, using outliers as $L_{\min}$ and $L_{\max}$ usually means the histogram is not expanded as much as it should.

A better approach is to select a portion of the pixels at either end of the histogram and set the highest of the low end portion to be $L_{\min}$, and the lowest of the high end portion to be $L_{\max}$. Some experimentation may be required to optimize for each print engine pipeline,

but taking the top and bottom 1/256 (or approximately 0.4%) of the histogram will usually account for any outliers. In other words, 0.4% of the samples are below $L_{\min}$ and 0.4% of the samples are above $L_{\max}$. This approach is likely to provide visually better results than simply taking the extreme levels however, it is slightly more computationally intensive.

- 5 There is also a risk of over expansion by taking a percentage from the top and bottom of the histogram. Too much expansion can introduce visible contouring in areas of with a ramped colour gradient (because of the large colour differences between adjacent pixels after expansion). To guard against this, the processor might impose a maximum allowable expansion.

10

Once $L_{\min}$ and $L_{\max}$ have been determined, any samples at $L_{\min}$ or below are mapped to 0, and any sample at $L_{\max}$ or more are mapped to 256. As shown in Figure 5, the remainder of the histogram is then expanded between 0 and 255. There are spikes at 0 and 256 because they now contain all the original $L_{\min}$ and $L_{\max}$ samples, as well as the

15 outliers. However, this is not likely to have any detrimental effect on image quality. Traditionally, histogram expansion has been done with the level mapping function given by:

$$L_{\text{new}} = 256 \cdot (L_{\text{old}} - L_{\min}) / (L_{\max} - L_{\min}) \quad \text{EQ(2)}$$

- 20 This function would be applied to the input levels to determine the expanded levels, the expanded levels then compared to the threshold values of the dither matrix. Therefore, it is equivalent and computationally easier to compress the range of threshold values by the inverse of the expansion that would otherwise be performed on the histogram. The algorithm for mapping each of the original threshold values to new threshold values is
- 25 given by equation 1 discussed above.

Figure 5 shows the dither matrix of Figure 2 compressed in accordance with equation 1.

Unfortunately, the compressed threshold values must be rounded or truncated to whole numbers as the hardware requires the dithering process to be a whole number comparison.

- 5 Hence the threshold values of Figure 2 are rounded to give the whole number thresholds shown in Figure 6.

Manipulating the input levels via the dither matrix also presents an opportunity to improve the image enhancement beyond that offered by histogram expansion.

10

During histogram expansion, mapping the old levels to new levels via equation 2 involves rounding as the new levels need to be integers (because of the printer hardware). Hence, some of the levels in the expanded histogram have no samples in them. As shown in Figure 4, these appear as gaps in the histogram. The colour difference between samples on either side of a gap is greater than the difference between the same samples in the unexpanded histogram. These increased colour differences are more likely to produce visible contouring in the printed image.

15

- Likewise, rounding (or truncating) the compressed threshold values in the dither matrix causes some of the thresholds to collide. For example, the shaded thresholds in the compressed matrix of Figure 6 are pairs of equal value whereas, the corresponding threshold in the original matrix (see Figure 2) were not equal. Consequently, the number of dots added when moving from one tone level to the next will not be smooth. Again, this increases the risk of visible contouring in the printed image.

25

Increasing the granularity of the dither matrix (that is, using threshold values greater than 8 bits), would avoid collisions and smooth out the number of dots added to the halftone image with each subsequent tone level. Unfortunately, in most printers, the comparison of the dither matrix with the contone CMY levels is a hardware function that requires the threshold values to be 8 bit integers only.

Alternatively, the software can reconstruct the original dither matrix at a higher granularity for a perceptually smooth transition between tone levels. This will effectively add sub-levels to each threshold so any collisions will be between 2 sub-levels and consequently involve far fewer dots.

