

INTEGRAL ESTIMATION IN QUANTUM PHYSICS

by
Jane Doe

A dissertation submitted to the faculty of
The University of Utah
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Mathematical Physics

Department of Mathematics
The University of Utah
May 2016

Copyright © Jane Doe 2016

All Rights Reserved

The University of Utah Graduate School

STATEMENT OF DISSERTATION APPROVAL

The dissertation of Jane Doe
has been approved by the following supervisory committee members:

<u>Cornelius L�nczos</u> ,	Chair(s)	<u>17 Feb 2016</u> Date Approved
<u>Hans Bethe</u> ,	Member	<u>17 Feb 2016</u> Date Approved
<u>Niels Bohr</u> ,	Member	<u>17 Feb 2016</u> Date Approved
<u>Max Born</u> ,	Member	<u>17 Feb 2016</u> Date Approved
<u>Paul A. M. Dirac</u> ,	Member	<u>17 Feb 2016</u> Date Approved

by Petrus Marcus Aurelius Featherstone-Hough , Chair/Dean of
the Department/College/School of Mathematics
and by Alice B. Toklas , Dean of The Graduate School.

ABSTRACT

Blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
 blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
 blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
 blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
 blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
 blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
 blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
 blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
 blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
 blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
 blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
 blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
 blah blah blah blah blah blah.

For my parents, Alice and Bob.

CONTENTS

ABSTRACT	iii
LIST OF FIGURES	vii
LIST OF TABLES	viii
NOTATION AND SYMBOLS	ix
TYPESETTING EXPERIMENTS	x
CHAPTERS	1
1. THE FIRST	1
1.1 The first section	1
1.1.1 The first subsection	2
1.1.2 The second subsection	2
1.1.3 The third subsection	2
1.1.3.1 The first subsubsection	2
1.1.3.2 The second subsubsection	3
1.1.3.2.1 The first numbered paragraph	3
1.1.3.2.2 The second numbered paragraph	3
1.2 The second section	3
1.3 The third section	5
1.4 Free software packages	6
1.5 Resizing figures	10
1.6 Summary and conclusions	14
1.7 References	14
2. THE SECOND	17
2.1 References	17
3. THE THIRD	19
3.1 References	19
4. THE FOURTH	21
4.1 More on the topic	21
4.2 Even more on the topic	21
4.3 Summary and conclusions	22
5. EXTENDING T_EX AND METAFONT WITH FLOATING POINT ARITHMETIC	23
5.1 Dedication	26

5.2	Introduction	26
5.3	Arithmetic in T _E X and METAFONT	26
5.4	Historical remarks	26
5.5	Why no floating-point arithmetic?	28
5.6	IEEE 754 binary floating-point standard	29
5.7	IEEE 754R precision and range	29
5.8	Remarks on floating-point arithmetic	30
5.9	Binary versus decimal	30
5.10	Problems with IEEE 754 arithmetic	31
5.11	How decimal arithmetic is different	31
5.12	Software floating-point arithmetic	31
5.13	How much work is needed	33
5.14	Summary	33
5.15	References	33
APPENDICES		37
A. THE FIRST		37
B. THE SECOND		39
C. THE THIRD		41
BINOMIAL NOMENCLATURE INDEX		45
FREE SOFTWARE INDEX		46
TOPIC INDEX		50

LIST OF FIGURES

1.1	The first figure.	9
1.2	The second figure.	9
1.3	The third figure.	9
1.4	The fourth figure (at 50% scale).	11
1.5	The fifth figure (at 75% scale).	11
1.6	The sixth figure (at native size).	11
1.7	The seventh figure (at 125% scale).	11
1.8	The eighth figure (at 175% scale).	11
1.9	The ninth figure (at 50% scale)	12
1.10	The tenth figure (at 75% scale)	12
1.11	Using \LaTeX picture mode	14

LIST OF TABLES

1.1 Lowercase Greek letters.	4
1.2 Uppercase Greek letters.	6

NOTATION AND SYMBOLS

α	fine-structure (dimensionless) constant, approximately $1/137$
α	radiation of doubly-ionized helium ions, He ⁺⁺
β	radiation of electrons
γ	radiation of very high frequency, beyond that of X rays
γ	Euler's constant, approximately 0.577 215 ...
δ	stepsize in numerical integration
$\delta(x)$	Dirac's famous function
ϵ	a tiny number, usually in the context of a limit to zero
$\zeta(x)$	the famous Riemann zeta function
...	...
$\psi(x)$	logarithmic derivative of the gamma function
ω	frequency

TYPESETTING EXPERIMENTS

In this section, we use color in several places. The `\colorbox` command takes two arguments — a named color and text to be in black on a background of that color — and sets the text in a box with a small margin of width `\fboxsep` (set to 3.0pt in this document).

Here, we want a tighter colored box that has a fixed height, and is independent of letter shape. We set the margin to zero inside a group so that the change is purely local, and so that height and depth of the line are not increased over what they would be if the colored box were not used. We prefix a T_EX `\strut` to the user-supplied text, because that command expands to a zero-width box of the height and depth of parentheses, which, in most fonts, delimit the extent of letter shapes.

```
\newcommand {\hilitebox} [1] {\fboxsep = 0pt\colorbox{pink}{\strut #1}}
```

Here is a fragment from the first chapter in another thesis, set in *emphasized text* to distinguish it from the rest of this section:

In light of the known results, the consistency of empirical semivariogram and related estimators is widely considered a settled matter. For example, Lahiri, Lee, and Cressie [?] state:

The simpler and more commonly used nonparametric estimators of the variogram, such as the method of moments estimator of Matheron (1962) and its robustified versions due to Cressie and Hawkins (1980) have many desirable properties like, unbiasedness, consistency, etc. ...

Regarding a kernel estimator of the covariance function, Hall and Patil [?] remarked:

It is not difficult to see that if, as n increases, the points t_i become increasingly dense in each bounded subset of \mathbb{R}^d , then the bandwidth h may be chosen so that $\check{\rho}(t) \rightarrow \rho(t)$ as $n \rightarrow \infty$, for each $t \in \mathbb{R}^d$.

However, in order to be true, such statements would need to be qualified by many assumptions on the random field as well as on the observation locations. We will see in §2.3 that even for well-behaved random fields (e.g., ρ^ -mixing Gaussian random fields), it is not enough to assume that the observation locations become increasingly dense in each bounded subset; a stronger assumption*

must be made to ensure that the observation locations do not become denser in one region too much faster than in others.

The text before the previous paragraph contained two quote environments separated by a line of prose. Here are some more tests of both kinds of L^AT_EX environments for showing text written by someone else.

This is a `quote` environment with one short line, following a fairly short paragraph of prose (in this, and following examples, the text is explicitly colored with a command like `\color{purple}` inside the environment before the text):

```
\begin{quote}
  \color{purple}
  14 March 2016 is  $\pi \approx 3.1416$  day in funny notation.
  \hfill \emph{Web news reports}
\end{quote}
```

14 March 2016 is $\pi \approx 3.1416$ day in funny notation. *Web news reports*

This is a `quote` environment with three short lines, each a separate paragraph, following a fairly short paragraph of prose.

```
\begin{quote}
  \color{forestgreen}
  14 March 2016 is  $\pi \approx 3.1416$  day in funny notation.
  \hfill \emph{Web news reports}

  14 March 2016 is  $\pi \approx 3.1416$  day in funny notation.
  \hfill \emph{Web news reports}

  14 March 2016 is  $\pi \approx 3.1416$  day in funny notation.
  \hfill \emph{Web news reports}
\end{quote}
```

14 March 2016 is $\pi \approx 3.1416$ day in funny notation. *Web news reports*

14 March 2016 is $\pi \approx 3.1416$ day in funny notation. *Web news reports*

14 March 2016 is $\pi \approx 3.1416$ day in funny notation. *Web news reports*

Here is another example, this time with separate colors for each paragraph:

```
\begin{quote}
```

```

\color{darkkhaki}
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
\hfill \emph{Web news reports}

\color{darkmagenta}
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
\hfill \emph{Web news reports}

\color{darkcyan}
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
\hfill \emph{Web news reports}

\color{darkorange}
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
14 March 2016 is $\pi \approx 3.1416$ day in funny notation.
\linebreak
\strut
\hfill \emph{Web news reports}
\end{quote}

14 March 2016 is  $\pi \approx 3.1416$  day in funny notation. Web news reports
14 March 2016 is  $\pi \approx 3.1416$  day in funny notation. Web news reports
14 March 2016 is  $\pi \approx 3.1416$  day in funny notation. Web news reports
14 March 2016 is  $\pi \approx 3.1416$  day in funny notation. 14 March 2016 is  $\pi \approx 3.1416$ 
day in funny notation. 14 March 2016 is  $\pi \approx 3.1416$  day in funny notation.
Web news reports

```

Notice that `quote` paragraphs are *not* indented, but the environment itself *is* indented on the left and right by the value of `\leftmargin` (set to 27.37506pt in this document, which should be identical to 2.5em, where 1em = 11.49739pt).

For debugging purposes, we also have `\leftmarginii` set to 27.37506pt, and we have `\leftmarginiii` set to 24.09003pt.

This is a `quotation` environment with one paragraph, following a fairly short paragraph of prose (notice that the quotation paragraphs *are* indented):

```

\begin{quotation}
\color{blue}
Algebra is concerned with manipulation in
\emph{time}, and geometry is concerned with
\emph{space}. These are two orthogonal aspects
of the world, and they represent two different

```


and on and on and on

Now all following text should be back in double-spaced mode, and just go on and on
and on and on and on and on and on and on and on and on and on and on and on and on
and on and on and on and on and on and on and on and on and on and on and on and on
and on and on and on and on and on and on and on and on and on and on and on and on
and on and on and on and on and on and on and on and on and on and on and on and on
and on and on and on and on and on and on and on and on and on and on and on and on
and on and on and on

Chapter 1

THE FIRST

This is a chapter. Remember that there should *always* be at least of few lines of prose after each sectional heading: failure to do so is a disservice to your readers, and also produces incorrect vertical spacing.

1.1 The first section

In [Figure 1.1](#) on page 9, we have a picture, and the \LaTeX markup to include it looks like this:

```
\begin{figure}[p]
  \centerline{\includegraphics{fig1}}
  \caption{The first figure.}%
  \figlabel{fig1}
\end{figure}
```

The convention in English-language typography is that captions are placed *under figures*, but *above tables*. \LaTeX does not know anything about that practice: it simply sets the caption where its macro occurs.

We intentionally omitted an extension on the filename, so that this document can be processed with `latex` to get an output `.dvi` file, or with `pdflatex` to get an output `.pdf` file. The first case uses the file `fig1.eps`, and the second uses `fig1.pdf` (or any of a number of bitmap file formats, such as `.jpg`, `.png`, `.tiff`, ...). The `distill` or `ps2pdf` commands can be used to convert from *Encapulated PostScript* files to *Portable Document Format* files.

The ImageMagick `convert` tool can transform large numbers of file formats, and by default, it uses file extensions to recognize the conversion needed. For example,

```
% convert infile.jpg outfile.pdf
```


would produce a PDF file from a JPEG bitmap image file. It has more than 200 options that can further control its behavior, but simple conversions rarely need them.

Please remember that with bitmap image formats, the resolution of the image *cannot be improved* by conversion to other formats. Indeed, unless it is reproduced with exact integer scaling in your output document, image quality is likely to suffer.

You also need to be aware that the discrete cosine transform used in the JPEG image compression algorithm represents each pixel value by a weighted sum of neighboring pixels, and truncation of that sum results in visual smearing of that image. For that reason, JPEG representation is strongly deprecated for images with sharp edges, such as line drawings in graphs, and for images that contain only two, or a few, colors, such as a black-and-white image scan.

