

Literature Review

David Gow

May 22, 2014

1 Introduction

Since mankind first climbed down from the trees, it is our ability to communicate that has made us unique. Once ideas could be passed from person to person, it made sense to have a permanent record of them; one which could be passed on from person to person without them ever meeting.

And thus the document was born.

Traditionally, documents have been static: just marks on paper, but with the advent of computers many more possibilities open up.

2 Rendering

Computer graphics comes in two forms: bit-mapped (or raster) graphics, which is defined by an array of pixel colours; and *vector* graphics, defined by mathematical descriptions of objects. Bit-mapped graphics are well suited to photographs and match how cameras, printers and monitors work.




Figure 1: A circle as a vector image and a 32×32 pixel raster image

However, bitmap devices do not handle zooming beyond their “native” resolution — the resolution where one document pixel maps to one display pixel —, exhibiting an artefact called pixelation where the pixel structure becomes evident. Attempts to use interpolation to hide this effect are never entirely successful, and sharp edges, such as those found in text and diagrams, are particularly affected.

Vector graphics avoid many of these problems: the representation is independent of the output resolution, and rather an abstract description of what it is being rendered, typically as a combination of simple geometric shapes like lines, arcs and glyphs.

As existing displays (and printers) are bit-mapped devices, vector documents must be *rasterized* into a bitmap at a given resolution. This bitmap is then displayed or printed. The resulting bitmap is then an approximation of the vector image at that resolution.

This project will be based around vector graphics, as these properties make it more suited to experimenting with zoom quality.

2.1 Rasterizing Vector Graphics

Before a vector document can be rasterized, the co-ordinates of any shapes must be transformed into *screen space* or *viewport space*[1]. On a typical display, many of these screen-space coordinates require very little precision or range. However, the co-ordinate transform must take care to ensure that precision is preserved during this transform.

After this transformation, the image is decomposed into its separate shapes, which are rasterized and then composited together. Most graphics formats support Porter-Duff compositing[2]. Porter-Duff compositing gives each element (typically a pixel) a “coverage” value, denoted α which represents the contribution of that element to the final scene. Completely transparent elements would have an α value of 0, and completely opaque elements have an α of 1. This permits arbitrary shapes to be layered over one another in the raster domain, while retaining soft-edges.

The rasterization process may then proceed on one object (or shape) at a time. There are special algorithms for rasterizing different shapes.

Line Segment Straight lines between two points are easily rasterized using Bresenham’s algorithm[3]. Bresenham’s algorithm draws a pixel for every point along the *long* axis of the line, moving along the short axis when the error exceeds $\frac{1}{2}$ a pixel.

Bresenham’s algorithm only operates on lines whose endpoints lie on integer pixel coordinates. Due to this, line “clipping” may be performed to find endpoints of the line segment such that the entire line will be on-screen. However, if line clipping is performed naïvely without also setting the error accumulator correctly, the line’s slope will be altered slightly, becoming dependent on the viewport.

Bézier Curve A Bézier curve is a smooth (i.e. infinitely differentiable) curve between two points, represented by a Bernstein polynomial. The coefficients of this Bernstein polynomial are known as the “control points.”

Line Segments are a first-order Bézier curve.

Bézier Spline A spline of order n is a C^{n-1} smooth continuous piecewise function composed of polynomials of degree $\leq n$. In a Bézier spline, these polynomials are Bernstein polynomials, hence the spline is a curve made by joining Bézier curves end-to-end (in a manner which preserves some level of smoothness).

Many vector graphics formats call Bézier splines of a given order (typically quadratic or cubic) “paths” and treat them as the fundamental type from which shapes are formed.

While traditionally, rasterization was done entirely in software, modern computers and mobile devices have hardware support for rasterizing lines and triangles designed for use rendering 3D scenes. This hardware is usually programmed with an API like `OpenGL`[4].

More complex shapes like Bézier curves can be rendered by combining the use of bitmapped textures (possibly using signed-distance fields[5][6][7]) with polygons approximating the curve’s shape[8][9].