The problem of threshold collision and the solution of higher granularity are illustrated in the tables below. Firstly, Table 1 shows the number of times some threshold values appear in the original matrix and the compressed matrix. The 64 by 64 matrix has 4096 elements so each of the 1 to 255 original threshold values will occur 4096/255, or about 16 times in the original matrix. When the matrix is compressed, rounding causes some of the original thresholds to map to the same compressed threshold. These colliding thresholds appear 32 times in the compressed matrix (e.g. compressed threshold 38 in Figure 6). So the number of dots added when moving from tone level 37 to 38, will be twice that of moving 38 to 39, or 36 to 37. Hence the increased risk of visually perceptible contouring.

Original Threshold	Occurrences in Original Matrix	Compressed Threshold (Rounded)	Occurrences in Compressed Matrix
34	16	46	16

18

35	16	47	32
36	16		
37	16	48	16
38	16	49	16
39	16	50	32
40	16		

Table 1: Inconsistent Occurrences of Compressed Thresholds because of Threshold Collision

- 5 For the purposes of this example, we will assume the thresholds in the compressed matrix should occur 21 times each ($4096/(L_{\max}-L_{\min}) = 21.005$ - so one threshold would occur 22 times) for least risk of visible contours. By adding granularity to the original matrix, the occurrences of the compressed thresholds can be more uniform. For example, if the original threshold values are 12 bit (or rather 8.2 bit) the original matrix essentially gains
- 10 extra threshold levels. Table 2 sets out the finer grained original thresholds and the smoothing effect this has on the compressed matrix.

Original Threshold	Occurrences in Original Matrix	Compressed Threshold (Rounded)	Occurrences in Compressed Matrix
35.00	4	47	20
35.25	4	47	
35.50	4	47	
35.75	4	47	

19

36.00	4	47	
36.25	4	48	20
36.50	4	48	
36.75	4	48	
37.00	4	48	
37.25	4	48	
37.50	4	49	20
37.75	4	49	
38.00	4	49	
38.25	4	49	
38.50	4	49	
38.75	4	50	24
39.00	4	50	
39.25	4	50	
39.50	4	50	
39.75	4	50	
40.00	4	50	

With the increased granularity, there are smaller inconsistencies in the number of times the compressed thresholds appear in the final matrix. Most compressed thresholds occur 20 times with about a quarter appearing 24 times. Therefore, moving from one tone level to

5 the next is smoother and there is far less likelihood of any visible contouring.

Instead of increasing the granularity of the original matrix, it may be simpler to multiply the number of times each threshold appears in the original matrix by the inverse of the

compression to determine the number of times the compressed threshold appears in the final matrix. Referring again to the above example, the inverse of the compression is $256/(L_{\max} - L_{\min}) = 1.313$. If an original threshold occurred 16 times in the matrix, the compressed threshold might ideally occur $1.313 \times 16 = 21$ (with one threshold occurring 22 times). Of course, if two original thresholds are colliding at a single compressed threshold, the occurrences of only one of the original thresholds should be used to determine how many times the compressed threshold occurs.

Enhancing the image by manipulating the dither matrix means that gathering image statistics can be reasonably heavily optimized because the impact of granularity issues on the final matrix (and therefore the printed image) is relatively minor. For example, in some applications, the histogram need not have 256 levels. 64 levels (6 bit) may be adequate. When building the histogram, it may not be necessary to collect statistics for every pixel. Sampling as little as 1 in 256 pixels can be within acceptable output errors. When calculating $L_{\max}$ and $L_{\min}$, portions of the image that can no longer have any impact on the end points of the histogram, may be ignored. Likewise, reasonable results may be obtained by computing or looking up a minimum and maximum value from each colour point, rather than each of the individual CMY values. All these optimizations serve to reduce the processing burden on the PEC, and so shorten the time between docking the camera and printing the images.

The invention has been described here by way of example only. Skilled workers in this field will readily recognize many variations and modifications that do not depart from the spirit and scope of the broad inventive concept.

CLAIMS

1. A method of manipulating contone image data to be halftoned with a dither matrix, the method comprising:

5 building a histogram of the colour levels of the contone image;
determining the number of colour levels falling within a predetermined distribution of the histogram;

deriving a secondary dither matrix from a predetermined primary dither matrix based on the number of colour levels falling within the predetermined distribution of the histogram; and,

10 halftoning the contone image data with the secondary dither matrix.