1.1.1 The first subsection

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

1.1.2 The second subsection

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

1.1.3 The third subsection

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

1.1.3.1 The first subsubsection

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Table 1.1. Lowercase Greek letters.

α	alpha
β	beta
γ	gamma
δ	delta
ϵ, ε	epsilon
ζ	zeta
η	eta
θ, ϑ	theta
ι	iota
κ	kappa
λ	lambda
μ	mu
ν	nu
ξ	xi
\omicron	omicron
π	pi
ρ	rho
σ, ς	sigma
τ	tau
υ	upsilon
ϕ, φ	phi
χ	chi
ψ	psi
ω	omega

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

In [Figure 1.2](#) on page 9, we have another picture.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah.

Table 1.2. Uppercase Greek letters. Notice that several have the same letter shapes as Latin letters, and for those, \TeX does not define macro names. For convenience, we supply our own definitions of these macros: `\Alpha`, `\Beta`, `\Epsilon`, `\Zeta`, `\Eta`, `\Iota`, `\Kappa`, `\Mu`, `\Nu`, `\Omicron`, `\Rho`, `\Tau`, and `\Chi`.

A	Alpha
B	Beta
Γ	Gamma
Δ	Delta
E	Epsilon
Z	Zeta
H	Eta
Θ	Theta
I	Iota
K	Kappa
Λ	Lambda
M	Mu
N	Nu
Ξ	Xi
O	Omicron
Π	Pi
P	Rho
Σ	Sigma
T	Tau
Y	Upsilon
Φ	Phi
X	Chi
Ψ	Psi
Ω	Omega

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
 Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
 blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

1.4 Free software packages

The Free Software Foundation offers almost 300 software packages, most easily portable to many different operating systems and CPU platforms. They include at least these:

`a2ps`, `acct`, `acm`, `adns`, `alive`, `anubis`, `apl`, `archimedes`, `aris`, `aspell`, `auctex`,
`autoconf-archive`, `autoconf`, `autogen`, `automake`, `avl`, `ballandpaddle`, `barcode`, `bash`,
`bayonne`, `bc`, `binutils`, `bison`, `bool`, `bpel2owfn`, `c-graph`, `ccaudio`, `ccd2cue`, `ccrtp`,

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

1.5 Resizing figures

In **Figure 1.4** on the next page through **Figure 1.8** on the facing page, we show how graphics files can be rescaled to convenient sizes.

```
\begin{figure}[p]
  \centerline{\includegraphics[scale = 0.5]{fig1}}
  \caption{The fourth figure (at 50\% scale).}%
  \figlabel{fig4}
\end{figure}

\begin{figure}[p]
  \centerline{\includegraphics[scale = 0.75]{fig1}}
  \caption{The fifth figure (at 75\% scale).}%
  \figlabel{fig5}
\end{figure}

\begin{figure}[p]
  \centerline{\includegraphics{fig1}}
  \caption{The sixth figure (at native size).}%
  \figlabel{fig6}
\end{figure}

\begin{figure}[p]
  \centerline{\includegraphics[scale = 1.25]{fig1}}
  \caption{The seventh figure (at 125\% scale).}%
  \figlabel{fig7}
\end{figure}

\begin{figure}[p]
  \centerline{\includegraphics[scale = 1.75]{fig1}}
  \caption{The eighth figure (at 175\% scale).}%
  \figlabel{fig8}
\end{figure}
```

You can include multiple images, each with its own caption inside a single *unbreakable*



This is Figure 1

Figure 1.4. The fourth figure (at 50% scale).



This is Figure 1

Figure 1.5. The fifth figure (at 75% scale).



This is Figure 1

Figure 1.6. The sixth figure (at native size).



This is Figure 1

Figure 1.7. The seventh figure (at 125% scale).



This is Figure 1

Figure 1.8. The eighth figure (at 175% scale).

figure environment, like this example shown in [Figure 1.9](#) and [Figure 1.10](#).

```
\begin{figure}[b]
  \centerline{\includegraphics[scale = 0.5]{fig1}}
  \caption{The fourth figure (at 50\% scale).}%
  \figlabel{fig9}
\end{figure}

\begin{figure}[b]
  \centerline{\includegraphics[scale = 0.75]{fig1}}
  \caption{The fifth figure (at 75\% scale).}%
  \figlabel{fig10}
\end{figure}
```

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

As a final example in this chapter, [Figure 1.11](#) on page 14 shows how you can use \LaTeX picture mode for annotating and positioning graphics images prepared outside \LaTeX .



This is Figure 1

[Figure 1.9](#). The ninth figure (at 50% scale), boxed with the tenth figure.



This is Figure 1

[Figure 1.10](#). The tenth figure (at 75% scale), boxed with the ninth figure.

The input that produced that figure looks like this:

```
\begin{figure}[b]
  %% The original image is 216bp wide by 72bp high, but we
  %% rescale it to 150 picture units divided by \unitlength:
  %% 150 / 0.75 = 112.5 mm
  \newcommand {\myfig} {\includegraphics[width = 112.5mm]{fig1}}

  \begin{center}
    %% The \unitlength is chosen to make the complete picture fit
    %% within the page margins

    \setlength{\unitlength}{0.75mm}

    %%%      insert (width,height)(lower-left-x,lower-left-y)
    \begin{picture}(170,70)(10,10)
      %% Place the included image FIRST!
      \put(10,10) {\myfig}

      %% Everything that follows OVERLAYS the original image!

      \graphpaper[10](0,0)(170,70)

      %% Mark the image center and corners by centered bullets
      \newcommand {\thedot} {\makebox (0,0) {$\bullet$}}
      \put( 85, 35) {\thedot}
      \put( 10, 10) {\thedot}
      \put( 10, 60) {\thedot}
      \put(160, 10) {\thedot}
      \put(160, 60) {\thedot}

      \put( 10, 10) {\makebox (0,0) [r] {lower-left}}
      \put(160, 10) {\makebox (0,0) [l] {lower-right}}
      \put( 10, 60) {\makebox (0,0) [r] {upper-left}}
      \put(160, 60) {\makebox (0,0) [l] {upper-right}}
    \end{picture}
  \end{center}

  \vspace{2\baselineskip}

  \caption[Using \LaTeX{} \texttt{picture} mode]
    {Using \LaTeX{} \texttt{picture} mode for figure labeling
     and positioning.}
  \figlabel{picture-mode}
\end{figure}
```

1.6 Summary and conclusions

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Here are three book references: [1, 2, 3]

1.7 References

- [1] G. FARMELO, *The Strangest Man: The Hidden Life of Paul Dirac, Mystic of the Atom*, Basic Books, New York, NY, USA, 2009.
- [2] F. W. J. OLVER, D. W. LOZIER, R. F. BOISVERT, AND C. W. CLARK, eds., *NIST Handbook of Mathematical Functions*, Cambridge University Press, Cambridge, UK, 2010.




Figure 1.11. Using \LaTeX picture mode for figure labeling and positioning.

- [3] S. SINGH, *Fermat's Enigma: The Epic Quest to Solve the World's Greatest Mathematical Problem*, Walker and Company, 435 Hudson Street, New York, NY 10014, USA, 1997.

Chapter 2

THE SECOND

This is a chapter.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Here are two article references: [1, 3].

The reference list at the end of this chapter is longer than just those two because of cross-referencing and citations within the BibTeX entries themselves.

2.1 References

- [1] H. P. BABBAE, *Babbage: Babbage's analytical engine*, Monthly Notices of the Royal Astronomical Society, 70 (1910), pp. 517–526, 645. Reprinted in [7, §2.3].
- [2] A. EINSTEIN, *Eine Neue Bestimmung der Moleküldimensionen*. (German) [A new determination of molecular dimensions], inaugural dissertation, Bern Wyss., Bern, Switzerland, 1905. Published in [3].
- [3] ———, *Eine neue Bestimmung der Moleküldimensionen*. (German) [A new determination of molecular dimensions], Annalen der Physik (1900) (series 4), 324 (1906), pp. 289–306. See corrections [4, 5]. This is a slightly revised version of Einstein's doctoral dissertation [2].
- [4] A. EINSTEIN, *Bemerkung zu meiner Arbeit: Eine Beziehung zwischen dem elastischen Verhalten*. (German) [Remark on my paper: "A relationship between the elastic behavior ..."], Annalen der Physik (1900) (series 4), 339 (1911), pp. 590–590. See [6].
- [5] A. EINSTEIN, *Berichtigung zu meiner Arbeit: Eine neue Bestimmung der Moleküldimensionen*. (German) [Corrections to my work: a new determination of molecular dimensions], Annalen der Physik (1900) (series 4), 339 (1911), pp. 591–592. See [3].
- [6] ———, *Eine Beziehung zwischen dem elastischen Verhalten und der spezifischen Wärme bei festen Körpern mit einatomigem Molekül*. (German) [A relationship between the elastic behavior and the specific heat of solid bodies with monatomic molecules], Annalen der Physik (1900) (series 4), 339 (1911), pp. 170–174, 590. See remarks [4, 5].

- [7] B. RANDELL, ed., *The Origins of Digital Computers: Selected Papers*, Texts and monographs in computer science, Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., third ed., 1982.

Chapter 3

THE THIRD

This is a chapter.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

We mention here three doctoral theses [1, 2, 3] and one corporate manual [8].

The reference list at the end of this chapter is longer than just those two because of cross-referencing and citations within the BibTeX entries themselves.

3.1 References

- [1] N. H. D. BOHR, *Studier over Metallernes Elektronteori. (Danish) [Studies on the electron theory of metals]*, doktor disputats, Københavns Universitet, København, Danmark, 1911. Afhandling for den filosofiske doktorgrad. [Thesis for the Doctor of Philosophy].
- [2] P. A. M. DIRAC, *Quantum Mechanics*, Ph.D. thesis, Cambridge University, Cambridge, UK, June 1926. According to [7, p. 101], this is the first thesis to be submitted anywhere on the subject of quantum mechanics.
- [3] A. EINSTEIN, *Eine Neue Bestimmung der Moleküldimensionen. (German) [A new determination of molecular dimensions]*, inaugural dissertation, Bern Wyss., Bern, Switzerland, 1905. Published in [4].
- [4] ———, *Eine neue Bestimmung der Moleküldimensionen. (German) [A new determination of molecular dimensions]*, *Annalen der Physik* (1900) (series 4), 324 (1906), pp. 289–306. See corrections [5, 6]. This is a slightly revised version of Einstein’s doctoral dissertation [3].
- [5] A. EINSTEIN, *Bemerkung zu meiner Arbeit: Eine Beziehung zwischen dem elastischen Verhalten. (German) [Remark on my paper: “A relationship between the elastic behavior ...”]*, *Annalen der Physik* (1900) (series 4), 339 (1911), pp. 590–590. See [?].
- [6] A. EINSTEIN, *Berichtigung zu meiner Arbeit: Eine neue Bestimmung der Moleküldimensionen. (German) [Corrections to my work: a new determination of molecular dimensions]*, *Annalen der Physik* (1900) (series 4), 339 (1911), pp. 591–592. See [4].

- [7] ———, *Eine Beziehung zwischen dem elastischen Verhalten und der spezifischen Wärme bei festen Körpern mit einatomigem Molekül. (German) [A relationship between the elastic behavior and the specific heat of solid bodies with monatomic molecules]*, *Annalen der Physik* (1900) (series 4), 339 (1911), pp. 170–174, 590. See remarks [5, 6].
- [8] G. FARMELO, *The Strangest Man: The Hidden Life of Paul Dirac, Mystic of the Atom*, Basic Books, New York, NY, USA, 2009.
- [9] POSTSCRIPT DEVELOPER TOOLS & STRATEGIES GROUP, ADOBE SYSTEMS INC., *Adobe font metric files specification — Version 3.0*, Mountain View, CA, USA, Mar. 1990.

