

* * * * *

Letters et alia

* * * * *

To the Editor:

I wonder how much serious interest there is amongst TUGboat readers in radically simplifying and speeding up the capture and presentation of complex mathematical text. I have designed an unexpected 'mutation' of the standard Qwerty keyboard which, in mock-up form, allowed—with generous allowance for making and correcting mistakes—simulating the capture of a standard book page full of equations with double levels sub/super-scripts, fractions, differentials and integrals, etc., in a few seconds over twelve minutes. Commercial typesetters estimate at least thirty minutes would be needed even on the most sophisticated equipment presently available. I have also found new and unexpected principles for displaying several typesizes and typefaces in an immediately and exactly recognisable way (without using graphics) that help increase error awareness. The speed of the keyboard comes from its physical design, rather than other 'tricks'; people quite inexperienced in typesetting have remarked how easy it is to use and set complex work that would have completely puzzled them beforehand—it is good for use by relatively unskilled people.

I would like to set up production for several kinds of units if there is sufficient interest of a serious nature. However, until the advent of \TeX the worldwide interest in special maths terminals has been minuscule; it is now uncertain, but probably greater. That is why I would like to hear of interest from TUGboat readers.

There are several useful units in mind. An "Idiot Keyboard" for those well versed in \TeX , that merely generates editable code and/or \TeX code as a direct output; a "Semi-Intelligent Keyboard", basically an 'Idiot Keyboard' with a video output and able to interact with a host; an "Intelligent Keyboard", with a monitor and two floppy discs, able to read editable code and generate \TeX code from that, and assistance facilities (e.g., spelling as well), also interacting with a host; a "Realistic Display Maths Terminal", having the performance of an 'Intelligent Keyboard', but giving realistic display of maths, that is also fully editable on screen (no coding shown) and showing typefaces and sizes in an immediately recognisable way, also assistance facilities and interaction with a host; a "T&M Host" also having a "Realistic Position Maths Display" (graphics) and a proofing printer. A 'Realistic Position Maths Display' ter-

minal would allow any maths office typist (or perhaps any other) easily to set the most complex material and automatically generate and edit \TeX code, for example.

I would very much like to hear from people who have a real interest in such units and/or systems. The nature of the response to this enquiry will very much determine whether there is significant interest. Items would be expected to give substantial savings in salaries of dedicated programmers, and compare well with currently available equipment.

J. M. Cole
17 St Mary's Mount
Leyburn
North Yorkshire DL8 5JB, England

* * * * *

FORMATTING A BOOK WITH \TeX : EXPERIENCES AND OBSERVATIONS

Michael Sannella
MIT Laboratory for Computer Science

In January of 1980 I was hired by Professors Harold Abelson and Andrea diSessa of MIT to format a book they had written: *Turtle Geometry, The Computer as a Medium for Exploring Mathematics*. They had used a computer to write and edit the book—the text was stored on-line—and they wanted me to use \TeX to produce the final formatted copy. In part, they hoped that using a computer formatter would be cheaper than traditional typesetting. However, they were also interested in the experiment of publishing a computer-formatted book, and curious about the ways that computer formatting systems could change the relationship between authors and publishers, giving an author more control over a book.

When I was hired, I didn't know anything about \TeX , or about book publishing. In the process of formatting the book (which took more than a year) I learned a lot about \TeX , about the problems associated with book-quality formatting, and about the interaction of computer formatters with the world of book publishing. This article is an attempt to record my experiences formatting *Turtle Geometry* and some thoughts I have had concerning computer formatters in general. I hope that information will be useful to other people who are trying to format large documents, such as books, using computer formatters. I will try to talk about general problems I encountered rather than about the details of how I fixed each problem. Also, I will try to be as simple as possible—you won't need to know anything about \TeX to read this article.

