

# Infinite Precision Document Formats

*Author:* Samuel Moore[1]

*Partners:* David Gow[2]

*Supervisor:* Prof Tim French



THE UNIVERSITY OF  
WESTERN AUSTRALIA

*Achieve International Excellence*

May 2, 2014

## Abstract

At the fundamental level, a document is a means to convey information. The limitations on a digital document format therefore restrict the types and quality of information that can be communicated. Whilst modern document formats are now able to include increasingly complex dynamic content, they still suffer from early views of a document as a static page; to be viewed at a fixed scale and position. In this report, we focus on the limitations of modern document formats (including PDF, PostScript, SVG) with regards to the level of detail, or precision at which primitives can be drawn. We propose a research project to investigate whether it is possible to obtain an “infinite precision” document format, capable of including primitives created at an arbitrary level of zoom.

**Move to introduction? But it discusses the Introduction :S**

In Chapter 1 we give an overview of the current state of the research in document formats, and the motivation for implementing “infinite precision” in a document format. We will outline our approach to research in collaboration with David Gow[]. In Chapter 2 we provide more detailed background examining the literature related to rendering, interpreting, and creating document formats, as well as possible techniques for increased and possibly infinite precision. In Chapter ?? gives the current state of our research and the progress towards the goals outlined in Chapter 1. In Chapter 4 we will conclude with a summary of our findings and goals.

**Keywords:** *document formats, precision, floating point, graphics, OpenGL, VHDL, PostScript, PDF, bootstraps*

**TODO:** Make document smaller; currently 16 pages with almost no content; limit is 20 with actual content

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Aim . . . . .	2
1.1.1	Clarification of Terms . . . . .	2
1.2	Methods . . . . .	2
1.3	Software and Hardware Requirements . . . . .	4
1.4	Timeline . . . . .	5
<b>2</b>	<b>Literature Review</b>	<b>6</b>
2.1	Raster and Vector Graphics . . . . .	7
2.2	Rendering Vector Images . . . . .	8
2.2.1	Bezier Curves . . . . .	8
2.2.2	Shapes . . . . .	9
2.2.3	Text . . . . .	10
2.2.4	Shading . . . . .	10
2.2.5	Other Things . . . . .	10
2.2.6	Rendering Vector Graphics on the GPU . . . . .	10
2.3	Document Representations . . . . .	10
2.3.1	Interpreted Model . . . . .	11
2.3.2	Crippled Interpreted Model . . . . .	11
2.3.3	Document Object Model . . . . .	12
2.3.4	Blurring the Line — Javascript . . . . .	12
2.3.5	Why do we still use static PDFs . . . . .	13
2.4	Precision in Modern Document Formats . . . . .	13
2.5	Representation of Numbers . . . . .	14
2.5.1	The IEEE Standard . . . . .	14
2.5.2	Floating Point Number Representations . . . . .	14
2.5.3	Limitations Imposed By CPU . . . . .	15
2.5.4	Limitations Imposed By Graphics APIs and/or GPUs . . . . .	15
2.5.5	Alternate Number Representations . . . . .	15

---

<b>3</b>	<b>Progress Report</b>	<b>17</b>
3.1	Development of Testbed Software . . . . .	17
3.2	Design and Implementation of “Tests” . . . . .	17
3.3	Document Format . . . . .	17
3.4	Floating Point Number Representations . . . . .	18
3.5	Virtual FPU . . . . .	18
3.6	Version Control . . . . .	18
<b>4</b>	<b>Conclusion</b>	<b>19</b>
4.1	Acheived Milestones . . . . .	19
4.2	Areas of further work . . . . .	19
4.3	Witty Conclusion Goes Here . . . . .	19
	<b>References</b>	<b>21</b>

# Chapter 1

## Introduction

Most of this chapter is copy pasted from the project proposal  
<<http://szmoore.net/ipdf/documents/ProjectProposalSam.pdf>>

Early electronic document formats such as PostScript were motivated by a need to print documents onto a paper medium. In the PostScript standard, this led to a model of the document as a program; a series of instructions to be executed by an interpreter which would result in “ink” being placed on “pages” of a fixed size[3]. The ubiquitous Portable Document Format (PDF) standard provides many enhancements to PostScript taking into account desktop publishing requirements[4], but it is still fundamentally based on the same imaging model[5]. This idea of a document as a static “page” has led to limited precision in these and other traditional document formats.