Indeed, there are several implementations of entire vector graphics systems using `OpenGL`: `OpenVG`[10] on top of `OpenGL ES`[11]; the `Cairo`[12] library, based around the `PostScript/PDF` rendering model, has the “Glitz” `OpenGL` backend[13] and the `SVG/PostScript GPU` renderer by `nVidia`[14] as an `OpenGL` extension[15].

3 Numeric formats

On modern computer architectures, there are two basic number formats supported: fixed-width integers and *floating-point* numbers. Typically, computers natively support integers of up to 64 bits, capable of representing all integers between 0 and $2^{64} - 1$, inclusive¹.

By introducing a fractional component (analogous to a decimal point), we can convert integers to *fixed-point* numbers, which have a more limited range, but a fixed, greater precision. For example, a number in 4.4 fixed-point format would have four bits representing the integer component, and four bits representing the fractional component:

$$\underbrace{0101}_{\text{integer component}} . \underbrace{1100}_{\text{fractional component}} = 5.75 \quad (1)$$

Floating-point numbers[16] are the binary equivalent of scientific notation: each number consisting of an exponent (e) and a mantissa (m) such that a number is given by

$$n = 2^e \times m \quad (2)$$

The IEEE 754 standard[17] defines several floating-point data types which are used² by most computer systems. The standard defines 32-bit (8-bit exponent, 23-bit mantissa, 1 sign bit) and 64-bit (11-bit exponent, 53-bit mantissa, 1 sign bit) formats³, which can store approximately 7 and 15 decimal digits of precision respectively.

Floating-point numbers behave quite differently to integers or fixed-point numbers, as the representable numbers are not evenly distributed. Large numbers are stored to a lesser precision than numbers close to zero. This can present problems in documents when zooming in on objects far from the origin.

¹Most machines also support *signed* integers, which have the same cardinality as their *unsigned* counterparts, but which represent integers in the range $[-(2^{63}), 2^{63} - 1]$

²Many systems’ implement the IEEE 754 standard’s storage formats, but do not implement arithmetic operations in accordance with this standard.

³The 2008 revision to this standard[18] adds some additional formats, but is less widely supported in hardware.

IEEE floating-point has some interesting features as well, including values for negative zero, positive and negative infinity, the “Not a Number” (NaN) value and *denormal* values, which trade precision for range when dealing with very small numbers. Indeed, with these values, IEEE 754 floating-point equality does not form an equivalence relation, which can cause issues when not considered carefully.[19]

There also exist formats for storing numbers with arbitrary precision and/or range. Some programming languages support “big integer”[20] types which can represent any integer that can fit in the system’s memory. Similarly, there are arbitrary-precision floating-point data types[21][22] which can represent any number of the form

$$\frac{n}{2^d} \quad n, d \in \mathbb{Z} \quad (3)$$

These types are typically built from several native data types such as integers and floats, paired with custom routines implementing arithmetic primitives.[23] These, therefore, are likely slower than the native types they are built on.

While traditionally, GPUs have supported some approximation of IEEE 754’s 32-bit floats, modern graphics processors also support 16-bit[24] and 64-bit[25] IEEE floats. Note, however, that some parts of the GPU are only able to use some formats, so precision will likely be truncated at some point before display. Higher precision numeric types can be implemented or used on the GPU, but are slow.[26]

Pairs of integers ($a \in \mathbb{Z}, b \in \mathbb{Z} \setminus 0$) can be used to represent rationals. This allows values such as $\frac{1}{3}$ to be represented exactly, whereas in fixed or floating-point formats, this would have a recurring representation:

$$\underbrace{0}_{\text{integer part}} . \underbrace{01}_{\text{recurring part}} 01 01 01 \dots \quad (4)$$

Whereas with a rational type, this is simply $\frac{1}{3}$. Rationals do not have a unique representation for each value, typically the reduced fraction is used as a characteristic element.