The first part of this article deals with the experiences I had formatting *Turtle Geometry*. Interacting with the people at MIT Press, learning TeX, and using TeX to format the book all presented problems. Many problems stemmed from the fact that I was using TeX, an extremely complex formatting program. However, all of my problems couldn't be blamed on TeX. *Turtle Geometry* was a very large, very complex document, with innumerable difficult figures, equations, etc. It would have been difficult to format this book using any method. Eventually, I dealt with all of these difficulties, and produced an extremely good-looking book. (*Turtle Geometry, The Computer as a Medium for Exploring Mathematics*, MIT Press, June 1981)

Since I finished formatting *Turtle Geometry* I have been thinking about computer text formatters. In the second part of this article, I discuss some of my ideas on the design of text formatting systems. My experiences formatting *Turtle Geometry* give me an interesting perspective for looking at this area. First, I talk about the problem of designing formatting languages which give the user the correct level of control over the formatting process. Then, I discuss the problem of integrating high-level aesthetic formatting goals into a computer formatting system. I don't claim to have any great answers to these problems—I just want to point out some problems that future text formatters will have to deal with.

I would appreciate receiving any comments that anyone may have concerning this article. I can be contacted by sending U.S. mail to:

Michael Sannella
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Or (for those of you with access to the ARPAnet) by sending computer mail to MJS at net site MIT-AI.

Formatting *Turtle Geometry*

I worked on this project part-time from January, 1980, through March, 1981. My job was defined very simply: to produce a camera-ready copy of the book, using TeX. Aside from that, nothing was specified. I was confident that this goal was achievable, since Knuth had designed TeX specifically to format his books. However, I knew that I had a lot to learn about producing book-quality copy. During the fifteen months I worked on this project, I had a good opportunity to observe the interaction between the two worlds of computer text formatting and book publishing. Computer formatters are being used increasingly for large commercial formatting projects, so it is necessary to investigate how they

can be used most effectively. I hope that this report of my experiences formatting *Turtle Geometry* will be useful to those people actively investigating this area.

DEALING WITH MIT PRESS

One of the first things I did after being asked to format *Turtle Geometry* was to talk with the people at MIT Press (the publishers of the book) and hear their views on the project. The book editor and the book designer were not computer scientists—they were mainly interested in producing a book. However, they were intrigued by my idea of using a computer formatting system, in contrast to traditional typesetting, and were willing to participate in the experiment of formatting the book by computer. They made a few things clear: first, the quality of *Turtle Geometry* must not suffer as a result of doing the formatting by computer. In the past, a few people have published books (mostly computer science textbooks) formatted using primitive text formatters and low quality printers, which looked horrible, and were impossible to read. MIT Press was not willing to publish a book of low quality. I was expected to produce computer-generated output fine enough so that, when it was photographed and printed, the book it would be virtually indistinguishable from a book typeset using the normal methods. They were not willing to sacrifice quality for the sake of our experiment. A second condition was that they did not want to put more work into this book than they would have put in if it was typeset normally. There would be little incentive to use the computer if it took more work on the part of the publisher. Within these restrictions, however, they were willing to be flexible. They were interested in seeing what we could do using a computer formatting system.

The book designer gave me a set of design specifications, written on a standard form. These included such information as the page dimensions, the placement of headings and page numbers, the type to be used with different types of text, etc. The design specifications had been chosen so as to make the book as beautiful and readable as possible. Book designers use their experience with book formatting and a long-cultivated sense of aesthetics, to choose a combination of design specifications such that the total effect is pleasing to the eye, as well as appropriate for the type of book being published. This is the art of book designing.

Turtle Geometry is a textbook dealing with mathematics, physics, and computer science. It contains mathematical equations, computer program listings, and figures, interspersed with paragraphs of text. The book designer decided how

of these objects should be formatted, so that the book would be readable, and would have the right "style." Each of these objects seemed to have fairly simple formatting—there were no fancy footnotes or other objects requiring complicated formatting—so I didn't anticipate any problems with the formatting of individual objects. However, the book as a whole was sure to present some problems, because there were literally hundreds of figures and equations that had to be formatted. I hoped that using a computer formatting system would allow me to cope with this problem better. In retrospect, formatting individual equations and figures didn't present many problems. The hard problems arose when I had to deal with the formatting of the book as a whole. In part this is because \TeX can handle low-level details beautifully, but cannot deal directly with higher-level formatting concepts. I will discuss this problem more thoroughly later.