The emergence of the internet, web browsers, XML/HTML, JavaScript and related technologies has seen a revolution in the ways in which information can be presented digitally, and the PDF standard itself has begun to move beyond static text and figures[6, 7]. However, the popular document formats are still designed with the intention of showing information at either a single, fixed level of detail, or a small range of levels.

As most digital display devices are smaller than physical paper medium, all useful viewers are able to “zoom” to a subset of the document. Vector graphics formats including PostScript and PDF support rasterisation at different zoom levels[3, 5], but the use of fixed precision floating point numbers causes problems due to imprecision either far from the origin, or at a high level of detail[8, 9].

We are now seeing a widespread use of mobile computing devices with touch screens, where the display size is typically much smaller than paper pages and traditional computer monitors; it seems that there is much to be gained by breaking free of the restricted precision of traditional document formats.

## 1.1 Aim

In this project, we will explore the state of the art of current document formats including PDF, PostScript, SVG, HTML, and the limitations of each in terms of precision. We will consider designs for a document format allowing graphics primitives at an arbitrary level of zoom with no loss of detail. We will refer to such a document format as “infinite precision”. A viewer and editor will be implemented as a proof of concept; we adopt a low level, ground up approach to designing this viewer so as to not become restricted by any single existing document format.

There are many possible applications for documents in which precision is unlimited. Several areas of use include: visualisation of extremely large or infinite data sets; visualisation of high precision numerical computations; digital artwork; computer aided design; and maps.

### 1.1.1 Clarification of Terms

It may be necessary to clarify what we mean by the terms “infinite precision” and “document formats”. Regarding the latter, we consider a document format to be any representation of visual information which is capable of being stored indefinitely. Regarding the former, we do not propose to be able to contain an infinite amount of information within such a document. The goal is to be able to render a primitive at the same level of detail it is specified by a document format, regardless of how precise this level is. For example, the precision of coordinates of primitives drawn in a graphical document editor will always be limited by the resolution of the display on which they are drawn, but not by the viewer.

## 1.2 Methods

Initial research and software development is being conducted in collaboration with David Gow[2]. Once a simple testbed application has been developed, we will individually explore approaches for introducing arbitrary levels of precision; these approaches will be implemented as alternate versions of the same software. The focus will be on drawing simple primitives (lines, polygons, circles). However, if time permits we will explore adding more complicated primitives (font glyphs, bezier curves, embedded bitmaps).

The process of rendering a document will be considered as a common area of research, whilst individual research will be conducted on means for allowing infinite precision. At this stage we have identified two possible areas for individual research:

1. **Arbitrary Precision real valued numbers** — Sam Moore

We plan to investigate the representation of real values to a high or arbitrary degree

of precision. Such representations would allow for a document to be implemented using a single global coordinate system. However, we would expect a decrease in performance with increased complexity of the data structure used to represent a real value. **Both software and hardware techniques will be explored.** We will also consider the limitations imposed by performing calculations on the GPU or CPU.

Starting points for research in this area are Priest’s 1991 paper, “Algorithms for Arbitrary Precision Floating Point Arithmetic”[10], and Goldberg’s 1992 paper “The design of floating point data types”[9]. **A more recent and comprehensive text book, “Handbook of Floating Point Arithmetic”[11], published in 2010, has also been identified as highly relevant.**

## 2. Local coordinate systems — David Gow [2]

An alternative approach involves segmenting the document into different regions using fixed precision floats to define primitives within each region. A quadtree or similar data structure could be employed to identify and render those regions currently visible in the document viewer. **Say more here?**

We aim to compare these and any additional implementations considered using the following metrics:

1. **Performance vs Number of Primitives**

As it is clearly desirable to include more objects in a document, this is a natural metric for the usefulness of an implementation. We will compare the performance of rendering different implementations, using several “standard” test documents.

2. **Performance vs Visible Primitives**

There will inevitably be an overhead to all primitives in the document, whether drawn or not. As the structure of the document format and rendering algorithms may be designed independently, we will repeat the above tests considering only the number of visible primitives.