4 Document Formats

Most existing document formats — such as the venerable PostScript and PDF — are, however, designed to imitate existing paper documents, largely to allow for easy printing. In order to truly take advantage of the possibilities operating in the digital domain opens up to us, we must look to new formats.

Formats such as HTML allow for a greater scope of interactivity and for a more data-driven model, allowing the content of the document to be explored in ways that perhaps the author had not anticipated.[27] However, these data-driven formats typically do not support fixed layouts, and the display differs from renderer to renderer.

4.1 A Taxonomy of Document formats

The process of creating and displaying a document is a rather universal one (4.1), though different document formats approach it slightly differently. A document

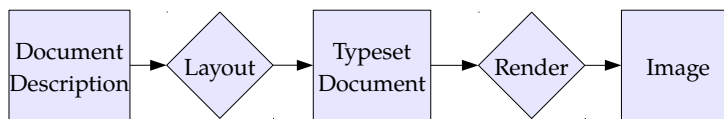


Figure 2: The lifecycle of a document

often begins as raw content: text and images (be they raster or vector) and it must end up as a set of photons flying towards the reader’s eyes.

There are two fundamental stages by which all documents — digital or otherwise — are produced and displayed: *layout* and *rendering*. The *layout* stage is where the positions and sizes of text and other graphics are determined. The text will be *flowed* around graphics, the positions of individual glyphs will be placed, ensuring that there is no undesired overlap and that everything will fit on the page or screen.

The *display* stage actually produces the final output, whether as ink on paper or pixels on a computer monitor. Each graphical element is rasterized and composited into a single image of the target resolution.

Different document formats cover documents in different stages of this project. Bitmapped images, for example, would represent the output of the final stage of the process, whereas markup languages typically specify a document which has not yet been processed, ready for the layout stage.

Furthermore, some document formats treat the document as a program, written in a (usually turing complete) document language with instructions which emit shapes to be displayed. These shapes are either displayed immediately, as in PostScript, or stored in another file, such as with $\text{T}_{\text{E}}\text{X}$ or $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$, which emit a DVI file. Most other forms of document use a *Document Object Model*, being a list or tree of objects to be rendered. DVI, PDF, HTML⁴ and SVG[28]. Of these, only HTML and $\text{T}_{\text{E}}\text{X}$ typically store documents in pre-layout stages, whereas even turing complete document formats such as PostScript typically encode documents which already have their elements placed.

$\text{T}_{\text{E}}\text{X}$ and $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ Donald Knuth’s typesetting language $\text{T}_{\text{E}}\text{X}$ is one of the older computer typesetting systems, originally conceived in 1977[29]. It implements a turing-complete language and is human-readable and writable, and is still popular due to its excellent support for typesetting mathematics. $\text{T}_{\text{E}}\text{X}$ only implements the “layout” stage of document display, and produces a typeset file, traditionally in DVI format, though modern implementations will often target PDF instead.

This document was prepared in $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X} 2_{\epsilon}$.

DVI $\text{T}_{\text{E}}\text{X}$ traditionally outputs to the DVI (“DeVice Independent”) format: a binary format which consists of a simple stack machine with instructions for drawing glyphs and curves[30].

⁴Some of these formats — most notably HTML — implement a scripting language such as JavaScript, which permit the DOM to be modified while the document is being viewed.

A DVI file is a representation of a document which has been typeset, and DVI viewers will rasterize this for display or printing, or convert it to another similar format like PostScript to be rasterized.

HTML The Hypertext Markup Language (HTML)[31] is the widely used document format which underpins the world wide web. In order for web pages to adapt appropriately to different devices, the HTML format simply defined semantic parts of a document, such as headings, phrases requiring emphasis, references to images or links to other pages, leaving the *layout* up to the browser, which would also rasterize the final document.

The HTML format has changed significantly since its introduction, and most of the layout and styling is now controlled by a set of style sheets in the CSS[32] format.