Chapter 4

THE FOURTH

This is a chapter.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

4.1 More on the topic

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

4.2 Even more on the topic

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

Chapter 5

EXTENDING T_EX AND METAFONT WITH FLOATING POINT ARITHMETIC

In this chapter, we demonstrate how to incorporate another document, such as a previously-published paper, into a dissertation or thesis.

The University of Utah Thesis Office encourages this practice for publications related to the thesis, with these requirements:

- You must have received permission to reproduce the document from its copyright holder, which, for journals is normally the journal publisher. Such permissions would normally be granted on written request, but might take a few weeks to process, so plan ahead accordingly.
- A PDF file of the publication is included as *single* chapter, with optional leading commentary, and suitably scaled for a maximal fit within the normal page boundaries (in Spring 2016, 6.00 inches wide by 8.65 inches high). The standard L^AT_EX package [pdfpages](#) provides the necessary support for inclusion of all, or part, of a multipage PDF file.
- The included document has its original (usually single) line spacing, even though most of a thesis uses double spacing.
- Sectional headings corresponding to the included PDF file must be generated for the table of contents.
- Even if the rest of thesis includes page headers with short chapter and section titles, the included document should have only page numbers at the top of each page, because the inclusion may already have running headers.

- Because the included document usually has its own bibliography, any remaining chapters that cite other publications must have separate bibliographies as well. There is then no end-of-volume reference list.

Because the inclusion is a *verbatim* copy of a previously-published document, the Thesis Office accepts it as-is, and does not proofread it, nor does it require that it conform to the normal thesis formatting rules.

Two additional packages are required in the top-level \LaTeX document preamble:

```
\usepackage {pdfpages}           % support for multipage PDF inclusion
...
\usepackage {uuthesis-2016-f}     % usual fixups after standard packages
...
\usepackage {uuthesis-chapterbib} % support for chapter bibliographies
```

The chapter might begin with some commentary on the publication, and an acknowledgment of copyright permission.

It is followed by a setup command, then commands to generate table-of-contents entries, using the page numbers of the *original* publication. Those numbers are adjusted automatically for use in the table of contents. Finally, a single command includes a scaled version of the document as the contents of the rest of the chapter. That inclusion automatically starts on a new page.

For this sample thesis, here is what the input might look like:

```
\chapter{Extending \TeX{} and METAFONT with floating point arithmetic}%
\label {tug-paper}
...
\setupuuchapterbib
```

```
The article in this chapter was originally published in
\emph{TUGboat}, {\bf 28(3)} (2007) 501--510. It is reproduced
here with permission of the publisher.
```

```
\uudummysection {Dedication}           {1}
\uudummysection {Introduction}          {1}
\uudummysection {Arithmetic in \TeX{} and METAFONT} {1}
\uudummysection {Historical remarks}    {1}
\uudummysection {Why no floating-point arithmetic?} {3}
\uudummysection {IEEE 754 binary floating-point standard} {4}
```

```

\udummysection {IEEE 754R precision and range}           {4}
\udummysection {Remarks on floating-point arithmetic}    {5}
\udummysection {Binary versus decimal}                   {5}
\udummysection {Problems with IEEE 754 arithmetic}        {6}
\udummysection {How decimal arithmetic is different}      {6}
\udummysection {Software floating-point arithmetic}       {6}
\udummysection {How much work is needed}                  {8}
\udummysection {Summary}                                  {8}
\udummysection {References}                                {8}

%%% NB: - means all pages.  Adjust scale to fit in thesis page box:
\includepdf [
    pages = -,          % want all document pages
    scale = 0.91,       % adjust to fit thesis page box
    pagecommand = {\pagestyle{plain}} % bare page numbers
]
{tug2007.pdf}

```

If the included document has subsections and subsubsections, you could generate table-of-contents values for them with additional commands like these:

```

\udummysubsection {Blah blah blah}           {3}
\udummysubsubsection {Blah blah blah}        {3}
\udummysubsubsection {Blah blah blah blah}    {4}

```

However, you need not include more sectional levels than you use for other chapters of your thesis.

At this point, we issue the block of commands above, beginning with `\setupuuchapterbib`:

The article in this chapter was originally published in *TUGboat*, **28(3)** (2007) 501–510. It is reproduced here with permission of the publisher.

Extending T_EX and METAFONT with floating-point arithmetic

Nelson H. F. Beebe

University of Utah

Department of Mathematics, 110 LCB

155 S 1400 E RM 233

Salt Lake City, UT 84112-0090

USA

WWW URL: <http://www.math.utah.edu/~beebe>

Telephone: +1 801 581 5254

FAX: +1 801 581 4148

Internet: beebe@math.utah.edu, beebe@acm.org, beebe@computer.org

Abstract

The article surveys the state of arithmetic in T_EX and METAFONT, suggests that they could usefully be extended to support floating-point arithmetic, and shows how this could be done with a relatively small effort, *without* loss of the important feature of platform-independent results from those programs, and *without* invalidating any existing documents, or software written for those programs, including output drivers.

Contents

1 Dedication	501
2 Introduction	501
3 Arithmetic in T _E X and METAFONT	501
4 Historical remarks	503
5 Why no floating-point arithmetic?	503
6 IEEE 754 binary floating-point standard	504
7 IEEE 754R precision and range	504
8 Remarks on floating-point arithmetic	505
9 Binary versus decimal	505
10 Problems with IEEE 754 arithmetic	506
11 How decimal arithmetic is different	506
12 Software floating-point arithmetic	506
13 How much work is needed?	508
14 Summary	508

1 Dedication

This article is dedicated to Professors Donald Knuth (Stanford University) and William Kahan (University of California, Berkeley), with thanks for their many scientific and technical contributions, and for their writing.

2 Introduction

Arithmetic is a fundamental feature of computer programming languages, and for some of us, the more we use computers, the more inadequate we find their computational facilities. The arithmetic in T_EX and METAFONT is particularly limiting, and this article explains why this is so, why it was not otherwise, and what can be done about it now that these two important programs are in their thirtieth year of use.

3 Arithmetic in T_EX and METAFONT

Before we look at issues of arithmetic in general, it is useful to summarize what kinds of numbers T_EX and METAFONT can handle, and how they do so.

T_EX provides binary integer and fixed-point arithmetic. Integer arithmetic is used to count things, such as with T_EX's \count registers. Fixed-point arithmetic is needed for values that have fractional parts, such as the \dimen dimension registers, the \muskip and \skip glue registers, and scale factors, as in 0.6\hsize.

For portability reasons, T_EX requires that the host computer support an integer data type of at least 32 bits. It uses that type for the integer arithmetic available to T_EX programs. For fixed-point numbers, it reserves the lower 16 bits for the fractional part, and all but two of the remaining bits for the integer part. Thus, on the 32-bit processors that are commonly found in personal computers, 14 bits are available for the integer part. One of the remaining two bits is chosen as a sign bit, and the other is used to detect *overflow*, that is, generation of a number that is too large to represent.

Nelson H. F. Beebe

When fractional numbers represent \TeX dimensions, the low-order fraction bit represents the value 2^{-16} pt. While printer's points have been a common unit of measurement since well before the advent of computer-based typesetting, this tiny value is new with \TeX , and has the special name *scaled point*. The value 1 sp is so small that approximately 100 sp is about the wavelength of visible light. This ensures that differences of a few scaled points in the positioning of objects on the printed page are completely invisible to human eyes.

The problem with fixed-point numbers in \TeX is at the other end: 14 integer bits can only represent numbers up to 16383. As a dimension, that many points is about 5.75 m, which is probably adequate for printed documents, but is marginal if you are typesetting a billboard. The PDP-10 computers on which \TeX and METAFONT were developed had 36-bit words: the four extra bits raised the maximum dimension by a factor of 16. Nevertheless, if \TeX 's fixed-point numbers are used for purposes other than page dimensions, then it is easy to exceed their limits.

\TeX is a macro-extensible language for typesetting, and arithmetic is expected to be relatively rare. \TeX has little support for numerical expressions, just verbose low-level operators, forcing the \TeX programmer to write code such as this fragment from layout.tex to accomplish the multiply-add operation noted in the comment:

```
% MRGNOTEYA = 0.75*TEXTHEIGHT + FOOTSKIP
\T = \TEXTHEIGHT
\multiply \T by 75 % possible overflow!
\divide \T by 100
\advance \T by \FOOTSKIP
\edef \MRGNOTEYA {\the \T}
```

Notice that the scale factor 0.75 could have been reduced from 75/100 to 3/4 in this example, but that is not in general possible. Similarly, here we could have written $\backslash T = 0.75 \backslash \text{TEXTHEIGHT}$, but that is not possible if the constant 0.75 is replaced by a variable in a register. The multiplication by 75 can easily provoke an overflow if $\backslash T$ is even as big as a finger length:

```
*\dimen1 = 220pt

*\dimen2 = 75\dimen1
! Dimension too large.

*\multiply \dimen1 by 75
! Arithmetic overflow.
```

See the \LaTeX calc package for more horrors of fixed-point arithmetic.

\TeX has, however, also seen use a scripting language, chosen primarily because of its superb quality,

stability, reliability, and platform independence. \TeX distributions now contain macro packages and utilities written in \TeX for generating complex font tables, for packing and unpacking document archives, for scanning PostScript graphics files, and even for parsing SGML and XML.

\TeX 's arithmetic does not go beyond the four basic operations of *add*, *subtract*, *multiply*, and *divide*. In particular, no elementary functions (square root, exponential, logarithm, trigonometric and hyperbolic functions, and so on) are provided in \TeX itself, even though, in principle, they can be provided with macro packages.

In \TeX , overflow is detected in division and multiplication but not in addition and subtraction, as I described in my TUG 2003 keynote address [4].

Input numbers in METAFONT are restricted to 12 integer bits, and the result of even trivial expressions can be quite surprising to users:

```
% mf expr
gimme an expr: 4095      >> 4095
gimme an expr: 4096
! Enormous number has been reduced.
>> 4095.99998
gimme an expr: infinity  >> 4095.99998
gimme an expr: epsilon   >> 0.00002
gimme an expr: 1/epsilon
! Arithmetic overflow.
>> 32767.99998
gimme an expr: 1/3       >> 0.33333
gimme an expr: 3*(1/3)   >> 0.99998
gimme an expr: 1.2 - 2.3 >> -1.1
gimme an expr: 1.2 - 2.4 >> -1.2
gimme an expr: 1.3 - 2.4 >> -1.09999
```

Notice that although 4096 is considered an overflow, internally METAFONT can generate a number almost eight times as large. Binary-to-decimal conversion issues produce the anomaly in $3 \times (1/3)$. The last line shows that even apparently simple operations are not so simple after all.

Overflows in METAFONT can also produce a report like this:

```
Uh, oh. A little while ago one of the
quantities that I was computing got
too large, so I'm afraid your answers
will be somewhat askew. You'll
probably have to adopt different
tactics next time. But I shall try to
carry on anyway.
```

METAFONT provides a few elementary functions: ++ (Pythagoras), abs, angle, ceiling, cosd, dir, floor, length, mexp, mlog, normaldeviate, round,

`sind`, `sqrt`, and `uniformdeviate`. They prove useful in the geometric operations required in font design.

4 Historical remarks

T_EX and METAFONT are not the only systems that suffer from the limitations of fixed-point arithmetic. Most early computers were inadequate as well:

It is difficult today to appreciate that probably the biggest problem facing programmers in the early 1950s was scaling numbers so as to achieve acceptable precision from a fixed-point machine.

