

Literature Notes

Sam Moore, David Gow
Faculty of Engineering, Computing and Mathematics, University of Western Australia
March 2014

Contents

1	Postscript Language Reference Manual[1]	1
2	Portable Document Format Reference Manual[2]	2
3	Portable Document Format (PDF) — Finally...[3]	3
4	Pixels or Perish [4]	4
5	Embedding and Publishing Interactive, 3D Figures in PDF Files[5]	5
6	27 Bits are not enough for 8 digit accuracy[?]	5
7	What every computer scientist should know about floating-point arithmetic[6]	5
8	Compositing Digital Images[7]	6
9	Bresenham’s Algorithm: Algorithm for computer control of a digital plotter[8]	6
10	Quad Trees: A Data Structure for Retrieval on Composite Keys[9]	7
11	Xr: Cross-device Rendering for Vector Graphics[10]	7
12	Glitz: Hardware Accelerated Image Compositing using OpenGL[11]	7
13	Boost Multiprecision Library[12]	7
14	A CMOS Floating Point Unit[13]	7
15	Simply FPU[14]	8
16	Floating Point Package User’s Guide[15]	8
17	Low-Cost Microarchitectural Support for Improved Floating-Point Accuracy[17]	8
18	Accurate Floating Point Arithmetic through Hardware Error-Free Transformations[18]	8

1 Postscript Language Reference Manual[1]

Adobe’s official reference manual for PostScript.

It is big.

2 Portable Document Format Reference Manual[2]

Adobe's official reference for PDF.

It is also big.

3 Portable Document Format (PDF) — Finally...[3]

This is not spectacularly useful, is basically an advertisement for Adobe software.

Intro

- Visual communications has been revolutionised by computing
- BUT there have always been problems in exchanging formats
- Filetypes like text, rich text, IGES, DXF, TIFF, JPEG, GIFF solve problems for particular types of files only
- PDF solves everything for everyone; can include text, images, animation, sound, etc

PDF Features

- Raster Image Process (RIP) — For printing (presumably also displaying on screen)
- Originally needed to convert to PS then RIP, with PS 3 can now RIP directly.
- Reduced filesize due to compression
- Four major applications - Stoy 1999[?]
 1. Download files from internet
 2. Files on CDs
 3. Files for outputting to printers
 4. Conventional [commercial scale?] printing
- List of various (Adobe) PDF related software
 - Includes software for PS that converts to/from PDF
 - So PS was obviously pretty popular before PDF
- Can Optimize for screen/printer [not clear how]
- Can compress for size

4 Pixels or Perish [4]

“The art of scientific illustration will have to adapt to the new age of online publishing” And therefore, JavaScript libraries (D³) are the future.

The point is that we need to change from thinking about documents as paper to thinking of them as pixels. This kind of makes it related to our paper, because it is the same way we are justifying our project. It does mention precision, but doesn't say we need to get more of it.

I get the feeling from this that Web based documents are a whole bunch of completely different design philosophies hacked together with JavaScript.

This paper uses Metaphors a lot. I never met a phor that didn't over extend itself.

Intro

- Drawings/Pictures are ornaments in science but they are not just ornamental
- Processes have changed a lot; eg: photographic plates → digital images
- “we are about to turn the page — if not close the book — on yet another chapter in publishing history.” (HO HO HO)
- It would be cool to have animated figures in documents (eg: Population pyramid; changes with time); not just as “supplements”
- In the beginning, there was PostScript, 1970s and 1980s, John Warnock and Charles Geschke, Adobe Systems
- PS is a language for vector graphics; objects are constructed from geometric primitives rather than a discrete array of pixels
- PS is a complete programming language; an image is also a program; can exploit this to control how images are created based on data (eg: Faces)
- PDF is “flattened” PS. No longer programable. Aspires to be “virtual paper”.
- But why are we using such powerful computing machines just to emulate sheets paper? (the author asks)