LEARNING \TeX

After talking with the book designer, I had to learn about \TeX . I knew that Donald Knuth at Stanford University had developed \TeX in order to format his series of books, *The Art of Computer Programming*. Supposedly, \TeX was a book-quality typesetting system particularly optimized for formatting mathematics. So I picked up a copy of Knuth's \TeX manual and proceeded to read it.

My first impression was that \TeX was very complicated. Knuth's manual tries to present the language "gently," a little at a time, so that the reader is not overwhelmed, but it only succeeds in being confusing. This style only served to make it difficult to use the manual as a quick reference. Eventually, I wound up reading the manual cover-to-cover three times, until I had a good idea about how the whole \TeX system worked.

\TeX can be used to produce very high-quality formatted output on a high-resolution printer. However, this comes at a price. The \TeX user has to specify exactly how he wants to format his document using a complicated formatting language. \TeX is an improvement over many other formatters in that it uses some of the vocabulary and concepts of printing and typesetting, but it still requires the user to specify the formatting in excruciating detail. In general, I would characterize \TeX by saying that its "output" is good, but its "input" is bad. It is possible to produce beautifully-formatted copy, but the specifications necessary to achieve this are very complex.

Rather than explicitly specifying the exact formatting for every object in a document, \TeX provides a macro facility which allows one to en-

capsulate and name a sequence of commonly-used formatting specifications, and to refer to them by using the macro name. \TeX 's macros are implemented using simple text substitution—when you reference a macro, \TeX acts as if the macro text were substituted for the macro reference. Macros can be defined to take arguments, which are inserted into the text of the macro in the same way that the macro is inserted into the text file. It was obvious that I should use macros to specify all of the different objects that would appear in *Turtle Geometry*. I began creating a file of macros. Initially, they were very simple, and I was able to get them working very quickly. I gradually refined them, changing the macro definitions until they fulfilled the specifications for the book.

I had a lot of problems using macros because of the way that they are implemented in \TeX . In most cases, simple text-substitution worked well, but sometimes I wished that it were possible to write "smarter" macros. One problem was that the exact effect of a set of formatting commands depends on where they occur. A \TeX macro has no way of detecting where it has been placed, so you can get very strange results by putting a macro in the wrong place. A similar problem occurred because the arguments to a macro are blindly inserted into the macro text, and there is no way to check whether they are appropriate. If the arguments are not of the correct form, the macro may do something quite unexpected. (I give an example of this second problem below, in the section on formatting figures.) Both of these problems happen as a result of people making mistakes—using macros incorrectly—so it could be claimed that this is not a deficiency of \TeX . However, this situation influences the way that people use \TeX . In order to use a macro correctly, you need to know exactly where it can be used, and the precise form of all of its arguments. Therefore, all macros need to be supplied with extensive documentation concerning their use. This is the same documentation problem that people have to deal with when maintaining programs written in computer languages, raised to a higher degree. Most computer languages allow programs to be somewhat isolated from each other, with only minimal communication between them—this allows one to consider the effects of one program apart from the others. Macros in \TeX get much of their meaning from the context within which they are inserted. A macro is meaningless by itself. Consequently, this makes the tough problem of documentation even tougher. While I was developing the macros for *Turtle Geometry*, I found that I had to write exten-

sive notes to remind myself exactly where and how to use my macros. In spite of this, I made a lot of mistakes which were difficult to discover and to fix. Luckily, the macros would not have to be used by anyone else. It would not have been easy to write macros that anyone could use. In general, I would not want anyone else to use my macros. Unless they understood them completely, it would just be too easy to make a mistake. Because of this problem, I predict that it will be difficult to set up “libraries” of commonly-used TeX macros, in an effort to prevent duplication of effort. Instead, most people using TeX will have to start from scratch, and write their own macros.

One unfortunate characteristic of TeX is that it is difficult to debug your formatting specifications if there is some type of error. TeX has a large set of error messages, but these only signal when you specify formatting illegally. Most of the time, TeX formats your document without complaining, and there is something strange in the output. What can you do? Usually you can figure out where something went wrong, but sometimes this is very difficult. This problem is compounded by the use of macros, which may do unpredictable things if you use them incorrectly. This makes TeX very hard to use.