Martin Campbell-Kelly
Programming the Mark I: Early Programming Activity at the University of Manchester
Annals of the History of Computing,
2(2) 130–168 (1980)

Scaling problems can be made much less severe if numbers carry an exponent as well as integer and fractional parts. We then have:

Floating Point Arithmetic ... The subject is not at all as trivial as most people think, and it involves a surprising amount of interesting information.

Donald E. Knuth
The Art of Computer Programming:
Seminumerical Algorithms (1998)

However, more than just an exponent is needed; the arithmetic system also has to be predictable:

Computer hardware designers can make their machines much more pleasant to use, for example by providing *floating-point arithmetic* which satisfies simple mathematical laws.

The facilities presently available on most machines make the job of rigorous error analysis *hopelessly difficult*, but properly designed operations would encourage numerical analysts to provide better subroutines which have certified accuracy.

Donald E. Knuth
Computer Programming as an Art
ACM Turing Award Lecture (1973)

5 Why no floating-point arithmetic?

Neither T_EX nor METAFONT have floating-point arithmetic natively available, and as I discussed in my *Practical T_EX 2005* keynote address [3], there is a very good reason why this is the case. Their output needs to be identical on all platforms, and when they were developed, there were many different computer vendors, some of which had several incompatible product lines.

This diversity causes several problems, some of which still exist:

- There is system dependence in *precision*, *range*, *rounding*, *underflow*, and *overflow*.
- The number base varies from 2 on most, to 3 (Setun), 4 (Illiad II), 8 (Burroughs), 10, 16 (IBM S/360), 256 (Illiad III), and 10000 (Maple).
- Floating-point arithmetic exhibits bizarre behavior on some systems:
 - $x \times y \neq y \times x$ (early Crays);
 - $x \neq 1.0 \times x$ (Pr1me);
 - $x + x \neq 2 \times x$ (Pr1me);
 - $x \neq y$ but $1.0/(x - y)$ gets zero-divide error;
 - wrap between underflow and overflow (e.g., C on PDP-10);
 - job termination on overflow or zero-divide (most).
- No standardization: almost every vendor had one or more distinct floating-point systems.
- Programming language dependence on available precisions:
 - Algol, Pascal, and SAIL (only `real`): recall that SAIL was the implementation language for the 1977–78 prototypes of T_EX and METAFONT;
 - Fortran (`REAL`, `DOUBLE PRECISION`, and on some systems, `REAL*10` or `REAL*16`);
 - C/C++ (originally only `double`, but `float` added in 1989, and `long double` in 1999);
 - C# and Java have only `float` and `double` data types, but their arithmetic is badly botched: see Kahan and Darcy's *How Java's Floating-Point Hurts Everyone Everywhere* [28].
- Compiler dependence: multiple precisions can be mapped to just one, without warning.
- BSD compilers on IA-32 still provide no 80-bit format after 27 years in hardware.
- Input/output problem requires base conversion, and is *hard* (e.g., conversion from 128-bit binary format can require more than 11500 decimal digits).
- Most languages do not guarantee exact base conversion.

Donald Knuth wrote an interesting article with the intriguing title *A simple program whose proof isn't* [29] about how T_EX handles conversions between fixed-point binary and decimal. The restriction to fixed-point arithmetic with 16-bit fractional parts simplifies the base-conversion problem, and allows T_EX to

Nelson H. F. Beebe

guarantee exact round-trip conversions of such numbers.

TeX produces the same line- and page-breaking across all platforms, because floating-point arithmetic is used only for interword glue calculations that could change the horizontal position of a letter by at most a few scaled points, but as we noted earlier, that is invisible.

METAFONT has no floating-point at all, and generates identical fonts on all systems.

6 IEEE 754 binary floating-point standard

With the leadership of William Kahan, a group of researchers in academic, government, and industry began a collaborative effort in the mid-1970s to design a new and much improved floating-point architecture. The history of this project is chronicled in an interview with Kahan [37, 38].

A preliminary version of this design was first implemented in the Intel 8087 chip in 1980, although the design was not finalized until its publication as IEEE Standard 754 in 1985 [23].

Entire books have been written about floating-point arithmetic: see, for example, Sterbenz [41] for historical systems, Overton [35] for modern ones, Omondi [34] and Parhami [36] for hardware, Goldberg [15, 17] for an excellent tutorial, and Knuth [30, Chap. 4] for theory. I am hard at work on writing two more books in this area. However, here we need only summarize important features of the IEEE 754 system:

- Three formats are defined: 32-bit, 64-bit, and 80-bit. A 128-bit format was subsequently provided on some Alpha, IA-64, PA-RISC, and SPARC systems.
- Nonzero normal numbers are *rational*: $x = (-1)^s f \times 2^p$, where the *significand*, f , lies in $[1, 2)$.
- Signed zero allows recording the direction from which an underflow occurred, and is particularly useful for arithmetic with complex (real + imaginary) numbers. The IEEE Standard requires that $\sqrt{-0}$ evaluate to -0 .
- The largest stored exponent represents Infinity if $f = 0$, and either quiet or signaling NaN (Not-a-Number) if $f \neq 0$. A vendor-chosen significand bit distinguishes between the two kinds of NaN.
- The smallest stored exponent allows leading zeros in f for *gradual underflow* to *subnormal* values.
- The arithmetic supports a model of fast *nonstop* computing. Sticky flags record exceptions, and Infinity, NaN, and zero values automatically

replace out-of-range values, without the need to invoke an exception handler, although that capability may also be available.

- Four rounding modes are provided:
 - *to nearest with ties to even* (default);
 - *to $+\infty$* ;
 - *to $-\infty$* ;
 - *to zero* (historical chopping).
- Values of $\pm\infty$ are generated from huge/tiny and finite/0.
- NaN values are generated from 0/0, $\infty - \infty$, ∞/∞ , and any operation with a NaN operand.
- A NaN is returned from functions when the result is undefined in real arithmetic (e.g., $\sqrt{-1}$), or when an argument is a NaN.
- NaNs have the property that they are unequal to anything, even themselves. Thus, the C-language inequality test $x != x$ is true *if, and only if*, x is a NaN, and should be readily expressible in *any* programming language. Sadly, several compilers botch this, and get the wrong answer.

7 IEEE 754R precision and range

In any computer arithmetic system, it is essential to know the available range and precision. The precisions of the four IEEE 754 binary formats are equivalent to approximately 7, 15, 19, and 34 decimal digits. The approximate ranges as powers of ten, including subnormal numbers, are $[-45, 38]$, $[-324, 308]$, $[-4951, 4932]$, and $[-4966, 4932]$. A future 256-bit binary format will supply about 70 decimal digits, and powers-of-ten in $[-315723, 315652]$.

A forthcoming revision of the IEEE Standard will include decimal arithmetic as well, in 32-, 64-, and 128-bit storage sizes, and we can imagine a future 256-bit size. Their precisions are 7, 16, 34, and 70 decimal digits, where each doubling in size moves from n digits to $2n + 2$ digits. Their ranges are wider than the binary formats, with powers of ten in $[-101, 96]$, $[-398, 384]$, $[-6176, 6144]$, and $[-1572932, 1572864]$.

In each case, the range and precision are determined by the number of bits allocated for the sign and the significand, and for the decimal formats, by restrictions imposed by the compact encodings chosen for packing decimal digits into strings of bits.

It is highly desirable that each larger storage size increase the exponent range (many older designs did not), and at least double the significand length, since that guarantees that products evaluated in the next higher precision *cannot overflow*, and are *exact*. For example, the Euclidean distance $\sqrt{x^2 + y^2}$ is then trivial

to compute; otherwise, its computation requires careful rescaling to avoid premature underflow and overflow.

8 Remarks on floating-point arithmetic

Contrary to popular misconception, even present in some books and compilers, floating-point arithmetic is *not* fuzzy.

- Results are *exact* if they are representable.
- Multiplication by a power of base is always exact, in the absence of underflow and overflow.
- Subtracting numbers of like signs and exponents is *exact*.

Bases other than 2 or 10 suffer from *wobbling precision* caused by the requirement that significands be normalized. For example, in hexadecimal arithmetic, $\pi/2 \approx 1.571 \approx 1.922_{16}$ has three fewer bits (almost one decimal digit) than $\pi/4 \approx 0.7854 \approx c.910_{16}$. Careful coders on such systems account for this in their programs by writing

```
y = (x + quarter_pi) + quarter_pi;
```

instead of

```
y = x + half_pi;
```

Because computer arithmetic systems have finite range and precision, they are not *associative*, so commonly-assumed mathematical transformations do not hold. In particular, it is often necessary to control evaluation order, and this may be at odds with what the compiler, or even a high-performance CPU with dynamic instruction reordering, does with the code.

The presence of multiple rounding modes also invalidates common assumptions. For example, the Taylor series for the sine function begins

$$\sin(x) = x - (1/3!)x^3 + (1/5!)x^5 - \dots$$

If x is small enough, because of finite precision, one might expect that $\sin(x)$ could be computed simply as x . However, that is only true for the default rounding mode; in other modes, the correct answer could be one *ulp* (unit in the last place) higher or lower, so at least two terms must be summed. Similarly, the mathematical equivalence $-(xy + z) \equiv (-xy - z)$ does not hold in some rounding modes. Except for some special numbers, it is not in general permissible to replace slow division with fast multiplication by the reciprocal, even though many optimizing compilers do that.

Some of the common elementary functions are *odd* ones: they satisfy $f(x) = -f(-x)$. This relation does not in general hold computationally if a rounding direction of other than *round-to-nearest* is in effect. Software designers are then forced to decide whether

obeying computer rounding modes is more important than preserving fundamental mathematical symmetries: in well-designed software, symmetry wins. Nevertheless, in some applications, like *interval arithmetic*, which computes upper and lower bounds for every numeric operation, precise control of rounding is imperative, and overrides symmetry.

See Monniaux [33] for a recent discussion of some of the many problems of floating-point evaluation. A good part of the difficulties described there arise because of higher intermediate precision in the Intel IA-32 architecture, the most common desktop CPU family today. Other problems come from unexpected instruction reordering or multiple threads of execution, and the incidence of these issues increases with each new generation of modern processors.

9 Binary versus decimal

Why should we care whether a computer uses binary or decimal arithmetic? Here are some reasons why a switch to decimal arithmetic has advantages:

- Humans are less uncomfortable with decimal arithmetic.
- In some case, binary arithmetic always gets the wrong answer. Consider this sales tax computation: 5% of 0.70 = 0.0349999... in *all* binary precisions, instead of the exact decimal 0.035. Thus, there can be significant cumulative rounding errors in businesses with many small transactions (food, music downloading, telephone, ...).
- Financial computations need fixed-point decimal arithmetic.
- Hand calculators use decimal arithmetic.
- Additional decimal rounding rules (eight instead of four) handle the financial and legal requirements of some jurisdictions.
- Decimal arithmetic eliminates most base-conversion problems.
- There is a specification of decimal arithmetic subsumed in the *IEEE 854-1987 Standard for Radix-Independent Floating-Point Arithmetic* [21].
- Older Cobol standards require 18D fixed-point.
- Cobol 2002 requires 32D fixed-point *and* floating-point.
- Proposals to add decimal arithmetic to C and C++ were submitted to the ISO language committees in 2005 and 2006.
- Twenty-five years of Rexx and NetRexx scripting languages give valuable experience in arbitrary-precision decimal arithmetic.