3. **Performance vs Zoom Level**

We will also consider the performance of rendering at zoom levels that include primitives on both small and large scales, since these are the cases under which floating point precision causes problems in the PostScript and PDF standards.

4. **Performance whilst translation and scaling**

Whilst changing the view, it is ideal that the document be re-rendered as efficiently as possible, to avoid disorienting and confusing the user. We will therefore compare the speed of rendering as the standard documents are translated or scaled at a constant rate.

5. **Artifacts and Limitations on Precision**

As we are unlikely to achieve truly “infinite” precision, qualitative comparisons of the accuracy of rendering under different implementations should be made.

## 1.3 Software and Hardware Requirements

Due to the relative immaturity and inconsistency of graphics drivers on mobile devices, our proof of concept will be developed for a conventional GNU/Linux desktop or laptop computer using OpenGL. However, the techniques explored could easily be extended to other platforms and libraries.



## 1.4 Timeline

Deadlines enforced by the faculty of Engineering Computing and Mathematics are *italicised*.<sup>1</sup>.

Date	Milestone
17 <sup>th</sup> April	Draft Literature Review completed. <b>This sort of didn't happen...</b>
1 <sup>st</sup> May	Testbed Software (basic document format and viewer) completed and approaches for extending to allow infinite precision identified.
26 <sup>th</sup> May	<i>Progress Report and Revised Literature Review due.</i>
9 <sup>th</sup> June	Demonstrations of limitations of floating point precision in the Testbed software.
1 <sup>st</sup> July	At least one implementation of infinite precision for basic primitives (lines, polygons, curves) completed. Other implementations, advanced features, and areas for more detailed research identified.
1 <sup>st</sup> August	Experiments and comparison of various infinite precision implementations completed.
1 <sup>st</sup> September	Advanced features implemented and tested, work underway on Final Report.
TBA	<i>Conference Abstract and Presentation due.</i>
10 <sup>th</sup> October	<i>Draft of Final Report due.</i>
27 <sup>th</sup> October	<i>Final Report due.</i>

<sup>1</sup>David Gow is being assessed under the 2014 rules for a BEng (Software) Final Year Project, whilst the author is being assessed under the 2014 rules for a BEng (Mechatronics) Final Year Project; deadlines and requirements as shown in Gow's proposal[2] may differ

## Chapter 2

# Literature Review

### 0. Here is a brilliant summary of the sections below

This chapter provides an overview of relevant literature. The areas of interest can be broadly grouped into two largely separate categories; Documents and Number Representations.

The first half of this chapter will be devoted to documents themselves, including: the representation and rendering of low level graphics primitives, how collections of these primitives are represented in document formats, and the various standards for documents in use today.

We will find that although there has been a great deal of research into the rendering, storing, editing, manipulation, and extension of document formats, all popular document standards are content to specify at best a single precision IEEE-754 floating point number representations.

The research on arbitrary precision arithmetic applied to documents is very sparse; however arbitrary precision arithmetic itself is a very active field of research. Therefore, the second half of this chapter will be devoted to considering the IEEE-754 standard, its advantages and limitations, and possible alternative number representations to allow for arbitrary precision arithmetic.

In Chapter ??, we will discuss our findings so far with regards to arbitrary precision arithmetic applied to document formats.

## 2.1 Raster and Vector Graphics

### 1. Here are the fundamentals of graphics (raster and vector, rendering)

At a fundamental level everything that is seen on a display device is represented as either a vector or raster image. These images can be stored as stand alone documents or embedded in a much more complex document format capable of containing many other types of information.

A raster image's structure closely matches its representation as shown on modern display hardware; the image is represented as a grid of filled square "pixels". Each pixel is the same size and contains information describing its colour. This representation is simple and also well suited to storing images as produced by cameras and scanners[?].

The drawback of raster images is that by their very nature there can only be one level of detail. Figures 2.1 and 2.2 attempt to illustrate this by comparing raster images to vector images in a similar way to Worth and Packard[12].

