Technical Information

Blend update

Linotype-Hell

An earlier document in this series, Blends and Shadestepping, covers the relationship between imagesetter addressability[1], screen ruling, halftone dot percentage change, steps in the blend and length of the blend. Through the use of a simple formula, you can predict the occurrence of shadestepping in blends created by PostScript** illustration programs. However in the time since that article was published, another method for creating blends has come into common usage within PostScript. And though the rules described in Blends and Shadestepping haven't changed, it is important to know how to apply them to the different types of blends.

[1]Addressability is a measure of the number of marks that an imagesetter can make in a given distance. It is often referred to as resolution. For an explanation of the difference between these terms, please refer to the Linotype-Hell technical information piece entitled Addressability and Spot Size.

In this document, we will look at three types of blends and how they are created. For background on some blend basics, please refer to the Linotype-Hell technical information piece entitled Blends and Shadestepping.

Blend terminology

What is commonly known as a blend may be called by a number of different names: degradés, fountains, gradations, graduated tints, and vignettes. While there are differences in the way these terms are used, for all practical purposes the term blend can be used generically. A reasonable definition for all of theses terms is 'a gradual transition from one gray shade or color to another gray shade or color.'

In the past, these types of effects were usually created through the use of airbrush illustrations or special photographic halftone screens. An airbrushed illustration requires a highly skilled artist, and the special photographic screens were limited because many important factors were pre-determined (i.e., screen ruling, screen angle, length, and the beginning and end gray). It is no surprise that people have tried to achieve effects with computer-generated blends that would be very difficult with conventional photographic ones.

Blends versus scans

Accurately reproducing a photograph is a difficult task, but producing a satisfying complex blend can be even more difficult. Why is this so?

- Many blends use every step in a tonal range. (Few scans do that.)
- Blends are often made up of data that is so mathematically pure that any physical aberration is easily visible. (There is usually enough variation in scans to discount this.)
- Users sometimes overlook the range of what is reasonably achievable and attempt to create blends of great length at low imagesetter addressabilities.
- While scanned images are usually larger in file size than most blends, many blends may be used in a single synthetic image. This increases overall complexity. (Remember that file size alone is not an accurate indicator of how fast a raster image processor (RIP) will process an image. Some complex synthetic images may take longer to process than much larger scanned images. Blends add considerable complexity to the complexity of a synthetic image.)

The page description language, the application software, and the output device must all work well together if the results are to be optimal.

Blend methods

There are three primary methods for doing blends: duplication blends, synthetic data blends, and sampled data blends.

Duplication blends - You may remember that when Adobe Illustrator** was first introduced, the advertisements stressed a tool that would allow you to gradually turn the outline shape of the letter ‘S’ into a drawing of a swan. This was achieved by duplication and slight changes through a number of steps from ‘S’ to swan. A much more common use for this tool is in the creation of a blend. By repeatedly duplicating an object, moving it slightly, and changing the dot percentage by a small amount, you can create the impression of a blending of gray or color. Applications that allow this usually allow you to chose the number of steps. The more steps you choose, the larger the file.

One common duplication technique is to create two boxes, one with the beginning color and one with the end color, and create a set of slightly changed duplicates between the two. The number of steps should not exceed 256, since that exceeds the number of grays that PostScript supplies for halftoning, and only creates more code than necessary. (A good general recommendation is to make the number of steps equal to the distance in printer’s points from the edge of the first box to the end of the next.)

Examples of duplication blends:
- Aldus FreeHand** blend menu selection[2]
- Adobe Illustrator blend tool.

[2]Olav Kern, in his book Real World FreeHand (Peachpit Press, 1991), says that in most cases, a FreeHand blend prints faster than a FreeHand graduated fill (see below). He adds that applying traps to abutting graduated fills is difficult.

File size chart:

Method	File size
Duplicated blends	
FreeHand blend	51K
Illustrator blend	27K
Synthetic blends	
FreeHand graduated fill	6K
QuarkXPress blend	6K
Sampled data blends	
Photoshop blend	
72 ppi	105K
150 ppi	444K
300 ppi	1,761K

Note: These file sizes are based on a 4" x 5" single color blend from 0% to 100%.

Multiple color blends, particularly sampled data blends, can be much larger in size. A CMYK sampled data blend will be four times as large as those values shown above.

And remember that if you created an 8" x 10" sampled data blend at the resolutions shown above, you would actually quadruple the file size. See the Linotype-Hell technical information piece entitled Scanned File Size for more information.