FORMATTING THE TEXT OF *Turtle Geometry*

Once I learned how to use TeX and developed the formatting macros, I began to edit the text of the book, inserting the macro invocations where appropriate. The book was first written and printed using a primitive text formatter called Tj6, so the text files (one per chapter) contained text interspersed with Tj6 commands. All of the Tj6 commands had to be removed, and TeX commands had to be inserted. I eventually deleted all of the Tj6 commands using a text editor, then scanned the straight text which remained, and inserted TeX commands as needed.

Most of the time formatting the text was no problem—the job was tedious, but not difficult. Occasionally I came across an object that had to be formatted somewhat strangely, and I had to do a lot of thinking and experimentation to make TeX do what I wanted. However, with most objects all I had to do was to insert the appropriate macro invocations, and TeX would format them correctly.

Using TeX's macro facility to specify formatting turned out to be very useful when I was not exactly sure how something should be formatted. For example, *Turtle Geometry* contained many vector equations. The authors and the book designer were not exactly sure how they wanted to represent vector variables. So I defined a macro called `\vect`,

and used it to specify all vector variables. I printed out a section of the book several times, defining `\vect` a different way each time. By changing that single macro definition, I could represent vectors in an entire section by letters with a dot on top of them, or by boldface characters. The way to judge which notation was better was to see how it looked in print. Using macros allowed me to experiment with very little effort.

FORMATTING FIGURES

Of all of the objects in *Turtle Geometry*, the hardest to format were the figures. MIT Press and the authors had decided that the figures would be drawn by a technical illustrator, and pasted into the camera-ready copy of the formatted document. Therefore, spaces had to be left for the drawings. Formatting a figure proved more complicated than simply leaving a specified amount of space on a page, however. The book designer had specified that each single figure should contain a caption, and a variable number of subfigures, each with an optional label, organized one or two or three in a row on the page. I had to develop macros to allow me to construct complex figure boxes.

One problem I ran into was determining exactly how much space to leave for a particular figure. If I put in too much space, the drawing would be surrounded by emptiness. If I didn't leave enough space, on the other hand, it would be impossible to paste the drawing into the final copy of the document. Given a drawing, I measured its height and gave its height as an argument to a figure-generating macro. In most cases, this produced a good-looking figure, with just enough blank space to paste the appropriate drawing. However, there were a number of figures which didn't have labels for their subfigures. Because of the way that TeX macros were implemented, there is no way for a macro to detect when it is given a null argument. When given a label argument, my figure formatting macros searched for the located space for a label, which resulted in a figure with too much space. Eventually, I solved this problem by subtracting an appropriate amount from the size I gave as an argument to the figure macro to compensate for the extra space added for the label. This was a rather inelegant solution. In retrospect, it would have been better for me to have defined separate macros for figures with labels and figures without labels. The best solution, if TeX had had the capability, would have been to define a figure macro that could detect when it is given a null label, and format the figure accordingly. This could be greatly improved by extending its macro facility so that a macro could test its arguments.

PAGE BREAKING AND HIGH-LEVEL FORMATTING

As I have mentioned before, I didn't encounter many problems with formatting most of the individual objects in the book. To format a figure, or a program listing, or an equation, I just had to use the appropriate macro from the set I developed. Admittedly, constructing and debugging that set of macros was a slow and painful process, but I only had to do that once. The real problems began after I had formatted the individual objects and was trying to bring them all together. The book designer at MIT Press had given me certain aesthetic rules to follow when formatting pages. Unfortunately, \TeX was not designed to deal with such high-level formatting concepts. As an example, consider the placement of figures within a document. Each figure is referred to at a specific point within the text. The readability of a document is improved if figures are placed so as to minimize the amount of page-flipping a reader has to go through to associate references with figures. For this reason, the book designer at MIT Press specified that a figure should be placed at the top of the page containing its reference, if possible. If that was not possible, the next best choice would be to have the figure on the facing page, to the left or to the right. If that too was not possible, then the figure should be on the next possible page later on in the document. Unfortunately, \TeX cannot deal with figure placement in these terms. For one thing, there is no notion of "facing pages" in \TeX . This is a serious failing, since commercial publishers design document formatting in terms of how each "spread" (facing left and right pages) looks. More generally, \TeX does not give the user very much control over which page a figure is placed upon.