Consider the right side of Figure 2.1. This is a raster image which should be recognisable as an animal defined by fairly sharp edges. Figure 2.2 shows that zooming on the animal's face causes these edges to appear jagged. There is no information in the original image as to what should be displayed at a larger size, so each square shaped pixel is simply increased in size. A blurring effect will probably be visible in most PDF viewers; the software has attempted to make the "edge" appear more realistic using a technique called "antialiasing" which averages neighbouring pixels in the original image in order to generate extra pixels in the scaled image[?].<sup>1</sup>

In contrast, the left sides of Figures 2.1 and 2.2 are a vector image. A vector image contains information about a number of geometric shapes. To display this image on modern display hardware, the coordinates are transformed according to the view and *then* the image is converted into a raster like representation. Whilst the raster image merely appears to contain edges, the vector image actually contains information about these edges, meaning they can be displayed "infinitely sharply" at any level of detail[?] — or they could be if the coordinates are stored with enough precision (see Section ??). Thus, vector images are well suited to high quality digital art<sup>2</sup> and text[?].

Woah, an entire page with only one citation ham fistted in after I had written the rest... and the "actually writing it" phase of the Lit Review is off to a great start.

---

<sup>1</sup>The exact appearance of the images at different zoom levels will depend greatly on the PDF viewer or printer used to display this report. On the author's display using the Atril (1.6.0) viewer, the top images appear to be pixel perfect mirror images at a 100% scale. In the bottom raster image, antialiasing is not applied at zoom levels above 125% and the effect of scaling is quite noticable.

<sup>2</sup>Figure 2.1 is not to be taken as an example of this.

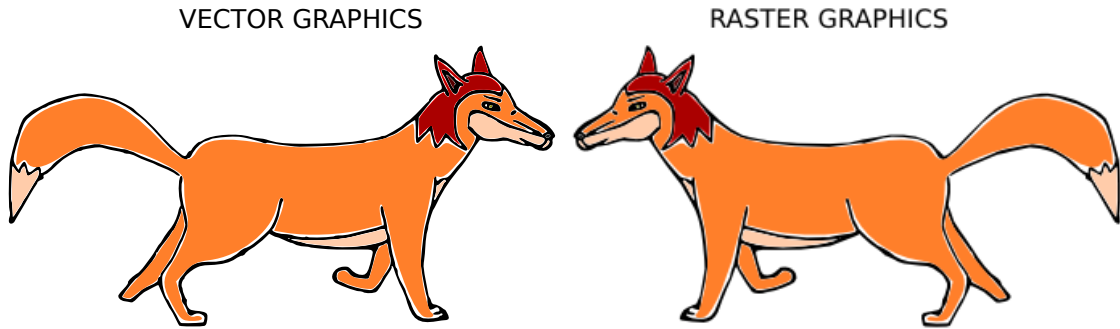


Figure 2.1: Original Vector and Raster Images




Figure 2.2: Scaled Vector and Raster Images

## 2.2 Rendering Vector Images

Throughout Section 2.1 we were careful to refer to “modern” display devices, which are raster based. It is of some historical significance that vector display devices were popular during the 70s and 80s, and so algorithms for drawing a vector image directly without rasterisation exist. An example is the shading of polygons which is somewhat more complicated on a vector display than a raster display[13, 14].

All modern displays of practical interest are raster based. In this section we explore the structure of vector graphics in more detail, and how different primitives are rendered.

After the wall of citationless text in Section 2.1 we should probably redeem ourselves a bit here

### 2.2.1 Bezier Curves

The bezier curve is of vital importance in vector graphics.

### Things this section lacks

- Who came up with them (presumably it was a guy named Bezier)
- Flesh out how they evolved or came into use?
- Naive algorithm
- De Casteljau Algorithm

Recently, Goldman presented an argument that Bezier's could be considered as fractal in nature, a fractal being the fixed point of an iterated function system[15]. Goldman's proof depends upon a modification to the De Casteljau Subdivision algorithm. Whilst we will not go the details of the proof, or attempt comment on its theoretical value, it is interesting to note that Goldman's algorithm is not only worse than the De Casteljau algorithm upon which it was based, but it also performs worse than a naive Bezier rendering algorithm. Figure 2.3 shows our results using implementations of the various algorithms in python.