Synthetic data blends - There are a number of ways to create blends without using duplication. For example, there is an operator in the PostScript page description language called the image operator[3] which not only allows you to print scanned (i.e. sampled) data, but also to create synthetic data. This operator can be used to make a blend by creating a synthetic image and stretching it over a given area. Instead of setting a number of steps, as in duplication blends, the image operator creates a synthetic image that extends from one gray (or color) value to another. Generally, synthetic data blends are much smaller in file size than duplication blends. This is because the boxes in a duplication blend are repeated over and over again, which takes up more space. (See chart to left.)

Examples of synthetic data blends:
- Aldus FreeHand graduated fill (either log or linear)
- QuarkXPress** 3.1 Linear Blend option (under Show Colors/View menu).

[3]See the following references for more information on using the image operator for creating a synthetic blend:

- PostScript Language Program Design, Adobe Systems Inc., Addison Wesley, 1988, pgs. 123-130. (This is the so-called ‘Green Book’)
- Real World PostScript, Steve Roth, ed., Addison Wesley, 1988, pgs. 211-213.

These references show a synthetic blend in its simplest form. Each manufacturer of illustration software may take a slightly different (and probably more sophisticated) approach to the creation of a synthetic blend.

Sampled data blends - Blends may also be created as sampled data. Data in this form is commonly known as a bitmap[4]. High-end systems (like Linotype-Hell’s ChromaCom*) create blends this way, as do image manipulation programs (like Adobe Photoshop** or Letraset ColorStudio**). Within PostScript, it is again the image operator that allows this function.

[4]A bitmap is a grid of sampled data. Since each spot on the grid is usually made up of a byte of information, some people have suggested the use of the term bytemap.

This method requires that data be created at a certain spatial resolution. The higher you set this resolution, the larger the file. (See chart to left.) Strictly speaking this data has been created, not sampled, but it can be manipulated just as if it had been scanned.

Sampled data may be filtered, which can come in handy for adding some precisely located noise that can blur the distinction between one gray step and the next (This ‘noise’ should not be applied indiscriminately throughout the image, for the result then would only be a grainy look. Ideally, it should be added very subtly at the edge between one gray and the next.)

Examples of sampled data blends:
• Adobe Photoshop blend tool and Letraset ColorStudio paint bucket tool
• High-end system blends

In general, synthetic blends are the smallest in file size while sampled data blends are the largest. And though you can see differences in file size or in the way the code is created, it may be very difficult to tell any of these blends apart solely by looking at the final output.

Mysterious blend problems

No matter which method you use to create a blend, at times it seems like nothing you do solves a shadestepping problem. Sometimes simply changing a single factor can result in a dramatic improvement. Here is a list of possible things to change:

Variables involved	**Possible solution**
• Steps in a duplication blend	Add steps (but not beyond 256)
• Length of blend	Decrease length of blend
• Percentage change in blend	Increase percentage change of color
• Addressability	Increase imagesetter addressability
• Screen ruling	Decrease screen ruling
• Screen angle	Change the screen angle[5]
• Halftone dot shape	Try elliptical dot to avoid mid-tone jump[6]
• Maximum density	Decrease dmax[7]

[5]Changing the screen angle of a single color blend may help for reasons that are as yet unclear.

[6]This will only help with a single step caused by mid-tone jump. Numerous repeating steps are unlikely to go away solely through a change of halftone dot shape.

[7]Films with a high dmax are more likely to exhibit artifacts. When setting a production standard, keep in mind that most imagesetter films perform best at dmaxes below 4.0.

0	1	2	3	4	5	6	7
8	9	10	11	12	13	14	15
16	**17**	18	19	20	21	22	23
24	25	**26**	27	28	29	30	31
32	33	34	**35**	36	37	38	39
40	41	42	43	**44**	45	46	47
48	49	50	51	52	**53**	54	55
56	57	58	59	60	61	**62**	63
64	65	66	67	68	69	70	**71**
80	81	82	83	84	85	86	87
72	73	74	75	76	77	78	79
88	**89**	90	91	92	93	94	95
96	97	**98**	99	100	101	102	103
104	105	106	**107**	108	109	110	111
112	113	114	115	**116**	117	118	119
120	121	122	123	124	**125**	126	127
128	129	130	131	132	133	**134**	135
136	137	138	139	140	141	142	**143**
144	145	146	147	148	149	150	151
152	153	154	155	156	157	158	159
160	**161**	162	163	164	165	166	167
168	169	**170**	171	172	173	174	175
176	177	178	**179**	180	181	182	183
184	185	186	187	**188**	189	190	191
192	193	194	195	196	**197**	198	199
200	201	202	203	204	205	**206**	207
208	209	210	211	212	213	214	**215**
216	217	218	219	220	221	222	223
224	225	226	227	228	229	230	231
232	**233**	234	235	236	237	238	239
240	241	**242**	243	244	245	246	247
248	249	250	**251**	252	253	254	255
256	257	258	259	**260**	261	262	263
264	265	266	267	268	**269**	270	271
272	273	274	275	276	277	**278**	279
280	281	282	283	284	285	286	**287**
288							