PostScript Much like DVI, PostScript[33] is a stack-based format for drawing vector graphics, though unlike DVI (but like \TeX), PostScript is text-based and turing complete. PostScript was traditionally run on a control board in laser printers, rasterizing pages at high resolution to be printed, though PostScript interpreters for desktop systems also exist, and are often used with printers which do not support PostScript natively.[34]

PostScript programs typically embody documents which have been typeset, though as a turing-complete language, some layout can be performed by the document.

PDF Adobe’s Portable Document Format (PDF)[35] takes the PostScript rendering model, but does not implement a turing-complete language. Later versions of PDF also extend the PostScript rendering model to support translucent regions via Porter-Duff compositing[2].

PDF documents represent a particular layout, and must be rasterized before display.

SVG Scalable Vector Graphics (SVG) is a vector graphics document format[?] which uses the Document Object Model. It consists of a tree of matrix transforms, with objects such as vector paths (made up of Bézier curves) and text at the leaves.

4.2 Precision in Document Formats

Existing document formats — typically due to having been designed for documents printed on paper, which of course has limited size and resolution — use numeric types which can only represent a fixed range and precision. While this works fine with printed pages, users reading documents on computer screens using programs with “zoom” functionality are prevented from working beyond a limited scale factor, lest artefacts appear due to issues with numeric precision.

\TeX uses a 14.16 bit fixed point type (implemented as a 32-bit integer type, with one sign bit and one bit used to detect overflow)[36]. This can represent values in the range $[-(2^{14}), 2^{14} - 1]$ with 16 binary digits of fractional precision.

The DVI files \TeX produces may use “up to” 32-bit signed integers[30] to specify the document, but there is no requirement that implementations support the full 32-bit type. It would be permissible, for example, to have a DVI viewer

support only 24-bit signed integers, though many files which require greater range may fail to render correctly.

PostScript[33] supports two different numeric types: *integers* and *reals*, both of which are specified as strings. The interpreter’s representation of numbers is not exposed, though the representation of integers can be divined by a program by the use of bitwise operations. The PostScript specification lists some “typical limits” of numeric types, though the exact limits may differ from implementation to implementation. Integers typically must fall in the range $[-2^{31}, 2^{31} - 1]$, and reals are listed to have largest and smallest values of $\pm 10^{38}$, values closest to 0 of $\pm 10^{-38}$ and approximately 8 decimal digits of precision, derived from the IEEE 754 single-precision floating-point specification.

Similarly, the PDF specification[35] stores *integers* and *reals* as strings, though in a more restricted format than PostScript. The PDF specification gives limits for the internal representation of values. Integer limits have not changed from the PostScript specification, but numbers representable with the *real* type have been specified differently: the largest representable values are $\pm 3.403 \times 10^{38}$, the smallest non-zero representable values are ⁵ $\pm 1.175 \times 10^{-38}$ with approximately 5 decimal digits of precision *in the fractional part*. Adobe’s implementation of PDF uses both IEEE 754 single precision floating-point numbers and (for some calculations, and in previous versions) 16.16 bit fixed-point values.

The SVG specification[?] specifies numbers as strings with a decimal representation of the number. It is stated that a “Conforming SVG Viewer” must have “all visual rendering accurate to within one device pixel to the mathematically correct result at the initial 1:1 zoom ratio” and that “it is suggested that viewers attempt to keep a high degree of accuracy when zooming.” A “Conforming High-Quality SVG Viewer” must use “double-precision floating point⁶” for computations involving coordinate system transformations.

5 Quadtrees

When viewing or processing a small part of a large document, it may be helpful to only process — or *cull* — parts of the document which are not on-screen.

The quadtree[37] is a data structure — one of a family of *spatial* data structures — which recursively breaks down space into smaller subregions which can be processed independently. Points (or other objects) are added to a single node, which if certain criteria are met — typically the number of points in a node exceeding a maximum, though in our case likely the level of precision required exceeding that supported by the data type in use — is split into four equal-sized subregions, and points attached to the region which contains them.

In this project, we will be experimenting with a form of quadtree in which each node has its own independent coordinate system, allowing us to store some spatial information⁷ within the quadtree structure, eliminating redundancy in the coordinates of nearby objects.