Nelson H. F. Beebe

- The excellent IBM `decNumber` library provides *open source* decimal floating-point arithmetic with a billion (10^9) digits of precision and exponent magnitudes up to 999 999 999.
- Preliminary support in `gcc` for `+`, `-`, `*`, and `/` became available in late 2006, based on a subset of the IBM `decNumber` library.
- The author's `mathcw` package [5] provides a C99-compliant run-time library for binary, and also for decimal, arithmetic (2005–2008), with hundreds of additional functions, and important and useful extensions of the I/O functions.
- IBM zSeries mainframes got IEEE 754 binary floating-point arithmetic in 1999, and decimal floating-point arithmetic in firmware in 2006.
- The IBM PowerPC version 6 chips announced on 21 May 2007 add hardware decimal arithmetic, probably the first mainstream processor to do so in more than four decades.
- Hardware support seems likely in future Intel IA-32 and EM64T (x86_64) processors, and the current family members are among the most widely-used in the world for general-purpose computing. Other chip vendors will have to offer similar facilities to remain competitive.

10 Problems with IEEE 754 arithmetic

Despite the many benefits of IEEE 754 floating-point arithmetic, there are many impediments to its effective use:

- Language access to features has been slow: more than 27 years have passed since the Intel 8087, and we are still waiting!
- Programmer unfamiliarity, ignorance, and inexperience.
- A deficient educational system, both in academia, and in textbooks, leaves most programmers with little or no training in floating-point arithmetic.
- Partial implementations by some vendors deny access to important features (e.g., subnormals may flush to zero, IA-32 has only one NaN, IA-32 and IA-64 have imperfect rounding, Java and C# lack rounding modes and higher precisions).
- Long internal registers are generally beneficial, but also produce many computational surprises and double rounding [33], compromising portability.
- Rounding behavior at underflow and overflow limits is unspecified by the IEEE standards, and thus, is vendor dependent.

- Overeager, or incorrect, optimizations by compilers may produce wrong results, and prevent obtaining similar results across different platforms, or between different compilers on the same system, or even from the same compiler with different options.
- Despite decades of availability of IEEE 754 arithmetic, some compilers still mishandle signed zeros and NaNs, and it can be difficult to convince compiler vendors of the significance of such errors (I know, because I've tried, and failed).

11 How decimal arithmetic is different

Programmers in science and engineering have usually only had experience with binary floating-point arithmetic, and some relearning is needed for the move to decimal arithmetic:

- Nonzero normal floating-point numbers take the form $x = (-1)^s f \times 10^p$, where f is an *integer*, allowing simulation of fixed-point arithmetic.
- Lack of normalization means multiple storage forms, but 1., 1.0, 1.00, 1.000, ... compare equal, as long as floating-point instructions, rather than bitwise integer comparisons, are used.
- *Quantization* is detectable (e.g., for financial computations, 1.00 differs from 1.000).
- Signed zero and infinity, plus quiet and signaling NaNs, are detectable from the first byte, whereas binary formats require examination of all bits.
- There are *eight* rounding modes because of legal and tax mandates.
- Compact storage formats—Densely-Packed Decimal (DPD) [IBM] and Binary-Integer Decimal (BID) [Intel]—need fewer than BCD's four bits per decimal digit.

12 Software floating-point arithmetic

It may be better in some applications to have floating-point arithmetic entirely in software, as Apple once did with the no-longer-supported SANE (Standard Apple Numerics Environment) system. Here are some reasons why:

- `TEX` and METAFONT must continue to guarantee identical results across platforms.
- Unspecified behavior of low-level arithmetic guarantees *platform dependence*.
- Floating-point arithmetic is not associative, so instruction ordering (e.g., compiler optimization) affects results.

- Long internal registers on some platforms, and not on others, alter precision, and results.
- Multiply-add computes $x \times y + z$ with *exact* product and single rounding, getting different result from separate operations.
- Conclusion: only a single *software* floating-point arithmetic system in T_EX and METAFONT can guarantee *platform-independent results*.

Software is often best enhanced by connecting two or more systems with a clean and simple interface:

What if you could provide a seamlessly integrated, fully dynamic language with a conventional syntax while increasing your application's size by less than 200K on an x86? You can do it with *Lua*!

Keith Fieldhouse

If we want to have floating-point arithmetic in T_EX and METAFONT, then rather than modify those stable and reliable programs, including adding convenient expression syntax, and a substantial function library, there is a cleaner, and easier, approach:

- There is no need to modify T_EX beyond what has already been done: LuaT_EX interfaces T_EX to a clean and well-designed scripting language — we just need to change the arithmetic and library inside *lua*.
- Scripting languages usually offer a single floating-point datatype, typically equivalent to IEEE 754 64-bit double (that is all that the C language used to have).
- *qawk* and *dnawk* are existing extensions by the author of *awk* for 128-bit binary and decimal arithmetic, respectively.
- Modern machines are fast and memories are big. We could adopt a 34D 128-bit format, or better, a 70D 256-bit format, instead as default numeric type.
- The author's *mathcw* package [5] is a highly-portable open-source library with support for *ten* floating-point precisions, including 256-bit binary and decimal.

Two more quotes from the father of the IEEE 754 design lead into our next points:

The convenient accessibility of double-precision in many Fortran and some Algol compilers indicates that double-precision will soon be universally acceptable as a substitute for ingenuity in the solution of numerical problems.

W. Kahan

Further Remarks on Reducing Truncation Errors

Comm. ACM **8**(1) 40, January (1965)

Nobody knows how much it would cost to compute y^w correctly rounded for every two floating-point arguments at which it does not over/underflow. Instead, reputable math libraries compute elementary transcendental functions mostly within slightly more than half an ulp and almost always well within one ulp. Why can't y^w be rounded within half an ulp like SQRT? Because nobody knows how much computation it would cost.... No general way exists to predict how many extra digits will have to be carried to compute a transcendental expression and round it correctly to some preassigned number of digits. Even the fact (if true) that a finite number of extra digits will ultimately suffice may be a deep theorem.

W. Kahan

Wikipedia entry

We need more than just the basic four operations of arithmetic: several dozen elementary functions, and I/O support, are essential.

- The *Table Maker's Dilemma* (Kahan) is the problem of always getting exactly-rounded results when computing the elementary functions. Here is an example of a hard case:
 $\log(+0x1.ac50b409c8aeep+8) =$
 $0x60f52f37aecfcfffffffffffffeeb...p-200$
 There are 62 consecutive 1-bits in that number, and at least $4 \times 13 + 62 + 1 = 115$ bits must be computed correctly in order to determine the correctly-rounded 53-bit result.
- Higher-than-needed-precision arithmetic provides a practical solution to the dilemma, as the Kahan quote observes.
- Random-number generation is a common portability problem, since algorithms for that computation are platform-dependent and vary in quality. Fortunately, several good ones are now known, and can be supplied in libraries, although careful attention still needs to be given to computer wordsize.
- The *mathcw* library gives *platform-independent* results for decimal floating-point arithmetic, since evaluation order is completely under programmer control, and identical everywhere, and the underlying decimal arithmetic is too.

Nelson H. F. Beebe

13 How much work is needed?

I argue that decimal floating-point arithmetic in software, isolated in a separate scripting language, is an effective and reasonable way to extend T_EX and METAFONT so that they can have access to floating-point arithmetic, and remove the limitations and nuisance of fixed-point arithmetic that they currently suffer.

It is therefore appropriate to ask what kind of effort would be needed to do this. In four separate experiments with three implementations of awk, and one of lua, I took two to four hours each, with less than 3% of the code requiring changes:

Program	Lines	Deleted	Added
dgawk	40 717	109	165
dlua	16 882	25	94
dmawk	16 275	73	386
dnawk	9 478	182	296
METAFONT in C	30 190	0	0
T _E X in C	25 215	0	0

14 Summary

Had the IEEE 754 design been developed before T_EX and METAFONT, it is possible that Donald Knuth would have chosen a software implementation of binary floating-point arithmetic, as he later provided for the MMIX virtual machine [2, 31, 32] that underlies the software analyses in newer editions of his famous book series, *The Art of Computer Programming*.

That did not happen, so in this article, I have shown how a different approach might introduce decimal floating-point arithmetic to T_EX and METAFONT through a suitable scripting language, for which Lua [26, 25, 27] seems eminently suited, and has already been interfaced to T_EX and is now in limited use for production commercial typesetting, and also for document style-file design. By selecting high working precision, at least 34 decimal digits and preferably 70, many numerical issues that otherwise compromise portability and reproducibility of typeset documents simply disappear, or at least, become highly improbable.

To make this workable, the compilers, the basic software arithmetic library, the elementary function library, and the I/O library need to be highly portable. The combination of the GNU gcc compiler family with the IBM decNumber library and the author's mathcw library satisfy all of these requirements. Within a year or two, we may therefore expect that decimal floating-point arithmetic in C could be available on all of the common platforms, allowing future T_EX Live releases to build upon that foundation, and LuaT_EX could become the T_EX version of choice in many environments. LuaMETAFONT and LuaMETAPOST could soon follow.

References

- [1] P. H. Abbott, D. G. Brush, C. W. Clark III, C. J. Crone, J. R. Ehrman, G. W. Ewart, C. A. Goodrich, M. Hack, J. S. Kapernick, B. J. Minchau, W. C. Shepard, R. M. Smith, Sr., R. Tallman, S. Walkowiak, A. Watanabe, and W. R. White. Architecture and software support in IBM S/390 Parallel Enterprise Servers for IEEE floating-point arithmetic. *IBM Journal of Research and Development*, 43(5/6):723–760, 1999. ISSN 0018-8646. URL <http://www.research.ibm.com/journal/rd/435/abbott.html>. Besides important history of the development of the S/360 floating-point architecture, this paper has a good description of IBM's algorithm for exact decimal-to-binary conversion, complementing earlier ones [39, 7, 29, 6, 40].
- [2] Heidi Anlauff, Axel Böttcher, and Martin Ruckert. *Das MMIX-Buch: ein praxisnaher Zugang zur Informatik. (German) [The MMIX Book: A practical introduction to computer science]*. Springer-Lehrbuch. Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 2002. ISBN 3-540-42408-3. xiv + 327 pp. EUR 24.95. URL <http://www.informatik.fh-muenchen.de/~mmix/MMIXBuch/>.
- [3] Nelson Beebe. The design of T_EX and METAFONT: A retrospective. *TUGboat*, 26(1):33–41, 2005. ISSN 0896-3207.
- [4] Nelson H. F. Beebe. 25 Years of T_EX and METAFONT: Looking back and looking forward — TUG 2003 keynote address. *TUGboat*, 25(1):7–30, 2004. ISSN 0896-3207.
- [5] Nelson H. F. Beebe. *The mathcw Portable Elementary Function Library*. 2008. In preparation.
- [6] Robert G. Burger and R. Kent Dybvig. Printing floating-point numbers quickly and accurately. *ACM SIGPLAN Notices*, 31(5):108–116, May 1996. ISSN 0362-1340. URL <http://www.acm.org:80/pubs/citations/proceedings/pldi/231379/p108-burger/>. This paper offers a significantly faster algorithm than that of [39], together with a correctness proof and an implementation in Scheme. See also [7, 1, 40, 8].
- [7] William D. Clinger. How to read floating point numbers accurately. *ACM SIGPLAN Notices*, 25(6):92–101, June 1990. ISBN 0-89791-364-7. ISSN 0362-1340. URL <http://www.acm.org:80/pubs/citations/proceedings/pldi/93542/p92-clinger/>. See also output algorithms in [29, 39, 6, 1, 40].
- [8] William D. Clinger. Retrospective: How to read floating point numbers accurately. *ACM SIGPLAN Notices*, 39(4):360–371, April 2004. ISSN 0362-1340. Best of PLDI 1979–1999. Reprint of, and retrospective on, [7].
- [9] William J. Cody, Jr. Analysis of proposals for the floating-point standard. *Computer*, 14(3):63–69, March 1981. ISSN 0018-9162. See [23, 24].