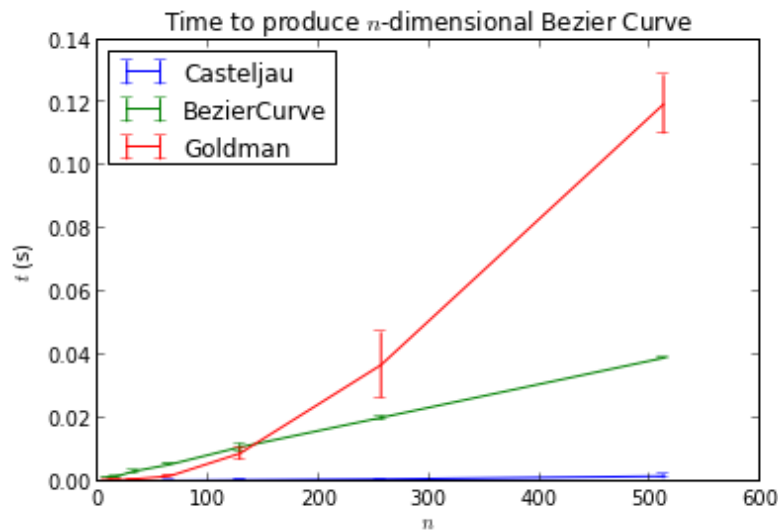


Figure 2.3: Performance of Bezier Subdivision Algorithms

Does the Goldman bit need to be here? Probably NOT. Do I need to check very very carefully that I haven't made a mistake before saying this? YES. Will I have made a mistake? Probably.

### 2.2.2 Shapes

Shapes are just bezier curves joined together.

## Approximating a Circle Using Cubic Beziers

An example of a shape is a circle. We used some algorithm on wikipedia that I'm sure is in Literature somewhere [?] and made a circle. It's in my magical ipython notebook with the De Casteljau stuff.

### 2.2.3 Text

Text is just Bezier Curves, I think we proved out point with the circle, but maybe find some papers to cite[?]

### 2.2.4 Shading

Shading is actually extremely complicated! [13, 14] *Sure, but do we care enough to talk about it? We will run out of space at this rate*

### 2.2.5 Other Things

We don't really care about other things in this report.

*6. Sort of starts here... or at least background does*

### 2.2.6 Rendering Vector Graphics on the GPU

Traditionally, vector graphics have been rasterized by the CPU before being sent to the GPU for drawing[16]. Lots of people would like to change this [12, 17, 18, 16, 19] ... *All of these are things David found except kilgard which I thought I found and then realised David already had it :S*

*2. Here are the ways documents are structured ... we got here eventually*

## 2.3 Document Representations

*The file format can be either human readable<sup>3</sup> or binary<sup>4</sup>. Can also be compressed or not. Here we are interested in how the document is interpreted or traversed in order to produce graphics output.*

---

<sup>3</sup>For some definition of human and some definition of readable

<sup>4</sup>So, our viewer is basically a DOM style but stored in a binary format

### 2.3.1 Interpreted Model

Did I just invent that terminology or did I read it in a paper? Is there actually existing terminology for this that sounds similar enough to “Document Object Model” for me to compare them side by side?

- This model treats a document as the source code program which produces graphics
- Arose from the desire to produce printed documents using computers (which were still limited to text only displays).
- Typed by hand or (later) generated by a GUI program
- PostScript — largely supersceded by PDF on the desktop but still used by printers<sup>5</sup>
- $\text{\TeX}$ — Predates PostScript!  $\text{\LaTeX}$  is being used to create this very document and until now I didn’t even have it here!
  - I don’t really want to go down the path of investigating the billion steps involved in getting  $\text{\LaTeX}$  into an actually viewable format
  - There are interpreters (usually WYSIWYG editors) for  $\text{\LaTeX}$  though
  - Maybe if  $\text{\LaTeX}$  were more popular there would be desktop viewers that converted  $\text{\LaTeX}$  directly into graphics
- Potential for dynamic content, interactivity; dynamic PostScript, enhanced Postscript
- Scientific Computing — Mathematica, Matlab, IPython Notebook — The document and the code that produces it are stored together
- Problems with security — Turing complete, can be exploited easily

### 2.3.2 Crippled Interpreted Model

I’m pretty sure I made that one up

- PDF is PostScript but without the Turing Completeness
- Solves security issues, more efficient

---

<sup>5</sup>Desktop pdf viewers can still cope with PS, but I wonder if a smartphone pdf viewer would implement it?