Web based Documents

- HTML, CSS, JavaScript - The Axis of Web Documents
 - HTML - Defines document structure
 - CSS - Defines presentation of elements in document
 - JavaScript - Encodes actions, allows dynamic content (change the HTML/CSS)
- <canvas> will let you draw anything (So in principle don't even need all of HTML/CSS)
 - Not device independent
 - “Coordinates can be specified with precision finer than pixel resolution” (**TODO: Investigate this?**)
 - JavaScript operators to draw things on canvas are very similar to the PostScript model
- SVG — Same structure (Document Object Model (DOM)) as HTML
 - “Noun language”
 - Nouns define lines/curves etc, rather than paragraphs/lists
 - Also borrows things from PostScript (eg: line caps and joints)
 - IS device independent, “very high precision” (**TODO: Investigate**)
 - JavaScript can be used to interact with SVG too
- D³ (Data Driven Documents) - A JavaScript library
 - Idea is to create or modify elements of a DOM document using supplied data
 - <https://github.com/mbostock/d3/wiki>
- We are in a new Golden Age of data visualisation
- Why do we still use PDFs?
 - PDFs are “owned” by the author/reader; you download it, store it, you can print it, etc
 - HTML documents are normally on websites. They are not self contained. They often rely on remote content from other websites (annoying to download the whole document).
- **Conclusion** Someone should open up PDF to accept things like D³ and other graphics formats (links nicely with [5])
- Also, Harry Potter reference

5 Embedding and Publishing Interactive, 3D Figures in PDF Files[5]

- Linkes well with [4]; I heard you liked figures so I put a figure in your PDF
- Title pretty much summarises it; similar to [4] except these guys actually did something practical

6 27 Bits are not enough for 8 digit accuracy[?]

Proves with maths, that rounding errors mean that you need at least q bits for p decimal digits. $10^p < 2^{q-1}$

- Eg: For 8 decimal digits, since $10^8 < 2^{27}$ would expect to be able to represent with 27 binary digits
- But: Integer part requires digits bits (regardless of fixed or floating point representation)
- Trade-off between precision and range
 - 9000000.0 \rightarrow 9999999.9 needs 24 digits for the integer part $2^{23} = 8388608$
- Floating point zero = smallest possible machine exponent
- Floating point representation:

$$y = 0.y_1y_2\dots y_q \times 2^n$$

- Can eliminate a bit by considering whether $n = -e$ for $-e$ the smallest machine exponent (???)
 - Get very small numbers with the same precision
 - Get large numbers with the extra bit of precision

7 What every computer scientist should know about floating-point arithmetic[6]

- Book: *Floating Point Computation* by Pat Sterbenz (out of print... in 1991)
- IEEE floating point standard becoming popular (introduced in 1987, this is 1991)
 - As well as structure, defines the algorithms for addition, multiplication, division and square root
 - Makes things portable because results of operations are the same on all machines (following the standard)
 - Alternatives to floating point: Floating slasi and Signed Logarithm (TODO: Look at these, although they will probably not be useful)
- Base β and precision p (number of digits to represent with) - powers of the base can be represented exactly.
- Largest and smallest exponents e_{min} and e_{max}
- Need bits for exponent and fraction, plus one for sign
- “Floating point number” is one that can be represented exactly.
- Representations are not unique! $0.01 \times 10^1 = 1.00 \times 10^{-1}$ Leading digit of one \implies “normalised”
- Requiring the representation to be normalised makes it unique, **but means it is impossible to represent zero.**
 - Represent zero as $1 \times \beta^{e_{min}-1}$ - requires extra bit in the exponent
- **Rounding Error**
 - “Units in the last place” eg: 0.0314159 compared to 0.0314 has ulp error of 0.159
 - If calculation is the nearest floating point number to the result, it will still be as much as 1/2 ulp in error
 - Relative error corresponding to 1/2 ulp can vary by a factor of β “wobble”. Written in terms of ϵ
 - Maths \implies **Relative error is always bounded by $\epsilon = (\beta/2)\beta^{-p}$**
 - Fixed relative error \implies ulp can vary by a factor of β . Vice versa
 - Larger $\beta \implies$ larger errors
- **Guard Digits**
 - In subtraction: Could compute exact difference and then round; this is expensive