- [10] Jerome T. Coonen. An implementation guide to a proposed standard for floating-point arithmetic. *Computer*, 13(1):68–79, January 1980. ISSN 0018-9162. See errata in [11]. See [23, 24].
- [11] Jerome T. Coonen. Errata: An implementation guide to a proposed standard for floating point arithmetic. *Computer*, 14(3):62, March 1981. ISSN 0018-9162. See [10, 23, 24].
- [12] Jerome T. Coonen. Underflow and the denormalized numbers. *Computer*, 14(3):75–87, March 1981. ISSN 0018-9162. See [23, 24].
- [13] Charles B. Dunham. Surveyor's Forum: "What every computer scientist should know about floating-point arithmetic". *ACM Computing Surveys*, 24(3):319, September 1992. ISSN 0360-0300. See [15, 16, 45].
- [14] W. H. J. Feijen, A. J. M. van Gasteren, D. Gries, and J. Misra, editors. *Beauty is our business: a birthday salute to Edsger W. Dijkstra*. Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1990. ISBN 0-387-97299-4. xix + 453 pp. LCCN QA76 .B326 1990.
- [15] David Goldberg. What every computer scientist should know about floating-point arithmetic. *ACM Computing Surveys*, 23(1):5–48, March 1991. ISSN 0360-0300. URL <http://www.acm.org/pubs/toc/Abstracts/0360-0300/103163.html>. See also [16, 13, 45].
- [16] David Goldberg. Corrigendum: "What every computer scientist should know about floating-point arithmetic". *ACM Computing Surveys*, 23(3):413, September 1991. ISSN 0360-0300. See [15, 13, 45].
- [17] David Goldberg. Computer arithmetic. In *Computer Architecture—A Quantitative Approach*, chapter H, pages H-1–H-74. Morgan Kaufmann Publishers, Los Altos, CA 94022, USA, third edition, 2002. ISBN 1-55860-596-7. LCCN QA76.9.A73 P377 2003. US\$89.95. URL <http://books.elsevier.com/companions/1558605967/appendices/1558605967-appendix-h.pdf>. The complete Appendix H is not in the printed book; it is available only at the book's Web site: <http://www.mkp.com/CA3>.
- [18] David Gries. Binary to decimal, one more time. In Feijen et al. [14], chapter 16, pages 141–148. ISBN 0-387-97299-4. LCCN QA76 .B326 1990. This paper presents an alternate proof of Knuth's algorithm [29] for conversion between decimal and fixed-point binary numbers.
- [19] David Hough. Applications of the proposed IEEE-754 standard for floating point arithmetic. *Computer*, 14(3):70–74, March 1981. ISSN 0018-9162. See [23, 24].
- [20] IEEE. IEEE standard for binary floating-point arithmetic. *ACM SIGPLAN Notices*, 22(2):9–25, February 1985. ISSN 0362-1340. See [23].
- [21] IEEE. *854-1987 (R1994) IEEE Standard for Radix-Independent Floating-Point Arithmetic*. IEEE, New York, NY, USA, 1987. ISBN 1-55937-859-X. 16 pp. US\$44.00. URL http://standards.ieee.org/reading/ieee/std_public/description/busarch/854-1987_desc.html. Revised 1994.
- [22] IEEE Computer Society Standards Committee. Working group of the Microprocessor Standards Subcommittee and American National Standards Institute. *IEEE standard for binary floating-point arithmetic*. ANSI/IEEE Std 754-1985. IEEE Computer Society Press, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1985. 18 pp. See [23].
- [23] IEEE Task P754. *ANSI/IEEE 754-1985, Standard for Binary Floating-Point Arithmetic*. IEEE, New York, NY, USA, August 12, 1985. ISBN 1-55937-653-8. 20 pp. US\$35.00. URL http://standards.ieee.org/reading/ieee/std_public/description/busarch/754-1985_desc.html; <http://standards.ieee.org/reading/ieee/std/busarch/754-1985.pdf>; <http://www.iec.ch/cgi-bin/procgi.pl/www/iecwww.p?wwwlang=E&wwwprog=cat-det.p&wartnum=019113>; <http://ieeexplore.ieee.org/iel1/2355/1316/00030711.pdf>. Revised 1990. A preliminary draft was published in the January 1980 issue of IEEE Computer, together with several companion articles [9, 12, 10, 11, 19, 42, 43]. The final version was republished in [20, 22]. See also [44]. Also standardized as *IEC 60559 (1989-01) Binary floating-point arithmetic for microprocessor systems*.
- [24] IEEE Task P754. *ANSI/IEEE 754-1985, Standard for Binary Floating-Point Arithmetic*. IEEE, New York, August 12 1985. A preliminary draft was published in the January 1980 issue of IEEE Computer, together with several companion articles [9, 12, 10, 11, 19, 42, 43]. Available from the IEEE Service Center, Piscataway, NJ, USA.
- [25] Roberto Ierusalimschy. *Programming in Lua*. Lua.Org, Rio de Janeiro, Brazil, 2006. ISBN 85-903798-2-5. 328 (est.) pp.
- [26] Roberto Ierusalimschy, Luiz Henrique de Figueiredo, and Waldemar Celes. *Lua 5.1 Reference Manual*. Lua.Org, Rio de Janeiro, Brazil, 2006. ISBN 85-903798-3-3. 112 (est.) pp.
- [27] Kurt Jung and Aaron Brown. *Beginning Lua programming*. Wiley, New York, NY, USA, 2007. ISBN (paperback), 0-470-06917-1 (paperback). 644 (est.) pp. LCCN QA76.73.L82 J96 2007. URL <http://www.loc.gov/catdir/toc/ecip074/2006036460.html>.
- [28] W. Kahan and Joseph D. Darcy. How Java's floating-point hurts everyone everywhere. Technical report, Department of Mathematics and Department of Electrical Engineering and Computer Science, University of California, Berkeley, Berkeley, CA, USA, June 18, 1998. 80 pp. URL <http://www.cs.berkeley.edu/~wkahan/JAVAhurt.pdf>; <http://www.cs.berkeley.edu/~wkahan/JAVAhurt.ps>.
- [29] Donald E. Knuth. A simple program whose proof isn't. In Feijen et al. [14], chapter 27, pages 233–242. ISBN 0-387-97299-4. LCCN QA76 .B326 1990.

Nelson H. F. Beebe

- This paper discusses the algorithm used in TeX for converting between decimal and scaled fixed-point binary values, and for guaranteeing a minimum number of digits in the decimal representation. See also [7, 8] for decimal to binary conversion, [39, 40] for binary to decimal conversion, and [18] for an alternate proof of Knuth's algorithm.
- [30] Donald E. Knuth. *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*. Addison-Wesley, Reading, MA, USA, third edition, 1997. ISBN 0-201-89684-2. xiii + 762 pp. LCCN QA76.6 .K64 1997. US\$52.75.
 - [31] Donald Ervin Knuth. *MMIXware: A RISC computer for the third millennium*, volume 1750 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1999. ISBN 3-540-66938-8 (softcover). ISSN 0302-9743. viii + 550 pp. LCCN QA76.9.A73 K62 1999.
 - [32] Donald Ervin Knuth. *The art of computer programming: Volume 1, Fascicle 1. MMIX, a RISC computer for the new millennium*. Addison-Wesley, Reading, MA, USA, 2005. ISBN 0-201-85392-2. 134 pp. LCCN QA76.6 .K64 2005.
 - [33] David Monniaux. The pitfalls of verifying floating-point computations. Technical report HAL-00128124, CNRS/École Normale Supérieure, 45, rue d'Ulm 75230 Paris cedex 5, France, June 29, 2007. 44 pp. URL <http://hal.archives-ouvertes.fr/docs/00/15/88/63/PDF/floating-point.pdf>.
 - [34] Amos R. Omondi. *Computer Arithmetic Systems: Algorithms, Architecture, and Implementation*. Prentice-Hall, Upper Saddle River, NJ 07458, USA, 1994. ISBN 0-13-334301-4. xvi + 520 pp. LCCN QA76.9.C62 O46 1994. US\$40.00.
 - [35] Michael Overton. *Numerical Computing with IEEE Floating Point Arithmetic, Including One Theorem, One Rule of Thumb, and One Hundred and One Exercises*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2001. ISBN 0-89871-482-6. xiv + 104 pp. LCCN QA76.9.M35 O94 2001. US\$40.00. URL <http://www.cs.nyu.edu/cs/faculty/overton/book/>; <http://www.siam.org/catalog/mcc07/ot76.htm>.
 - [36] Behrooz Parhami. *Computer Arithmetic: Algorithms and Hardware Designs*. Oxford University Press, Walton Street, Oxford OX2 6DP, UK, 2000. ISBN 0-19-512583-5. xx + 490 pp. LCCN QA76.9.C62P37 1999. US\$85.00.
 - [37] C. Severance. An interview with the old man of floating-point: Reminiscences elicited from William Kahan. World-Wide Web document., 1998. URL <http://www.cs.berkeley.edu/~wkahan/ieee754status/754story.html>. A shortened version appears in [38].
 - [38] Charles Severance. Standards: IEEE 754: An interview with William Kahan. *Computer*, 31(3): 114–115, March 1998. ISSN 0018-9162. URL <http://pdf.computer.org/co/books/co1998/pdf/r3114.pdf>.
 - [39] Guy L. Steele Jr. and Jon L. White. How to print floating-point numbers accurately. *ACM SIGPLAN Notices*, 25(6):112–126, June 1990. ISSN 0362-1340. See also input algorithm in [7, 8], and a faster output algorithm in [6] and [29], IBM S/360 algorithms in [1] for both IEEE 754 and S/360 formats, and a twenty-year retrospective [40]. In electronic mail dated Wed, 27 Jun 1990 11:55:36 EDT, Guy Steele reported that an intrepid pre-SIGPLAN 90 conference implementation of what is stated in the paper revealed 3 mistakes:
 1. Table 5 (page 124):
insert `k <-- 0` after assertion, and also delete `k <-- 0` from Table 6.
 2. Table 9 (page 125):
for `-1:USER!("")`;
substitute `-1:USER!("0")`;
and delete the comment.
 3. Table 10 (page 125):
for `fill(-k, "0")`
substitute `fill(-k-1, "0")`
 - [40] Guy L. Steele Jr. and Jon L. White. Retrospective: How to print floating-point numbers accurately. *ACM SIGPLAN Notices*, 39(4):372–389, April 2004. ISSN 0362-1340. Best of PLDI 1979–1999. Reprint of, and retrospective on, [39].
 - [41] Pat H. Sterbenz. *Floating-point computation*. Prentice-Hall series in automatic computation. Prentice-Hall, Upper Saddle River, NJ 07458, USA, 1973. ISBN 0-13-322495-3. xiv + 316 pp. LCCN QA76.8.I12 S77 1974.
 - [42] David Stevenson. A proposed standard for binary floating-point arithmetic. *Computer*, 14(3):51–62, March 1981. ISSN 0018-9162. See [23, 24].
 - [43] David Stevenson. *A proposed standard for binary floating-point arithmetic: draft 8.0 of IEEE Task P754*. IEEE Computer Society Press, 1109 Spring Street, Suite 300, Silver Spring, MD 20910, USA, 1981. 36 pp. See [23, 24].
 - [44] Shlomo Waser and Michael J. Flynn. *Introduction to Arithmetic for Digital Systems Designers*. Holt, Reinhart, and Winston, New York, NY, USA, 1982. ISBN 0-03-060571-7. xvii + 308 pp. LCCN TK7895 A65 W37 1982. Master copy output on Alphatype CRS high-resolution phototypesetter. This book went to press while the IEEE 754 Floating-Point Standard was still in development; consequently, some of the material on that system was invalidated by the final Standard (1985) [23].
 - [45] Brian A. Wichmann. Surveyor's Forum: "What every computer scientist should know about floating-point arithmetic". *ACM Computing Surveys*, 24(3):319, September 1992. ISSN 0360-0300. See [15, 16, 13].