Another formatting rule I was given was that the pages should all be of the same length, if possible, but a page could be a line short or a line long, if that would help the formatting of the page. In \TeX , the height of a page is set to a specific distance using a `\vsize` command. Page breaks are generated to fit as closely as possible to that limit, without going over. There is no provision in \TeX for changing this number dynamically, on a page by page basis, subservient to other formatting needs. Within \TeX all formatting decisions have equal weight, which sometimes unnecessarily restricts the possibilities for formatting something, and reduces \TeX 's power.

If I only had to deal with a small, simple document, the aesthetic goals mentioned above would not be very important. It is not difficult to place a figure in a small document consisting of a few pages, and \TeX would probably position it correctly, us-

ing its simple figure-placement algorithms. However, *Turtle Geometry* was a very large document, with many figures and equations. When you have a large number of figures in a document, placing these figures becomes a much more complicated problem, and it is harder for \TeX to place them correctly. It can't deal with the aesthetic concepts that come into play with large documents. The first time I formatted *Turtle Geometry* using \TeX , it formatted the individual figures and equations correctly but the total effect was horrendous. The book was unreadable—of much lower quality than you would get using traditional human typesetting.

I had to find some way to correct this situation. First, I tried coercing \TeX into doing the correct page formatting by adding formatting commands at places where it broke pages and placed figures incorrectly. This proved to be very difficult. The problem was that I could easily prevent \TeX from breaking a page at a particular bad place, but it was harder to ensure that it would break it at a good place. I was never sure of the effect of adding any particular formatting command until I saw the text reformatted. Coercing \TeX involved endless experimentation, just to get around the inadequacies of \TeX 's page formatting algorithm. I realized quickly that this method was unsuitable. I was spending my time fighting \TeX , trying to merge its page formatting with my aesthetics, rather than worrying about the formatting of the book.

Rather than having both \TeX and myself handle the page breaking and figure placement, I decided to simplify things by doing all of the page breaking by myself. To be exact, I inserted an `\eject` command every place that I wanted a page break, and specified to \TeX that it shouldn't do any page breaks itself. So that I would not have to figure out by myself where all of the page breaks should go, I developed a macro called `\bookmark` which I inserted into each file containing a chapter of the book. The `\bookmark` macro told \TeX that all page breaks before the point where it occurs will be handled by explicit `\eject` commands, but that \TeX should attempt to generate page breaks from that point on. I eventually developed a routine for handling the page breaks in a chapter: I would insert a few `\eject` commands, move the `\bookmark` macro call forward to the end of my `\ejects`, and run the file through \TeX . Seeing how \TeX would format the next few pages helped me decide where I wanted to insert page breaks. Often, I didn't have to do anything more than insert `\eject` macros where \TeX had broken the page before, but most of the time I wanted to move figures around, and change the

exact placement of the page breaks. It was necessary to do this work a few pages at a time, because each formatting change I made could send effects percolating through the rest of the document in unpredictable ways.

By doing the page breaking and figure placement myself, I was able to optimize the page formatting with respect to aesthetic rules that I could not express with \TeX . Consequently, I produced a book of very high quality.

LOOKING BACK

Formatting *Turtle Geometry* was a long, difficult job. Everything did not go as smoothly as it may seem from what I have said above. One problem was that the macro development, the figure formatting, and the editing of the actual text of the book were all happening at the same time. It was very difficult to coordinate these separate activities. Another problem, which is sure to occur whenever a document is available on a computer system, was that of "freezing" the text of the document. When it is easy to change the text, authors are very reluctant to finish editing. There is always a great temptation to fix one more error. At some point, however, it is necessary to decide that the text will not change any more.

A typical example of the organizational problems I had to deal with was that associated with putting the figures in the book. It took a long time to get the figures drawn and checked. Until the drawings were in final form, I couldn't measure them and give these measurements to the figure macros. Therefore I couldn't do the page formatting, since the sizes of the figures critically affected the location of page breaks. To cope with this situation, I worked on each chapter separately, so that at any one time different chapters could be at different stages of production—the authors might still be editing the text of chapter 5 while I was sizing the figures for chapter 1, and formatting its pages.