Figure 1 - When more than 256 grays are available for a particular screen ruling and imagesetter addressability, the PostScript RIP must choose 256. For example, in this case removing every ninth available gray step does the trick. (Removed values are shown in bold.) Because of this, some neighboring halftone dots in a blend will increase in size by two laser spots (for example, 187 to 189) rather than by one. This may be visible.

Most of these solutions increase the number of grays in some fashion or another. Generally it is sufficient for an imagesetter addressability and screen ruling combination to be able to support 256 levels of gray. But in some cases, the ratio between the number of available grays and the 256 values allowed by PostScript can play a subtle role in shadestepping. For example, if the screen ruling and imagesetter addressability are capable of producing 289 gray levels, that means that the halftone cell for that combination of factors contains 288 laser spots. The smallest halftone dot would be made up of one laser spot. Each successive halftone dot gets larger as more laser spots are added. Solid black (a 100% halftone dot) would be made up of 288 laser spots. Paper white (0%) would have none.

Since PostScript allows 256 gray levels, the RIP chooses 256 of the 288 available grays. Thirty-two can’t be used. In this case, if the RIP removes every ninth value, it is left with 256 gray levels (288 - 256 = 32, 288/32 = 9). This means that in the resulting blend, every eighth gray step increases by two pixels rather than one. (See Figure 1.) In some cases, depending on a variety of factors, the two pixel steps may be visible. Changing the screen ruling or imagesetter addressability may help solve this problem.

Gray levels

What can be done to improve the general quality of blends? One key area is the number of gray levels that the page description language allows. This can be divided into two parts: gray levels that can be used for image data, and gray levels that can be used for halftoning. Adobe has announced that it has increased the top number of gray levels for image data from 256 (8 bit) to 4096 (12 bit) in PostScript Level 2. However, no change has been made in the number of grays available for halftoning. It remains at 256 in PostScript Level 2. Should this have been increased? Probably not. Blends would benefit from such an increase, but the costs of slower processing might not be worth it. And, there are other ways to improve shadestepping problems that may be better in the long run, for example:

- Noise, dithering, or error diffusion techniques may be used to mask the shadestepping in a blend. For example, the RIP 60 has a feature called the Equalizer which, when turned on, adds a small amount of noise to the halftone screening process.
- Properly created sampled data for a blend can reduce shadestepping. (This is one reason why high-end blends from ChromaCom generally have a quality edge over desktop blends.)
- Sampled data blends may be filtered via software. Effective filtering will improve the look of a blend.[8]
- Improved blending techniques by software manufacturers may increase quality. (Both synthetic data and sampled data blends may be able to take advantage of the higher image data ceiling in PostScript Level 2.)
- Increases in demand for higher screen ruling require higher imagesetter output addressabilities. Higher addressability and RIP speeds may allow higher addressabilities to be more commonly used.

[7]Adobe Systems has recently published a book called Design Essentials which covers many design aspects of using Adobe Illustrator and Adobe Photoshop. This book contains a section on blends, as well as a lot of other useful information. The book is available through Hayden, a division of Prentice Hall Computer Publishing. They may be reached at 1-800-428-5331.

Conclusion

As a final point, remember that what shows up as shadestepping on film may not appear in the printed piece, particularly in the case of lighter colors like cyan and yellow. Shadestepping can be affected by factors in printing such as ink dmax, color combinations, and paper surface.

Comments

Please direct any questions or comments to:

Jim Hamilton, Marketing Department
Linotype-Hell Company
425 Oser Avenue
Hauppauge, NY 11788

Acknowledgements

I would like to thank Gus Barbuto, Danny Johnson, Dave Klippel, Ken Larsen, Jim Mauro, Marco Mularoni, and Steve Ting of Linotype-Hell Company for their help in producing this document. Also special thanks to Grant Ruiz of Adobe Systems, and graphic arts consultant Brian Lawler.

Part Number 3315, 8/92