Appendix A

THE FIRST

This is an appendix. Notice that the `\LaTeX` markup for an appendix is, surprisingly, `\chapter`. The `\appendix` command does not produce a heading; instead, it just changes the numbering style from numeric to alphabetic, and it changes the heading prefix from **CHAPTER** to **APPENDIX**.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah.

Appendix B

THE SECOND

This is an appendix.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
Blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah
blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
blah blah blah blah blah blah blah blah blah. Blah blah blah blah blah blah
blah blah blah blah blah blah blah blah.

Appendix C

THE THIRD

This is an appendix.

There are several books [13, 20, 21, 22, 24, 25, 26, 28, 29, 30, 31] listed in our bibliography.

We also reference several journal articles [1, 2, 3, 5, 9, 10, 11, 14, 15, 16, 17, 18, 19, 23, 32, 33] and three famous doctoral theses of later winners [4, 7, 8] of the Nobel Prize in Physics (1922, 1933, and 1921):

Mention should also be made of this famous computer scientist's first publication [6].

Font metrics are an important, albeit low-level, aspect of typesetting. See the *Adobe Systems* manual about that company’s procedures [27].

The world's first significant computer program may have been that written in 1842 by Lady Augusta Ada Lovelace (1815–1852) for the computation of Bernoulli numbers [17, 19]. She was the assistant to Charles Babbage (1791–1871), and they are the world's first computer programmers. The programming language *Ada* is named after her, and is defined in the ANSI/MIL-STD-1815A Standard; its number commemorates the year of her birth.

We do not discuss mathematical *transforms* much in this dissertation, but you can find that phrase in the index.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.
Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah.

[illegible]

blah blah blah blah blah blah blah blah blah blah blah blah. Blah blah blah blah
 blah blah blah blah blah blah blah blah blah blah blah blah.

Blah blah blah blah blah blah blah blah blah blah blah blah blah blah blah. Blah
 blah blah blah blah blah blah blah blah blah blah blah blah. Blah
 blah blah blah blah blah blah blah blah blah blah blah blah.

C.1 References

- [1] H. P. BABBAGE, *Babbage: Babbage's analytical engine*, Monthly Notices of the Royal Astronomical Society, 70 (1910), pp. 517–526, 645. Reprinted in [28, §2.3].
- [2] N. H. F. BEEBE, *Extending T_EX and METAFONT with floating-point arithmetic*, TUGboat, 28 (2007), pp. 319–328.
- [3] N. H. F. BEEBE AND R. P. C. RODGERS, *<PLOT79>: a comprehensive portable Fortran scientific line graphics system, as applied to biomedical research*, Computers in Biology and Medicine, 19 (1989), pp. 385–402.
- [4] N. H. D. BOHR, *Studier over Metallernes Elektronteori. (Danish) [Studies on the electron theory of metals]*, doktor disputats, Københavns Universitet, København, Danmark, 1911. Afhandling for den filosofiske doktorgrad. [Thesis for the Doctor of Philosophy.].
- [5] W. J. CODY, JR., *Analysis of proposals for the floating-point standard*, Computer, 14 (1981), pp. 63–69.
- [6] E. W. DIJKSTRA, *Functionele beschrijving van de ARRA. (Dutch) [Functional description of the ARRA]*, Tech. Rep. 12, Mathematisch Centrum, Amsterdam, The Netherlands, 1953.
- [7] P. A. M. DIRAC, *Quantum Mechanics*, Ph.D. thesis, Cambridge University, Cambridge, UK, June 1926. According to [13, p. 101], this is the first thesis to be submitted anywhere on the subject of quantum mechanics.
- [8] A. EINSTEIN, *Eine Neue Bestimmung der Moleküldimensionen. (German) [A new determination of molecular dimensions]*, inaugural dissertation, Bern Wyss., Bern, Switzerland, 1905. Published in [9].
- [9] —, *Eine neue Bestimmung der Moleküldimensionen. (German) [A new determination of molecular dimensions]*, Annalen der Physik (1900) (series 4), 324 (1906), pp. 289–306. See corrections [10, 11]. This is a slightly revised version of Einstein's doctoral dissertation [8].
- [10] A. EINSTEIN, *Bemerkung zu meiner Arbeit: Eine Beziehung zwischen dem elastischen Verhalten. (German) [Remark on my paper: "A relationship between the elastic behavior ..."]*, Annalen der Physik (1900) (series 4), 339 (1911), pp. 590–590. See [12].
- [11] A. EINSTEIN, *Berichtigung zu meiner Arbeit: Eine neue Bestimmung der Moleküldimensionen. (German) [Corrections to my work: a new determination of molecular dimensions]*, Annalen der Physik (1900) (series 4), 339 (1911), pp. 591–592. See [9].

- [12] ———, *Eine Beziehung zwischen dem elastischen Verhalten und der spezifischen Wärme bei festen Körpern mit einatomigem Molekül. (German) [A relationship between the elastic behavior and the specific heat of solid bodies with monatomic molecules]*, *Annalen der Physik* (1900) (series 4), 339 (1911), pp. 170–174, 590. See remarks [10, 11].
- [13] G. FARMELO, *The Strangest Man: The Hidden Life of Paul Dirac, Mystic of the Atom*, Basic Books, New York, NY, USA, 2009.
- [14] H. H. GOLDSTINE AND A. GOLDSTINE, *The Electronic Numerical Integrator and Computer (ENIAC)*, *Mathematical Tables and Other Aids to Computation*, 2 (1946), pp. 97–110. Reprinted in [28, §7.7].
- [15] P. HALL AND P. PATIL, *Properties of nonparametric estimators of autocovariance for stationary random fields*, *Probability Theory and Related Fields*, 99 (1994), pp. 399–424.
- [16] J. L. HEILBRON AND T. S. KUHN, *The genesis of the Bohr atom*, *Historical Studies in the Physical Sciences*, 1 (1969), pp. vi, 211–290.
- [17] V. R. HUSKEY AND H. D. HUSKEY, *Lady Lovelace and Charles Babbage*, *Annals of the History of Computing*, 2 (1980), pp. 299–329.
- [18] S. C. JOHNSON AND M. E. LESK, *Language development tools*, *The Bell System Technical Journal*, 57 (1978), pp. 2155–2176.
- [19] E. E. KIM AND B. A. TOOLE, *Ada and the first computer: The collaboration between ada, countess of lovelace, and computer pioneer Charles Babbage resulted in a landmark publication that described how to program the world’s first computer*, *Scientific American*, 280 (1999), pp. 76–81.
- [20] D. E. KNUTH, *The T_EXbook*, vol. A of *Computers and Typesetting*, Addison-Wesley, Reading, MA, USA, 1986.
- [21] ———, *The METAFontbook*, vol. C of *Computers and Typesetting*, Addison-Wesley, Reading, MA, USA, 1986.
- [22] ———, *Digital Typography*, CSLI Publications, Stanford, CA, USA, 1999.
- [23] S. N. LAHIRI, Y. LEE, AND N. CRESSIE, *On asymptotic distribution and asymptotic efficiency of least squares estimators of spatial variogram parameters*, *Journal of Statistical Planning and Inference*, 103 (2002), pp. 65–85.
- [24] L. LAMPORT, *L^AT_EX—A Document Preparation System—User’s Guide and Reference Manual*, Addison-Wesley, Reading, MA, USA, 1985.
- [25] F. MITTELBACH, M. GOOSSENS, J. BRAAMS, D. CARLISLE, C. ROWLEY, C. DETIG, AND J. SCHROD, *The L^AT_EX Companion*, *Tools and Techniques for Computer Typesetting*, Addison-Wesley, Reading, MA, USA, second ed., 2004.
- [26] F. W. J. OLVER, D. W. LOZIER, R. F. BOISVERT, AND C. W. CLARK, eds., *NIST Handbook of Mathematical Functions*, Cambridge University Press, Cambridge, UK, 2010.
- [27] POSTSCRIPT DEVELOPER TOOLS & STRATEGIES GROUP, ADOBE SYSTEMS INC., *Adobe font metric files specification — Version 3.0*, Mountain View, CA, USA, Mar. 1990.

- [28] B. RANDELL, ed., *The Origins of Digital Computers: Selected Papers*, Texts and monographs in computer science, Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., third ed., 1982.
- [29] A. ROBBINS AND N. H. F. BEEBE, *Classic Shell Scripting*, O'Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA 95472, USA, 2005.
- [30] D. SALOMON, *The Advanced T_EXbook*, Springer-Verlag, Berlin, Germany / Heidelberg, Germany / London, UK / etc., 1995.
- [31] S. SINGH, *Fermat's Enigma: The Epic Quest to Solve the World's Greatest Mathematical Problem*, Walker and Company, 435 Hudson Street, New York, NY 10014, USA, 1997.
- [32] R. TAYLOR AND A. WILES, *Ring-theoretic properties of certain Hecke algebras*, *Annals of Mathematics*, 142 (1995), pp. 553–572. This paper is a companion to [33], providing the remedy for the flaw in Wiles' 1993 proof of Fermat's Last Theorem. See also [31].
- [33] A. WILES, *Modular elliptic curves and Fermat's Last Theorem*, *Annals of Mathematics*, 142 (1995), pp. 443–551. This paper contains the bulk of the author's proof of the Taniyama-Shimura conjecture and Fermat's Last Theorem, carried out at Princeton University. The companion paper [32] contains the solution to the flaw discovered in the proof that Wiles announced on June 23, 1993, in Cambridge, England. See also [31]. In March 2014, now Royal Society Research Professor Sir Andrew John Wiles of Oxford University was awarded the prestigious Abel Prize in Mathematics for this proof — an award that also carries a cash prize of six million Norwegian crowns, or about US\$722,000.