Other spatial data structures exist, such as the KD-tree[38], which partitions the space on any axis-aligned line; or the BSP tree[39], which splits along an

⁵The PDF specification mistakenly leaves out the negative in the exponent here.

⁶Presumably the 64-bit IEEE 754 “double” type.

⁷One bit per-coordinate, per-level of the quadtree

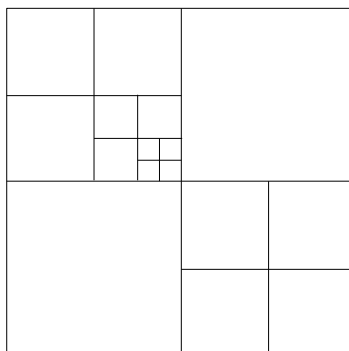


Figure 3: A simple quadtree.

arbitrary line which need not be axis aligned. We believe, however, that the simpler conversion from binary coordinates to the quadtree’s binary split make it a better avenue for initial research to explore.

References

- [1] James Blinn. A trip down the graphics pipeline: Grandpa, what does “viewport” mean? *Computer Graphics and Applications, IEEE*, 12(1):83–87, Jan 1992.
- [2] Thomas Porter and Tom Duff. Compositing digital images. In *ACM SIGGRAPH Computer Graphics*, volume 18, pages 253–259. ACM, 1984.
- [3] Jack E Bresenham. Algorithm for computer control of a digital plotter. *IBM Systems journal*, 4(1):25–30, 1965.
- [4] Mark Segal, Kurt Akely, and Jon Leech. *The OpenGL® Graphics System: A Specification*. The Kronos Group, Inc, 2014.
- [5] F Leymarie and Martin D Levine. Fast raster scan distance propagation on the discrete rectangular lattice. *CVGIP: Image Understanding*, 55(1):84–94, 1992.
- [6] Sarah F Frisken, Ronald N Perry, Alyn P Rockwood, and Thouis R Jones. Adaptively sampled distance fields: a general representation of shape for computer graphics. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 249–254. ACM Press/Addison-Wesley Publishing Co., 2000.
- [7] Chris Green. Improved alpha-tested magnification for vector textures and special effects. In *ACM SIGGRAPH 2007 courses*, pages 9–18. ACM, 2007.
- [8] Charles Loop and Jim Blinn. Resolution independent curve rendering using programmable graphics hardware. *ACM Transactions on Graphics (TOG)*, 24(3):1000–1009, 2005.

- [9] Charles Loop and Jim Blinn. Rendering vector art on the gpu. *GPU gems*, 3:543–562, 2007.
- [10] Mathieu Robart. OpenVG paint subsystem over OpenGL ES shaders. In *Consumer Electronics, 2009. ICCE'09. Digest of Technical Papers International Conference on*, pages 1–2. IEEE, 2009.
- [11] Aekyung Oh, Hyunchan Sung, Hwanyong Lee, Kujin Kim, and Nakhoon Baek. Implementation of OpenVG 1.0 using OpenGL ES. In *Proceedings of the 9th international conference on Human computer interaction with mobile devices and services*, pages 326–328. ACM, 2007.
- [12] Carl Worth and Keith Packard. Xr: Cross-device rendering for vector graphics. In *Linux Symposium*, page 480, 2003.
- [13] Peter Nilsson and David Reveman. Glitz: Hardware accelerated image compositing using OpenGL. In *USENIX Annual Technical Conference, FREENIX Track*, pages 29–40, 2004.
- [14] Mark J Kilgard and Jeff Bolz. GPU-accelerated path rendering. *ACM Transactions on Graphics (TOG)*, 31(6):172, 2012.
- [15] Mark J Kilgard. Programming with NV path rendering: An annex to the SIGGRAPH paper GPU-accelerated path rendering. *heart*, 300:300.
- [16] David Goldberg. The design of floating-point data types. *ACM Lett. Program. Lang. Syst.*, 1(2):138–151, June 1992.
- [17] IEEE standard for binary floating-point arithmetic. *ANSI/IEEE Std 754-1985*, 1985.
- [18] IEEE standard for floating-point arithmetic. *IEEE Std 754-2008*, pages 1–70, Aug 2008.
- [19] David Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Comput. Surv.*, 23(1):5–48, March 1991.
- [20] Oracle Corporation. java.math.BigInteger. <http://docs.oracle.com/javase/6/docs/api/java/math/BigInteger.html>.
- [21] Oracle Corporation. java.math.BigDecimal. <http://docs.oracle.com/javase/7/docs/api/java/math/BigDecimal.html>.
- [22] John Maddock and Christopher Kormanyos. Boost multiprecision library. http://www.boost.org/doc/libs/1_53_0/libs/multiprecision/doc/html/boost_multiprecision/.
- [23] D.M. Priest. Algorithms for arbitrary precision floating point arithmetic. In *Computer Arithmetic, 1991. Proceedings., 10th IEEE Symposium on*, pages 132–143, Jun 1991.
- [24] Pat Brown. NV_half_float. http://www.opengl.org/registry/specs/NV/half_float.txt, 2002.