I made many mistakes, and had to do some things several times. I was learning what to do, and developing new methods for doing things, while I was formatting the book. Everything would have been much better if all of the people working on *Turtle Geometry* had sat down before the work was started, and developed a routine for producing the book. We could have avoided a lot of problems, and produced a better book with less effort, if we had been more organized.

Thoughts on Text Formatters

Current computer text formatting systems suffer from many deficiencies, which become increas-

ingly apparent as people attempt to use them in more and more situations. Some of these problems are due to formatting languages that cannot cope with the aesthetic ideas that book designers have developed from long experience with document production. Other problems result because formatters are designed with small documents in mind and cannot deal with the formatting concerns of large documents. I don't claim to have "the solution" to all of these formatting problems. However, I can come up with a few ideas about how text formatting could be improved.

FORMATTING LANGUAGES AND CONTROLS

Computer-controlled printers deal with documents at a very low level. To a graphics-oriented printer, a document is a set of commands which specify the exact positions of individual dots, lines, and characters on a page. Obviously, people do not want to have to specify the formatting of their documents in such excruciating detail. Instead, they use a computer text formatter, which could be considered as a translator between an abstract, high-level formatting language and the specific detailed language of a printer. A formatter allows its user to specify document formatting in terms which are more abstract and general, such as words and paragraphs and lines of text. After giving these specifications, the formatter decides what each character of the document should be placed on each page.

An important characteristic of a formatting language is the level of detail with which the user has to specify the formatting of a document. Suppose a formatting language deals only with high-level formatting concepts, and the user wishes to specify the formatting of a particular object a little more explicitly. This could happen if the formatter does not format something exactly right automatically, and the user wants to fix this. Another possibility is that the user may wish to create a one-time special object, which needs to be formatted strangely, and the user is willing to take the extra time to specify it in detail. (As an example of this, in *Turtle Geometry* there were only two or three tables, and it was necessary to specify them in detail, rather than develop general macros for formatting tables.) In this case, the only solution is for the user to try to convince the system into formatting correctly. This may involve tricking the formatter, making it treat special objects like other objects, and trying to second-guess the formatting algorithm. This is very difficult to do, it requires an intimate knowledge of the internal workings of the formatting system which most users may not have, and it is extremely inelegant.

problem is really that the way such formatters are designed, there are simply some documents which cannot be produced. In the interest of making a formatter easier to use, the set of possible documents that could have been formatted with it has been constrained.

Now, let us consider the opposite extreme, where you have a very low-level formatting language, and you wish to do high-level formatting. In other words, you don't care exactly how your document looks, as long as it conforms to certain high-level formatting goals, such as having lines filled to a constant length, etc. Using a low-level formatting language, it should be possible to produce a document conforming to these high-level constraints, but you will be required to specify a lot of low-level details that you are not really interested in. Aside from the inconvenience of specifying the formatting in this much detail, this makes things very difficult if you wish to change your document later. The more detailed your specifications, the more work has to be done in order to change them. When you work with specifications at a low level of detail, you essentially have to calculate the formatting yourself, rather than letting a computer do it.

There is another, rather subtle, problem associated with formatting languages which use descriptions at a lower level than you need for a particular document. The greater the detail in which you specify the formatting of a document, the more you constrain the possible ways that it could be formatted. If you have to use a low-level language, you wind up making a lot of fairly arbitrary decisions about how something should be formatted at a low level, since you have to specify something. However, all of the formatting specifications are heeded by a computer formatting system, whether they are important to the user, or whether they were just specified arbitrarily. The more constrained a formatter is, the smaller the space of possible ways a document could be formatted, and the more likely that there will be no formatting which is acceptable. When a user is specifying how his document should be formatted, he should use the least possible amount of detail, while still mentioning those things important to him. In this way, the possibility is increased that the formatter can find an acceptable way to format his document.