### 2.3.3 Document Object Model

- DOM = Tree of nodes; node may have attributes, children, data
- XML (SGML) is the standard language used to represent documents in the DOM
- XML is plain text
- SVG is a standard for a vector graphics language conforming to XML (ie: a DOM format)
- CSS style sheets represent more complicated styling on objects in the DOM

### 2.3.4 Blurring the Line — Javascript

- The document is expressed in DOM format using XML/HTML/SVG
- A Javascript program is run which can modify the DOM
- At a high level this may be simply changing attributes of elements dynamically
- For low level control there is canvas2D and even WebGL which gives direct access to OpenGL functions
- Javascript can be used to make a HTML/SVG interactive
  - Overlooking the fact that the SVG standard already allows for interactive elements...
- Javascript is now becoming used even in desktop environments and programs (Windows 8, GNOME 3, Cinnamon, Game Maker Studio) (**shudder**)
- There are also a range of papers about including Javascript in PDF “Pixels or Perish” being the only one we have actually read[6]
  - I have no idea how this works; PDF is based on PostScript... it seems very circular to be using a programming language to modify a document that is modelled on being a (non turing complete) program
  - This is yet more proof that people will converge towards solutions that “work” rather than those that are optimal or elegant
  - I guess it’s too much effort to make HTML look like PDF (or vice versa) so we could phase one out



### 2.3.5 Why do we still use static PDFs

Despite their limitations, we still use static, boring old PDFs. Particularly in scientific communication.

- They are portable; you can write an amazing document in Mathematica/Matlab but it
- Scientific journals would need to adapt to other formats and this is not worth the effort
- No network connection is required to view a PDF (although DRM might change this?)
- All resources are stored in a single file; a website is stored accross many seperate files (call this a “distributed” document format?)
- You can create PDFs easily using desktop processing WYSIWYG editors; WYSIWYG editors for web based documents are worthless due to the more complex content
- Until Javascript becomes part of the PDF standard, including Javascript in PDF documents will not become widespread
- Once you complicate a PDF by adding Javascript, it becomes more complicated to create; it is simply easier to use a series of static figures than to embed a shader in your document. Even for people that know WebGL.

3. Here are the ways document standards specify precision (or don't)

## 2.4 Precision in Modern Document Formats

All the above is very interesting and provides important context, but it is not actually directly related to the problem of infinite precision which we are going to try and solve. Sorry to make you read it all.

- Implementations of PostScript and PDF must by definition restrict themselves to IEEE binary32 “single precision”<sup>6</sup> floating point number representations in order to conform to the standards[3, 5].
- Implementations of SVG are by definition required to use IEEE binary32 as a **minimum**. “High Quality” SVG viewers are required to use at least IEEE binary64.[20]

---

<sup>6</sup>The original IEEE-754 defined single, double and extended precisions; in the revision these were renamed to binary32, binary64 and binary128 to explicitly state the base and number of bits

- Numerical computation packages such as Mathematica and Maple use arbitrary precision floats
  - Mathematica is not open source which is an issue when publishing scientific research (because people who do not fork out money for Mathematica cannot verify results)
  - What about Maple? [11] and [21] both mention it being buggy.
  - Octave and Matlab use fixed precision doubles

The use of IEEE binary32 floats in the PostScript and PDF standards is not surprising if we consider that these documents are oriented towards representing static pages. They don't actually need higher precision to do this; 32 bits is more than sufficient for A4 paper.

4. Here is IEEE-754 which is what these standards use

## 2.5 Representation of Numbers

Although this project has been motivated by a desire for more flexible document formats, the fundamental source of limited precision in vector document formats is the restriction to IEEE floating point numbers for representation of coordinates.

Whilst David Gow will be focusing on structures **and the use of multiple coordinate systems** to represent a document so as to avoid or reduce these limitations[2], the focus of our own research will be **increased precision in the representation of real numbers so as to get away with using a single global coordinate system.**

### 2.5.1 The IEEE Standard

### 2.5.2 Floating Point Number Representations

$$x = (-1)^s \times m \times B^e$$

$B = 2$ , although IEEE also defines decimal representations for  $B = 10$  — these are useful in financial software[22].