-
- Keep fixed number of digits but shift operand right; discard precision. Lead to relative error up to $\beta - 1$
 - Guard digit: Add extra digits before truncating. Leads to relative error of less than 2ϵ . This also applies to addition
 - **Catastrophic Cancellation** - Operands are subject to rounding errors - multiplication
 - **Benign Cancellation** - Subtractions. Error $< 2\epsilon$
 - Rearrange formula to avoid catastrophic cancellation
 - Historical interest only - speculation on why IBM used $\beta = 16$ for the system/370 - increased range? Avoids shifting
 - Precision: IEEE defines extended precision (a lower bound only)
 - Discussion of the IEEE standard for operations (TODO: Go over in more detail)
 - NaN allow continuing with underflow and Infinity with overflow
 - “Incidentally, some people think that the solution to such anomalies is never to compare floating-point numbers for equality but instead to consider them equal if they are within some error bound E . This is hardly a cure all, because it raises as many questions as it answers.” - On equality of floating point numbers

8 Compositing Digital Images[7]

Peter and Duff’s classic paper “Compositing Digital Images” lays the foundation for digital compositing today. By providing an “alpha channel,” images of arbitrary shapes and images with soft edges or sub-pixel coverage information can be overlaid digitally, allowing separate objects to be rasterized separately without a loss in quality.

Pixels in digital images are usually represented as 3-tuples containing (red component, green component, blue component). Nominally these values are in the [0-1] range. In the Porter-Duff paper, pixels are stored as (R, G, B, α) 4-tuples, where alpha is the fractional coverage of each pixel. If the image only covers half of a given pixel, for example, its alpha value would be 0.5.

To improve compositing performance, albeit at a possible loss of precision in some implementations, the red, green and blue channels are premultiplied by the alpha channel. This also simplifies the resulting arithmetic by having the colour channels and alpha channels use the same compositing equations.

Several binary compositing operations are defined:

- over
- in
- out
- atop
- xor
- plus

The paper further provides some additional operations for implementing fades and dissolves, as well as for changing the opacity of individual elements in a scene.

The method outlined in this paper is still the standard system for compositing and is implemented almost exactly by modern graphics APIs such as `OpenGL`. It is all but guaranteed that this is the method we will be using for compositing document elements in our project.

9 Bresenham’s Algorithm: Algorithm for computer control of a digital plotter[8]

Bresenham’s line drawing algorithm is a fast, high quality line rasterization algorithm which is still the basis for most (aliased) line drawing today. The paper, while originally written to describe how to control a particular plotter, is uniquely suited to rasterizing lines for display on a pixel grid.

Lines drawn with Bresenham’s algorithm must begin and end at integer pixel coordinates, though one can round or truncate the fractional part. In order to avoid multiplication or division in the algorithm’s inner loop,

The algorithm works by scanning along the long axis of the line, moving along the short axis when the error along that axis exceeds 0.5px. Because error accumulates linearly, this can be achieved by simply adding the per-pixel error (equal to (short axis/long axis)) until it exceeds 0.5, then incrementing the position along the short axis and subtracting 1 from the error accumulator.

As this requires nothing but addition, it is very fast, particularly on the older CPUs used in Bresenham’s time. Modern graphics systems will often use Wu’s line-drawing algorithm instead, as it produces antialiased lines, taking sub-pixel coverage into account. Bresenham himself extended this algorithm to produce Bresenham’s circle algorithm. The principles behind the algorithm have also been used to rasterize other shapes, including Bézier curves.

10 Quad Trees: A Data Structure for Retrieval on Composite Keys[9]

This paper introduces the “quadtree” spatial data structure. The quadtree structure is a search tree in which every node has four children representing the north-east, north-west, south-east and south-west quadrants of its space.