Early text formatters were designed around the idea that documents were fairly simple, consisting of text organized into paragraphs, with an occasional heading or footnote. These programs would take the user's specifications, and format a document using rules internal to the formatter. It was possible

to modify exactly how the formatting was done by specifying the values of certain parameters, but the formatting was essentially done by one "smart" program. These early formatters present a good example of formatting languages which work at too high a level for many formatting tasks. As long as the shape of your document conformed to the ideas explicit in the formatter's abstract paragraphs and words, it would be easy to format your document. However, if your document contained something out of the ordinary that the program was not equipped to handle, there was little that you could do. No matter how much effort you were willing to expend to specify the detail, these early formatting languages wouldn't allow you to. The user simply had very little control over the formatting process. The shape of a document had to be specified in terms of the objects the formatter was built to deal with, then the user would just have to "throw it to the winds," and hope the formatter produced something acceptable.

\TeX is similar to these early formatters, in that it deals with abstract ideas such as paragraphs and pages, and the user only has indirect control over how these objects are formatted. However, \TeX also can deal with the idea of "boxes and glue," which allow a user to specify an object at a very low, but manageable, level. This facility is very useful for handling "special case" objects, such as complex figures, which would be too complicated for \TeX to deal with as built-in objects. In essence, \TeX allows the user to deal with formatting at two distinct levels of detail, or control. Either you can "throw your formatting to the wind" and let \TeX take care of it all, or you can specify the formatting of the object in low-level detail. This is an improvement over the earlier text formatters, which gave you no choice, but it still proves inadequate in many situations. The problem is, a user sometimes want to have control over the formatting, without having to specify *all* of the formatting explicitly. For example, consider the problem I had while formatting the Turtle Geometry book with breaking pages, and placing figures. When I let \TeX do all of the formatting by itself, it would do it incorrectly, in part because it doesn't consider some of the aesthetic considerations I am concerned with. For a while, I tried tricking \TeX 's formatting algorithms into doing the correct thing, but this involved going through all sorts of complicated contortions. Finally, I was forced to deal with the problem at a very low level, by specifying all of the page breaks and figure placements myself, explicitly. I was able to get what I wanted, but it required a lot of work

on my part. Another issue in this situation was that once I had done the page breaking by hand, that essentially “froze” that part of the document. If any changes had needed to be made to the document after I had done the page breaks, it would have been necessary for me to do all of that work over again. It would have been better if I had the ability in \TeX to specify formatting at a higher level, while still retaining control over the formatting.

People who design text formatting systems need to worry about this issue. One way that formatters could be improved is by giving more control to the user, at different levels of detail, so that it is not necessary to specify document formatting at more than the minimum amount of specificity. Looking at \TeX in particular, one way to improve it would be to introduce “smarter” boxes—objects that are not quite as explicit and low-level as boxes are in \TeX . For example, you could have an object called a “paragraph box,” which contains text that should be formatted as a paragraph. Eventually, this may not be formatted as a single rectangular box—the paragraph may eventually extend across several pages. Without explicitly specifying the exact formatting of every character in the paragraph, you could express parameters appropriate to the level of abstraction of a paragraph, such as line length, degree of raggedness, whether the paragraph can be broken, etc. To format any particular object in a document, simply express it in terms of a formatter object which deals with the correct level of abstraction.

One problem with implementing this scheme is that it could lead to a complicated formatting language with an unmanageably large number of objects defined. A way to deal with this is to only define a few objects, which can be used to specify formatting at differing levels of detail. For example, consider a “box” object that acts in the following way: given a set of objects as arguments, it formats them like a paragraph, using default width, raggedness, etc. values. If these values are given as parameters to the box, it will use them instead. You can imagine a whole set of parameters and defaults set up so that a single box object could be used for formatting a paragraph (at the most abstract level) or used exactly like a box in \TeX (if everything, including vertical and horizontal sizes, are specified). A single object, with many parameters and defaults, could be used at many different levels of abstraction.

HIGH-LEVEL FORMATTING GOALS

One problem I have mentioned regarding current text formatters is that they do not deal with some of the aesthetic rules that book designers have

developed after years of experience with form documents. I had to explicitly specify page breaking for *Turtle Geometry*, because \TeX could not deal with some of these concepts. Let us examine this problem more closely. When breaking pages by hand, I had to keep in mind several aesthetic considerations. First, figures had to be positioned as near to their references in the text as possible (on the same page, or on the facing page, or on the next page). Second, I didn't want a page break in the middle of a computer program listing, if possible. Third, the length of each page should be nearly constant, plus or minus a line. In certain special situations, I had to worry about other aesthetic considerations, such as not having an equation directly under a figure, but this didn't come up often. Mainly, I was concerned with choosing page breaks, and positioning figures, so as to produce a document optimized with respect to those three aesthetic rules given above.