- [25] Pat Brown, Barthold Lichtenbelt, Bill Licea-Kane, Bruce Merry, Chris Dodd, Eric Werness, Graham Sellers, Greg Roth, Jeff Bolz, Nick Haemel, Pierre Boudier, and Piers Daniell. ARB_gpu_shader_fp64. http://www.opengl.org/registry/specs/ARB/gpu_shader_fp64.txt, 2010.
- [26] Niall Emmart and Charles Weems. High precision integer multiplication with a graphics processing unit. In *2010 IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum (IPDPSW)*, pages 1–6. IEEE, 2010.
- [27] Brian Hayes. Pixels or perish. *American Scientist*, 100(2):106 – 111, 2012.
- [28] Erik Dahlstóm, Patric Dengler, Anthony Grasso, Chris Lilley, Cameron McCormack, Doug Schepers, Jonathon Watt, Jon Ferraiolo, Fujisawa Jun, and Dean Jackson. Scalable vector graphics (svg) 1.1 (second edition). *WC3 Recommendation*, August 2011.
- [29] Donald Knuth. Preliminary preliminary description of T_EX. [http://www.saildart.org/TEXDR.AFT\[1,DEK\]1](http://www.saildart.org/TEXDR.AFT[1,DEK]1), 1977.
- [30] David Fuchs. The format of T_EX’s DVI files. *TUGBoat*, 3(2), 1982.
- [31] Tim Berners-Lee and Daniel Connolly. Hypertext markup language – 2.0. *Internet RFC 1866*, 1995.
- [32] Bert Bos, Håkon Wium Lie, Chris Lilley, and Jacobs Ian. Cascading style sheets, level 2, CSS2 specification. <http://www.w3.org/TR/1998/REC-CSS2-19980512/>.
- [33] Adobe Systems Incorporated. *PostScript Language Reference*. Addison-Wesley Publishing Company, 3rd edition, 1985 - 1999.
- [34] Artifex Software. Ghostscript, an interpreter for the postscript language and pdf. <http://www.ghostscript.com/>, 1988.
- [35] Adobe Systems Incorporated. *PDF Reference*. Adobe Systems Incorporated, 6th edition, 2006.
- [36] Nelson Beebe. Extending T_EX and METAFONT with floating-point arithmetic. *TUGboat*, 28(3), 2007.
- [37] Raphael A. Finkel and Jon Louis Bentley. Quad trees a data structure for retrieval on composite keys. *Acta informatica*, 4(1):1–9, 1974.
- [38] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, September 1975.
- [39] Henry Fuchs, Zvi M. Kedem, and Bruce F. Naylor. On visible surface generation by a priori tree structures. In *Proceedings of the 7th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH ’80*, pages 124–133, New York, NY, USA, 1980. ACM.