11 Xr: Cross-device Rendering for Vector Graphics[10]

Xr (now known as Cairo) is an implementation of the PDF v1.4 rendering model, independent of the PDF or PostScript file formats, and is now widely used as a rendering API. In this paper, Worth and Packard describe the PDF v1.4 rendering model, and their PostScript-derived API for it.

The PDF v1.4 rendering model is based on the original PostScript model, based around a set of *paths* (and other objects, such as raster images) each made up of lines and Bézier curves, which are transformed by the “Current Transformation Matrix.” Paths can be *filled* in a number of ways, allowing for different handling of self-intersecting paths, or can have their outlines *stroked*. Furthermore, paths can be painted with RGB colours and/or patterns derived from either previously rendered objects or external raster images. PDF v1.4 extends this to provide, amongst other features, support for layering paths and objects using Porter-Duff compositing[7], giving each painted path the option of having an α value and a choice of any of the Porter-Duff compositing methods.

The Cairo library approximates the rendering of some objects (particularly curved objects such as splines) with a set of polygons. An `XrSetTolerance` function allows the user of the library to set an upper bound on the approximation error in fractions of device pixels, providing a trade-off between rendering quality and performance. The library developers found that setting the tolerance to greater than 0.1 device pixels resulted in errors visible to the user.

12 Glitz: Hardware Accelerated Image Compositing using OpenGL[11]

This paper describes the implementation of an OpenGL based rendering backend for the Cairo library.

The paper describes how OpenGL’s Porter-Duff compositing is easily suited to the Cairo/PDF v1.4 rendering model. Similarly, traditional OpenGL (pre-version 3.0 core) support a matrix stack of the same form as Cairo.

The “Glitz” backend will emulate support for tiled, non-power-of-two patterns/textures if the hardware does not support it.

Glitz can render both triangles and trapezoids (which are formed from pairs of triangles). However, it cannot guarantee that the rasterization is pixel-precise, as OpenGL does not provide this consistently.

Glitz also supports multi-sample anti-aliasing, convolution filters for raster image reads (implemented with shaders).

Performance was much improved over the software rasterization and over XRender accelerated rendering on all except nVidia hardware. However, nVidia’s XRender implementation did slow down significantly when some transformations were applied.

13 Boost Multiprecision Library[12]

- “The Multiprecision Library provides integer, rational and floating-point types in C++ that have more range and precision than C++’s ordinary built-in types.”
- Specify number of digits for precision as a template argument.
- Precision is fixed... **possible approach to project:** Use `boost::mpf_float<N>` and increase N as more precision is required?

14 A CMOS Floating Point Unit[13]

The paper describes the implementation of a FPU for PowerPC using a particular Hewlett Packard process (HP14B 0.5 μ m, 3M, 3.3V). It implements a “subset of the most commonly used double precision floating point instructions”. The unimplemented operations are compiled for the CPU.

The paper gives a description of the architecture and design methods. This appears to be an entry to a student design competition.

Standard is IEEE 754, but the multiplier tree is a 64-bit tree instead of a 54 bit tree. “The primary reason for implementing a larger tree is for future additions of SIMD [Single Instruction Multiple Data (?)] instructions similar to Intel’s MMX and Sun’s VIS instructions”.

HSPICE simulations used to determine transistor sizing.

Paper has a block diagram that sort of vaguely makes sense to me. The rest requires more background knowledge.

15 Simply FPU[14]

This is a webpage at one degree of separation from wikipedia.

It talks about FPU internals, but mostly focuses on the instruction sets. It includes FPU assembly code examples (!)

It is probably not that useful, I don't think we'll end up writing FPU assembly?

FPU's typically have 80 bit registers so they can support REAL4, REAL8 and REAL10 (single, double, extended precision).

16 Floating Point Package User's Guide[15]

This is a technical report describing floating point VHDL packages <http://www.vhdl.org/fphdl/vhdl.html>

In theory I know VHDL (cough) so I am interested in looking at this further to see how FPU hardware works. It might be getting a bit sidetracked from the "document formats" scope though.