**Aside: Are decimal representations for a document format eg: CAD also useful because you can then use metric coordinate systems?**

### Precision

The floats map an infinite set of real numbers onto a discrete set of representations.

Figure: 8 bit “minifloats” (all 255 of them) clearly showing the “precision vs range” issue

The most a result can be rounded in conversion to a floating point number is the units in last place;  $m_N \times B^e$ .

Even though that paper that claims double is the best you will ever need because the error can be as much as the size of a bacterium relative to the distance to the moon[] there are many cases where increased number of bits will not save you.[11]

5. Here are limitations of IEEE-754 floating point numbers on compatible hardware

### 2.5.3 Limitations Imposed By CPU

CPU’s are restricted in their representation of floating point numbers by the IEEE standard.

### 2.5.4 Limitations Imposed By Graphics APIs and/or GPUs

Traditionally algorithms for drawing vector graphics are performed on the CPU; the image is rasterised and then sent to the GPU for rendering[]. Recently there has been a great deal of literature relating to implementation of algorithms such as bezier curve rendering[] or shading[] on the GPU. As it seems the trend is to move towards GPU

6. Here are ways GPU might not be IEEE-754 — This goes \*somewhere\* in here but not sure yet

- Internal representations are GPU dependent and may not match IEEE[23]
- OpenGL standards specify: binary16, binary32, binary64
- OpenVG aims to become a standard API for SVG viewers but the API only uses binary32 and hardware implementations may use less than this internally[18]
- It seems that IEEE has not been entirely successful; although all modern CPUs and GPUs are able to read and write IEEE floating point types, many do not conform to the IEEE standard in how they represent floating point numbers internally.

7. Sod all that, let’s just use an arbitrary precision library (AND THUS WE FINALLY GET TO THE POINT)

### 2.5.5 Alternate Number Representations

They exist[11].

Do it all using MFPR[], she'll be right.

8. Here is a brilliant summary of sections 7- above

Dear reader, thankyou for your persistance in reading this mangled excuse for a Literature Review. Hopefully we have brought together the radically different areas of interest together in some sort of coherant fashion. In the next chapter we will talk about how we have succeeded in rendering a rectangle. It will be fun. I am looking forward to it.

## Chapter 3

# Progress Report

This chapter outlines the current state of our research in relation to the aims outlined in Chapter 1. *It will serve as an explanation for where the Figures in Chapter 2 came from. It will just be a short summary of the implementation details.*

### 3.1 Development of Testbed Software

We wrote a very simple OpenGL 1.1 program to experiment with, and then David Gow converted it to OpenGL 3.1 and I have no idea how it works anymore.

### 3.2 Design and Implementation of “Tests”

- Compile by swapping out `main()` for a tester
- There are tests for doing some of the things in Chapter 1 but most still aren’t written yet.

### 3.3 Document Format

Currently we effectively have a DOM format but with the following non-features:

- Binary file format (non standard; not XML)
- Only rectangles.

## 3.4 Floating Point Number Representations

I have<sup>1</sup> some figures that I would prefer to include in Chapter 2 when I am talking about the papers that inspired them. This section will probably briefly talk about how they were created and just refer back to them.

- `calculatepi.test`
- `typedef` of `Real`

## 3.5 Virtual FPU

Techniques for dealing with FP numbers can be implemented in software (CPU) or on dedicated hardware (FPU). We are able to run FP arithmetic on arbitrary simulations of FPUs created using VHDL. *Hopefully explore this a bit in Chapter 2.*

## 3.6 Version Control

Git is a distributed version control system widely used in the development of open source software<sup>1</sup>. All resources created for or used by this project have been placed in git repositories on several servers. The repositories are publically accessible at <http://git.ucc.asn.au>

---

<sup>1</sup>Ok... “will have”

# Chapter 4

## Conclusion

This report has provided motivation for considering approaches to achieving an infinite level of zoom in a document.