Certain problems arise when dealing with high-level aesthetic formatting goals. It is difficult to explicitly specify such goals in a form that a computer can work with. Consider what a human has to go through to format a document while considering these goals. It is seldom the case, given a text formatting situation with a lot of figures, that all of these formatting goals are compatible. Often they are contradictory. For example, suppose you have a page with a computer program listing in the middle, and a figure reference near the top of the page. Suppose that if you place the figure at the top of the page, this pushes the text on the page down far enough so that the computer program listing is broken between this page and the next page. Here you have a conflict between the rule that figures should be on the same page as their references and the rule that doesn't want computer program listings to be broken. Exactly what I would do depends on the exact situation. If the page in question is a left-hand one, it might be possible to move the figure to the facing right-hand page. Alternatively, if the place where the program was to be broken is only a few lines into the program, I might just make that page short. (notice that in this second case, I may be breaking the third aesthetic rule.) Making an aesthetic decision of this sort involves finding a compromise between several contradictory goals.

The problem of formatting a document with respect to certain aesthetic formatting goals can be reduced to the problem of considering two alternative solutions, where a part of a document is formatted in two different ways, and deciding which is “better.” Usually there are only a few possibilities to con-

in any formatting problem, and a computer could afford to try the various possibilities, and pick the "best" one. The key to this problem is obviously determining what the word "better" means with respect to a set of high-level aesthetic formatting goals. It is instructive to examine a mechanism that \TeX has for dealing with this problem on a lower level. In order to optimally generate line breaks and page breaks within a paragraph, \TeX associates an integer with each paragraph known as its "badness." Every line break generates a certain "penalty," another integer, which indicates how "bad" that break is. Line breaks have large penalties associated with them if the break requires hyphenating a word, or if a line is too long, among other things. Page breaks in a paragraph have large penalties associated with them if they generate widows (the first or last line in the paragraph alone on a separate page). The badness of a paragraph is calculated as the sum of all of the penalties in the paragraph. The formatting algorithms inside \TeX are designed to make line and page breaks so as to minimize the total badness of a paragraph. This is a fairly flexible system—users can specify penalties explicitly at certain points to indicate places that are particularly good or bad to break at—and it formats individual paragraphs very well. However, this type of system would not be suitable for dealing with more complex formatting constraints. "Badness" is a simple, one-dimensional "score" which can be easily calculated and used for comparisons, but it is not powerful enough to express some of the more complex combinations of higher-level formatting goals. Consider the aesthetic goals given above, dealing with figure placement and computer program listings. You can't simply take a

badness score with respect to one goal, and a badness score with respect to another, and sum them to get total badness. This is like adding apples and oranges. Aesthetic formatting goals combine in more complex ways than that. In a particular situation, it might be possible to represent a situation in terms of carefully-adjusted penalties, but this wouldn't work in general. You need a more complex, multidimensional, way of comparing two situations.

The fundamental problem with implementing high-level abstract formatting goals in a text formatting system is the problem of achieving a goal. Supposing that you have a way of unambiguously stating fairly complex aesthetic rules, how is the computer to make sure that they are achieved? Exactly what can the formatter do if the formatting it chooses disobeys one of the aesthetic rules? Consider the low-level decisions that a formatter makes when formatting a document. The decisions about exactly which word to break a line at, or exactly how much to stretch a line, are fairly arbitrary. Unless the user specifically specifies it, the formatter has a fair amount of leeway as to exactly how the document is formatted. Ideally, the low-level arbitrary decisions could be influenced by the high-level aesthetic rules. I admit that I don't know how this could be accomplished, but this is obviously what text formatters should work towards. As an intermediate step, formatting systems should be designed that can express the concepts of high-level aesthetic goals. Even if they couldn't use these goals to influence the formatting process, it would be worthwhile simply to flag violations of these rules, pointing out to the user places in the document where one or another of the rules are broken. This by itself would lead to better formatters, producing better-looking documents.