The report does talk briefly about the IEEE standard and normalised / denormalised numbers as well.

See also: Java Optimized Processor[16] (it has a VHDL implementation of a FPU).

17 Low-Cost Microarchitectural Support for Improved Floating-Point Accuracy[17]

Mentions how GPUs offer very good floating point performance but only for single precision floats.

Has a diagram of a Floating Point adder.

Talks about some magical technique called "Native-pair Arithmetic" that somehow makes 32-bit floating point accuracy "competitive" with 64-bit floating point numbers.

18 Accurate Floating Point Arithmetic through Hardware Error-Free Transformations[18]

From the abstract: "This paper presents a hardware approach to performing accurate floating point addition and multiplication using the idea of error-free transformations. Specialized iterative algorithms are implemented for computing arbitrarily accurate sums and dot products."

The references for this look useful.

It also mentions VHDL.

So whenever hardware papers come up, VHDL gets involved... I guess it's time to try and work out how to use the Opensource VHDL implementations.

References

- [1] Adobe Systems Incorporated. *PostScript Language Reference*. Addison-Wesley Publishing Company, 3rd edition, 1985 - 1999.
- [2] Adobe Systems Incorporated. *PDF Reference*. Adobe Systems Incorporated, 6th edition, 2006.
- [3] Michael A. Wan-Lee Cheng. Portable document format (pdf) – finally, a universal document exchange technology. *Journal of Technology Studies*, 28(1):59 – 63, 2002.
- [4] Brian Hayes. Pixels or perish. *American Scientist*, 100(2):106 – 111, 2012.
- [5] David G. Barnes, Michail Vidiassov, Bernhard Ruthensteiner, Christopher J. Fluke, Michelle R. Quayle, and Colin R. McHenry. Embedding and publishing interactive, 3-dimensional, scientific figures in portable document format (pdf) files. *PLoS ONE*, 8(9):1 – 15, 2013.
- [6] David Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Comput. Surv.*, 23(1):5–48, March 1991.
- [7] Thomas Porter and Tom Duff. Compositing digital images. In *ACM Siggraph Computer Graphics*, volume 18, pages 253–259. ACM, 1984.
- [8] Jack E Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems journal*, 4(1):25–30, 1965.
- [9] Raphael A. Finkel and Jon Louis Bentley. Quad trees a data structure for retrieval on composite keys. *Acta informatica*, 4(1):1–9, 1974.
- [10] Carl Worth and Keith Packard. Xr: Cross-device rendering for vector graphics. In *Linux Symposium*, page 480, 2003.
- [11] Peter Nilsson and David Reveman. Glitz: Hardware accelerated image compositing using OpenGL. In *USENIX Annual Technical Conference, FREENIX Track*, pages 29–40, 2004.
- [12] John Maddock and Christopher Kormanyos. Boost multiprecision library. http://www.boost.org/doc/libs/1_53_0/libs/multiprecision/doc/html/boost_multiprecision/.
- [13] Michael J. Kelley, Matthew A. Postiff, Advisor Richard, and B. Brown. A cmos floating point unit, 1997.
- [14] Raymond Filiatreault. Simply fpu. <http://www.website.masmforum.com/tutorials/fptute/index.html>, 2003.
- [15] David Bishop. Floating point package user’s guide. *EDA Industry Working Groups*, 2008. Technical Report.
- [16] jop devel. Java optimized processor. <https://github.com/jop-devel/jop>.
- [17] William R. Dieter, Akil Kaveti, and Henry G. Dietz. Low-cost microarchitectural support for improved floating-point accuracy. *IEEE Comput. Archit. Lett.*, 6(1):13–16, January 2007.
- [18] Edin Kadric, Paul Gurniak, and André DeHon. Accurate parallel floating-point accumulation. In *Computer Arithmetic (ARITH), 2013 21st IEEE Symposium on*, pages 153–162. IEEE, 2013.