### 4.1 Acheived Milestones

### 4.2 Areas of further work

- Continue looking for relevant literature
- Implement all those tests mentioned in Chapter 1
- **Actually identify the techniques I will use THIS ONE SHOULD BE DONE BEFORE I HAND IN THE LITERATURE REVIEW!**
- Possible Ultimate Goal: Implement (a subset) of SVG and then show an SVG document that we can render but a browser can't
  - This means extending our viewer to be able to read (a subset) SVG
  - Can already read XML, so this shouldn't actually be too bad
    - \* Emphasis on **subset**
    - \* (I've seen the SVG standard; I'm talking about implementing the 18 pages under "Basic Shapes". The other 818 pages can complain to someone who cares.)
  - Suggestion to David that he probably won't like (or read): Make his octree structure specifiable as an SVG extension

### 4.3 Witty Conclusion Goes Here

# References

- [1] Sam Moore. Infinite precision document formats (project proposal). <http://szmoore.net/ipdf/documents/ProjectProposalSam.pdf>, 2014.
- [2] David Gow. Infinite-precision document formats (project proposal). <http://davidgow.net/stuff/ProjectProposal.pdf>, 2014.
- [3] Adobe Systems Incorporated. *PostScript Language Reference*. Addison-Wesley Publishing Company, 3rd edition, 1985 - 1999.
- [4] Michael A. Wan-Lee Cheng. Portable document format (pdf) – finally, a universal document exchange technology. *Journal of Technology Studies*, 28(1):59 – 63, 2002.
- [5] Adobe Systems Incorporated. *PDF Reference*. Adobe Systems Incorporated, 6th edition, 2006.
- [6] Brian Hayes. Pixels or perish. *American Scientist*, 100(2):106 – 111, 2012.
- [7] 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.
- [8] David Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Comput. Surv.*, 23(1):5–48, March 1991.
- [9] David Goldberg. The design of floating-point data types. *ACM Lett. Program. Lang. Syst.*, 1(2):138–151, June 1992.
- [10] D.M. Priest. Algorithms for arbitrary precision floating point arithmetic. In *Computer Arithmetic, 1991. Proceedings., 10th IEEE Symposium on*, pages 132–143, Jun 1991.
- [11] Jean-Michel Muller, Nicolas Brisebarre, Florent de Dinechin, Claude-Pierre Jeannerod, Vincent Lefèvre, Guillaume Melquiond, Nathalie Revol, Damien Stehlé, and Serge Torres. *Handbook of Floating-Point Arithmetic*. Birkhäuser Boston Inc., Cambridge, MA, USA, 2010.



- 
- [12] Carl Worth and Keith Packard. Xr: Cross-device rendering for vector graphics. In *Linux Symposium*, page 480, 2003.
- [13] Kurt E. Brassel and Robin Fegeas. An algorithm for shading of regions on vector display devices. *SIGGRAPH Comput. Graph.*, 13(2):126–133, August 1979.
- [14] J. M. Lane and R. and M. Rarick. An algorithm for filling regions on graphics display devices. *ACM Trans. Graph.*, 2(3):192–196, July 1983.
- [15] Ron Goldman. The fractal nature of bezier curves. The de Casteljau subdivision algorithm is used to show that Bezier curves are also attractors (ie: fractals). A new rendering algorithm is derived for Bezier curves.
- [16] Mark J Kilgard and Jeff Bolz. Gpu-accelerated path rendering. *ACM Transactions on Graphics (TOG)*, 31(6):172, 2012.
- [17] Charles Loop and Jim Blinn. Rendering vector art on the gpu. *GPU gems*, 3:543–562, 2007.
- [18] Daniel Rice and RJ Simpson. Openvg specification, version 1.1. *Khronos Group*, 2008.
- [19] Chris Green. Improved alpha-tested magnification for vector textures and special effects. In *ACM SIGGRAPH 2007 courses*, pages 9–18. ACM, 2007.
- [20] 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.
- [21] Laurent Fousse, Guillaume Hanrot, Vincent Lefèvre, Patrick Pélissier, and Paul Zimmermann. Mpf: A multiple-precision binary floating-point library with correct rounding. *ACM Trans. Math. Softw.*, 33(2), June 2007.
- [22] Ieee standard for floating-point arithmetic. *IEEE Std 754-2008*, pages 1–70, Aug 2008.
- [23] Karl E Hillesland and Anselmo Lastra. Gpu floating-point paranoia. *Proceedings of GP 2004*, 2004.