2. A method according to claim 1, wherein the secondary dither matrix is derived from the primary dither matrix using the number of colour levels falling within the predetermined distribution of the histogram, divided by the total number of colour levels in a predetermined range
15 of discrete colour levels.

3. A method according to claim 1, wherein the predetermined distribution is a distribution containing at least 90% of the colour levels in the histogram.

20 4. A method according to claim 2, wherein the secondary dither matrix is derived from the primary dither matrix according to the formula:

$$(L_{\max} - L_{\min})/L_{\text{total}}$$

where:

25 $L_{\max}$ is the highest colour level when the colour level values of the contone image outside of the predetermined distribution are disregarded;

$L_{\min}$ is the lowest colour level when the colour level values of the contone image outside of the predetermined distribution are disregarded; and,

L_{total} is the total number of colour levels in the predetermined range of discrete colour levels.

30 5. A method according to claim 4, wherein the highest 5% of the colour level values of the contone image data and the lowest 5% of the colour level values of the contone image data are disregarded.

6. A method according to claim 4, wherein the highest 1% of the colour level values of the contone image data and the lowest 1% of the colour level values of the contone image data are disregarded.

7. A method according to claim 4, wherein the secondary dither matrix is populated with compressed threshold values derived from the primary dither matrix according to the following formula:

$$T_{\text{new}} = L_{\text{min}} + T_{\text{old}} \cdot (L_{\text{max}} - L_{\text{min}}) / L_{\text{total}}$$

where:

T_{new} is the compressed threshold values in the secondary dither matrix; and,
 T_{old} is a threshold value in the primary dither matrix.

8. A method according to claim 7, wherein the compressed threshold values in the secondary matrix are rounded or truncated to the nearest whole numbers.

9. A print engine controller for an inkjet printer, the print engine controller comprising:
 a processor for receiving contone image data; and
 memory for storing a primary dither matrix, wherein,
 the processor is configured to build a histogram of the colour levels of the contone image data, determine the number of colour levels falling within a predetermined distribution of the histogram, and derive a secondary dither matrix from the primary dither matrix based on the number of colour levels falling within the predetermined distribution of the histogram, and
 the processor is further configured to halftone the contone image data using the secondary dither matrix.

10. A print engine controller according to claim 9, wherein the processor calculates the number of colour levels falling within the predetermined distribution of the histogram, divided by the total number of colour levels in a predetermined range of discrete colour levels.

11. A print engine controller according to claim 10, wherein the processor performs the calculation in accordance with the following formula:

$$(L_{\text{max}} - L_{\text{min}}) / L_{\text{total}}$$

where:

L_{max} is the highest colour level when the colour level values of the contone image outside of the predetermined distribution are disregarded;

$L_{\min}$ is the lowest colour level when the colour level values of the contone image outside of the predetermined distribution are disregarded; and,

L_{total} is the total number of colour levels in the predetermined range of discrete colour levels.

5

12. A print engine controller according to claim 11, wherein the processor disregards the highest 5% of the colour level values of the contone image data and the lowest 5% of the colour level values of the contone image data to obtain the predetermined distribution.

10

13. A print engine controller according to claim 11, wherein the processor populates secondary dither matrix with compressed threshold values derived from the dither matrix according to the following formula:

$$T_{\text{new}} = L_{\min} + T_{\text{old}} \cdot (L_{\max} - L_{\min}) / L_{\text{total}}$$

where:

15

T_{new} is the compressed threshold values in the secondary dither matrix; and,

T_{old} is a threshold value in the primary dither matrix.

14. A print engine controller according to claim 13, wherein the processor truncates or rounds the compressed threshold values in the secondary matrix to the nearest whole numbers.

20

15. A print engine controller according to claim 9, wherein the processor samples a limited portion of the pixels of the contone image data to determine $L_{\min}$ and $L_{\max}$.