BINOMIAL NOMENCLATURE INDEX

A

Alces alces 14, 17, 19, 22
Antilocapra americana 14, 17, 19, 22

B

bobcat (*Lynx rufus*) 19, 22
 bongo (*Tragelaphus eurycerus*) 19, 22
Bos taurus indicus 14, 22

C

Cervus canadensis 14, 19, 22
Connochaetes gnou 14, 37, 39, 42
Puma concolor 19, 22

E

E. coli bacterium 14, 37, 39, 42
 elk (*Cervus canadensis*) 19, 22

J

jaguar (*Panthera onca*) 19, 22

L

Lama glama 14, 19, 22
 llama (*Lama glama*) 19, 22
 lynx (*Lynx canadensis*) 19, 22
Lynx canadensis 19, 22
Lynx rufus 22

M

manatee (*Trichechus inunguis*) 17, 19, 22
 margay (*Leopardus wiedii*) 17, 19, 22
 marmot (*Marmota marmota*) 17, 19, 22
Marmota marmota 17, 19, 22
Marmota monax 14, 22
 moose (*Alces alces*) 17, 19, 22

O

ocelot (*Leopardus pardalis*) 17, 19, 22
Odocoileus virginianus 14, 22

P

Panthera onca 19, 22

Panthera tigris 14, 22
 pronghorn (*Antilocapra americana*) 17, 19, 22

T

Panthera tigris) 14, 22
Tragelaphus eurycerus 19, 22

V

Vicugna vicugna 14, 22
 vicuña (*Vicugna vicugna*) 14, 22

W

wapiti (*Cervus canadensis*) 14
 white-tailed deer (*Odocoileus virginianus*)
 14, 22
 woodchuck (*Marmota monax*) 14, 22

Z

zebu (cattle) (*Bos taurus indicus*) 14

FREE SOFTWARE INDEX

A

a2ps 6
acct 6
acm 6
adns 6
alive 6
anubis 6
apl 6
archimedes 6
aris 6
aspell 6
auctex 6
autoconf 6
autoconf-archive 6
autogen 6
automake 6
avl 6

B

ballandpaddle 6
barcode 6
bash 6
bayonne 6
bc 6
binutils 6
bison 6
bool 6
bpel2owfn 6

C

c-graph 6
ccaudio 6
ccd2cue 6
ccrtp 6
ccscript 7
cfengine 7
cflow 7
cgicc 7
chess 7
cim 7
classpath 7

classpathx 7
clisp 7
combine 7
commoncpp 7
complexity 7
config 7
coreutils 7
cpio 7
cppi 7
cssc 7
cursynth 7

D

dap 7
datamash 7
ddd 7
ddrescue 7
dejagnu 7
denemo 7
dico 7
diction 7
diffutils 7
dionysus 7
direvent 7
dismal 7
dominion 7

E

easejs 7
ed 7
edma 7
electric 7
emacs 7
emms 7
enscript 7

F

fdisk 7
ferret 7
findutils 7
fisicalab 7

flex 7
fontutils 7
freedink 7
freefont 7
freeipmi 7

G

gama 7
garpd 7
gawk 7
gcal 7
gcc 7
gcide 7
gcl 7
gcompris 7
gdb 7
gdbm 7
gengen 7
gengetopt 7
gettext 7
gforth 7
ggradebook 7
ghostscript 7
gift 7
gleem 7
glibc 7
global 7
glpk 7
gmp 7
gnash 7
gnats 7
gnatsweb 7
gnu-c-manual 7
gnu-crypto 7
gnu-pw-mgr 7
gnubatch 7
gnubik 7
gnucap 7
gnucobol 7
gnudos 7
gnue 7
gnugo 7
gnuit 7
gnujump 7
gnukart 7
gnumach 7
gnun 7
gnunet 7

gnupod 7
gnuprologjava 7
gnuradio 7
gnurobots 7
gnuschool 7
gnushogi 7
gnusound 7
gnuspeech 7
gnuspool 7
gnustep 7
gnutls 7
gnutrition 7
gnuzilla 7
goptical 7
gperf 7
gprolog 7
greg 7
grep 7
groff 7
grub 7
gsasl 7
gsegrafix 7
gsl 7
gslip 7
gsrc 7
gss 7
gtypist 7
guile 7
guile-gnome 7
guile-gtk 7
guile-ncurses 7
guile-opengl 7
guile-rpc 7
guile-sdl 7
gv 7
gvpe 7
gxmessage 7
gzip 7

H

halifax 7
health 7
hello 7
help2man 7
hp2xx 7
httptunnel 7
hurd 7
hyperbole 7

I

idutils 7
ignuit 7
indent 7
inetutils 7
intlfonts 7

J

jacal 7
jel 7
jwhois 7

K

kawa 7

L

less 7
libcdio 7
libextractor 7
libffcall 7
libiconv 7
libidn 7
libmatheval 7
libmicrohttpd 7
librejs 7
libsigsegv 7
libtasn1 7
libtool 7
libunistring 7
libxmi 7
lightning 7
lilypond 7
liquidwar6 7
lsh 7

M

m4 7
macchanger 7
mailman 7
mailutils 7
make 7
marst 7
maverik 7
mc 7
mcron 7
mcsim 7
mdk 7
metahtml 7

mifluz 7
mig 7
miscfiles 7
mit-scheme 7
moe 7
motti 7
mpc 7
mpfr 7
mpria 7
mtools 7
myserver 7

N

nano 7
ncurses 7
nettle 7
non-gnu 7

O

ocrad 7
octave 7
oleo 7
orgadoc 7
osip 7

P

paperclips 7
parallel 7
parted 7
patch 7
pem 7
pexec 7
phantom 7
pies 7
plotutils 7
proxyknife 7
pspp 7
psychosynth 7
pth 7
pyconfigure 7

R

radius 7
rcs 7
readline 7
recutils 7
reftex 7
remotecontrol 7

rottlog 7
rpge 7
rush 7

S

sather 7
sauce 7
savannah 7
scm 7
screen 7
sed 7
serveez 7
sharutils 7
shishi 7
shmm 7
shtool 7
sipwitch 7
slib 7
smalltalk 7
solfege 7
spacechart 7
spell 7
sqltutor 7
src-highlite 7
stow 7
superopt 7
swbis 7

T

tar 7
termcap 7
termutils 7
teseq 7
teximpatient 7
texinfo 7
thales 7
time 7
tramp 7
trueprint 7

U

unifont 7
units 7
unrtf 7
userv 7
uucp 7

V

vc-dwim 7

vcdimager 7
vera 7

W

wb 7
wdiff 7
websocket4j 7
wget 7
which 7
windows 7

X

xaos 7
xboard 7
xhippo 7
xlogmaster 7
xnee 7
xorriso 7

Z

zile 7

TOPIC INDEX

A

aardvark 19, 22
aardwolf 19, 22
addax 19, 22
African ungulate *see* gnu
Alces alces 14, *see* moose, 17, 19, 22
antelope 19, 22
Antilocapra americana 14, *see* pronghorn, 17, 19, 22
aoudad 19, 22
azymous (unleavened) 19, 22

B

Babbage, Charles (1791–1871) 41
beaver 19, 22
bison 19, 22
bobcat (*Lynx rufus*) 19, 22
bongo (*Tragelaphus eurycerus*) 19, 22
Bos taurus indicus 14, 22

C

caribou 19, 22
Cervus canadensis 14, *see* elk, *see* wapiti, 19, 22
cheetah 19, 22
Connochaetes gnou 14, *see* gnu, 37, 39, 42
Puma concolor) 19, 22
coyote 19, 22
crocodile 19, 22

D

DCT *see* discrete cosine transform
deer 19, 22
discrete cosine transform 2
DWT *see* discrete wavelet transform

E

E. coli bacterium 14, 37, 39, 42
elephant 19, 22
elk (*Cervus canadensis*) 19, 22
Encapsulated PostScript (EPS) 1

Escherichia coli *see* *E. coli*

F

ferret 19, 22
fox 19, 22
free software *see also* GNU Project

G

gazelle 19, 22
gecko 19, 22
gila monster 19, 22
giraffe 19, 22
gnu 14, 17, 19, 22, 37, 39, 42
 diet 14, 17, 19, 39, 42
 dry season 14, 17, 19, 39, 42
 wet season 14, 17, 19, 39, 42
hair 14, 37, 42
 color 14, 37, 42
 DNA analysis 14, 37, 42
 texture 14, 37, 42
 thickness 14, 37, 42
nonpredators
 aardvarks 22
 blesbok 22
 elephants 22
 gazelles 22
 giraffes 22
 zebras 22
predators 14, 19, 22, 39, 42
 crocodiles 14, 19, 22, 39, 42
 hyenas 14, 19, 22, 39, 42
 lions 14, 19, 22, 39, 42
gnu hair 14, 37, 39, 42
GNU Project 14, 37, 39, 42

H

hartebeest 19, 22
hippopotamus 19, 22
hyena 19, 22

I

ibex 19, 22

impala 19, 22

J

jackal 19, 22
 jackass 19, 22
 jackrabbit 19, 22
 jaguar (*Panthera onca*) 19, 22
 jerboa 19, 22

K

kangaroo 19, 22
 koala 19, 22
 kudu 19, 22

L

Lama glama 14, *see* llama, 19, 22
 lemming 19, 22
 lemur 19, 22
 leopard 19, 22
Leopardus pardalis *see* ocelot
Leopardus wiedii *see* margay
 lion 19, 22
 llama (*Lama glama*) 19, 22
 Lovelace, Lady Augusta Ada (1815–1852)
 41
 lynx (*Lynx canadensis*) 19, 22
Lynx canadensis *see* lynx, 19, 22
Lynx rufus *see* bobcat, 22

M

mammoth 17, 19, 22
 manatee (*Trichechus inunguis*) 17, 19, 22
 margay (*Leopardus wiedii*) 17, 19, 22
 marmot (*Marmota marmota*) 17, 19, 22
Marmota marmota 17, 19, 22
Marmota monax 14, *see* woodchuck, 22
 mastiff 17, 19, 22
 mastodon 17, 19, 22
 moose (*Alces alces*) 17, 19, 22
 μ (mu) *see* Borel measure
 musk ox 17, 19, 22
 muskrat 17, 19, 22

N

narwhal 17, 19, 22
 nautilus 17, 19, 22
 Neanderthal 17, 19, 22
 nilgai 17, 19, 22

O

ocelot (*Leopardus pardalis*) 17, 19, 22
 octopus 17, 19, 22
Odocoileus virginianus 14, 22
 okapi 17, 19, 22
 opossum 17, 19, 22

P

pachyderm *see* elephant, *see*
 hippopotamus, *see* rhinoceros
 panda 17, 19, 22
 panther 17, 19, 22
Panthera onca *see* jaguar, 19, 22
Panthera tigris 14, *see* tiger, 22
 peccary 17, 19, 22
 Portable Document Format (PDF) 1
 pronghorn (*Antilocapra americana*) 17, 19, 22
 pterodactyl 17, 19, 22
 puffin 17, 19, 22
 puma *see* cougar
Puma concolor *see* cougar

Q

quagga 17, 19, 22
 quail 17, 19, 22
 Queensland viper 17, 19, 22

R

raccoon 17, 19, 22
 rat 17, 19, 22
 rhea 17, 19, 22
 Rhine berry 17, 19, 22
 Rhinegrave 17, 19, 22
 Rhineland 17, 19, 22
 rhinestone 17, 19, 22
 rhino *see* rhinoceros
 rhinocerite 17, 19, 22
 rhinoceros 17, 19, 22
 rhinoceros horn 17, 19, 22
 rhinodon 17, 19, 22
 rhinodont 17, 19, 22

S

saber-toothed cat 14, 22
 salamander 14, 22
 sapajou 14, 22
 skink 14, 22
 skunk 14, 22

sunfish 14, 22
 suslik 14, 22
 swordfish 14, 22

T

tangun 14, 22
 tapir 14, 22
Panthera tigris) 14, 22
 toucan 14, 22
Tragelaphus eurycerus see bongo, 19, 22
 transform 41, see also Discrete DCT
 Transform, see also Fast Fourier
 Transform
Trichechus inunguis see manatee
 tuna 14, 22
 turbot 14, 22

U

ungulate see gnu, see impala, see kudu, see
 springbok
 unicorn 14, 22

V

vampire bat 14, 22
Vicugna vicugna 14, see vicuña, 22
 vicuña (*Vicugna vicugna*) 14, 22
 vulture 14, 22

W

wallaby 14, 22
 walrus 14, 22
 wanderoo 14, 22
 wapiti 22
 wapiti (*Cervus canadensis*) 14
 wart hog 14, 22
 water buffalo 14, 22
 weasel 14, 22
 whale 14, 22
 whippet 14, 22
 white-tailed deer (*Odocoileus virginianus*)
 14, 22
 wildebeest see gnu
 wolf 14, 22
 wolverine 14, 22
 woodchuck (*Marmota monax*) 14, 22

X

X ray 14, 22

X-radiation 14, 22
 xenon (noble gas) 14, 22
 xylem (woody tissue of a plant) 14, 22
 xylophone 14, 22

Y

yak 14, 22
 yucca 14, 22

Z

zebra 14, 22
 zebu (cattle) 22
 zebu (cattle) (*Bos taurus indicus*) 14
 zoetrope 14, 22
 zythum (ancient Egyptian malt beverage)
 14, 22