Principles

¶

Book Design for TEX Users Part 1: Theory

Philip Taylor

Abstract

Book design cannot be taught; it can only be learned, preferably by critical study of as many books as possible. Of all the elements which make up a book, white space is frequently the least considered and the most important. *Avant garde* designs are compared and contrasted with more conservative and traditional approaches. Three key elements: *uniformity*, *information* and *structure* are identified, and 'good design practice' discussed in terms of each of these.

Keywords. Design, typography, layout

motto: There can never be too little space below headings, only too much!

Introduction

The widespread use of TEX and other typesetting or DTP packages by tens of thousands of scientists, researchers and other academics has resulted in two rather disturbing phenomena: (1) more and more people are spending ever longer trying to get their publications to *look* right, rather than worrying about whether such publications are factually correct or are well written, and (2) fewer and fewer people, on opening a book for the first time, think first about the content, but instead commence by judging the book on its form, or to be more precise, on the appearance of the design and typesetting. We are, in fact, becoming a generation of self-taught designers and typographers, but in so doing we are tacitly avoiding the many years of training, apprenticeship and indenture which previous generations have deemed necessary.

This is, in itself, no bad thing — there are far too many self-appointed 'experts' ever ready to initiate neophytes into the arcane mysteries of their craft, in exchange for not inconsiderable sums of money — but in order for learning by osmosis to be effective, the beginner has to be exposed both to good and to bad examples of the art, and to think critically about what it is that differentiates the former from the latter. In Departments of Typography and Design, such examples abound, and the professors daily compare and contrast good with bad to the great benefit of their students; but in the incestuous world of TEX, good examples are rare whilst bad examples abound.

But why should this be? What is it about TEX, which in skilled hands is capable of producing results equalling the very best examples of hot lead composition, that encourages the production of second- and even third-rate designs? I suggest that there are two main answers to this: (1) in *The TEXbook*, which is presumably the first (if not the only) book on typesetting that users of TEX [1] encounter, there is extraordinarily little guidance given concerning document design, as opposed to document formatting, and (2) the standard styles which accompany LATEX generate results which even the staunchest LATEX adherent would have difficulty arguing represent "the state of the art" in document design, and which if considered dispassionately might justifiably be said to lack subtlety and finesse. [2]

Thus the lack of explicit guidance, together with the rather poor examples generated by the standard LATEX styles, has resulted in a proliferation of poorly designed books all of which shriek "TEX" (or "LATEX"). [3] It would not be fair on the authors to adduce particular examples of this creed of mediocrity, but a glance at any reasonably complete library of TEX-related (or even TEX-set) books will shew what I mean...

However, all is not lost: a new generation of TEX-setters are emerging who appear to have studied the typesetter's craft, and several of the more recent books on TEX shew every evidence of having been *designed*, rather than having been ripped untimely from their progenitor's womb.

1: as opposed to LATEX.

2: The Dutch, always sensitive to such issues, have produced a sub-style 'Sober' which attempts to tone down the worst excesses of the default LATEX styles.

3: Knuth, in his closing exhortation, wrote: "GO FORTH now and create *masterpieces of the publishing art*." Nowhere, so far as I can trace, did he write: "and let every one of them shriek 'TEX' from every page"...

In this paper, then, I propose to discuss what it is that differentiates a well-designed book from one that is ill-designed (or worse, one that has not been designed at all); and in so doing, I hope that I can in some small way contribute to a more universal adoption of TEX. For all the while that professional publishing houses see only bad examples of TEX-set books, they are unlikely to consider adopting TEX as a house standard; but if the general standard of TEX-set books can be raised to a point where they are either indistinguishable from, or even better than, books produced by traditional means or by commercial typesetting packages, then simple economics will ensure that such publishing houses give TEX the consideration it so richly deserves.

The Book

We all know what a book is, for we handle them every day; assuming a Western culture, it is basically a set of uniformly sized sheets of paper, joined at their left edge in some way, sandwiched between two slightly larger sheets of a more rigid or robust material that wraps around the left edge. It is differentiated from a magazine primarily by virtue of its cover: that of a magazine is only slightly more robust (although frequently more glossy) than the pages which it protects, whilst that of a book is almost invariably either thicker, or more rigid, or both; a magazine cover, too, has only one basic point of articulation, whilst most book covers articulate independently of each other. One other aspect separates the all but the thinnest book from all but the thickest magazine: a book is typically bound in *signatures*, whilst a magazine is usually stapled through its spine as a single entity.

But open a book and open a magazine, and we see that these differences are only superficial; for there are far more fundamental differences which manifest themselves once inside the cover. The magazine is characterised by variation — each page is clearly different from the preceding and the next; whilst the book is characterised by uniformity — each page, seen from a distance, is virtually indistinguishable from the next (special pages apart). And in this uniformity lies the basis of successful book design; for readers have come to *expect* this uniformity, and anything which detracts from it will serve only to distract the reader.

Yet uniformity of itself is not enough: we could achieve uniformity by leaving each page blank, or by simply placing a large black rectangle within the margins of each page; but this will not satisfy our reader, who looks not only for uniformity but for *information*. Indeed, information is the very *raison d'être* of a book: without it, the book serves no purpose at all, and is at best a work of art (and at worst is totally valueless).

So the book exists to supply information; and anything which inhibits or interrupts the flow of information from book to reader will diminish its value. If the flow of information is too badly affected, the reader will simply cast the book aside (how many of us, on attempting to read a page of reversed-out Bodoni in some otherwise traditional magazine, have simply given up and left the material unread? I have, many times, and cursed the designer for his/her stupidity in putting form before function).

Uniformity, information: what else? Well, if the book is in any sense *technical* (by which I exclude the novel but include almost everything else), then it is also *structured* (indeed, as we shall see, even a novel is structured in many senses, but not in the one which I am using here); and, possessing structure, it is capable of being accessed in a structured manner. It will have, at the very least, a table of contents; it *should* have an index (although far too many books that would benefit enormously from an index are lacking in this respect), and it may also possess an internal structure, in that the reader may be asked from time to time to *see Chapter 3*, or *see also Section 2.4.2*.

And these three elements, I suggest, lie at the heart of successful design: uniformity, information, and structure. We will look at each in turn to see how it may best be achieved, implemented or accomplished.

Uniformity

Take a book (a traditional book, not one hot off the presses of the DTP revolution), and flick the pages, rather as if there were one of those old-fashioned animated cartoon characters lurking in the corner of each page. What do you see? Most people perceive a regular grey grid: not black and white — you only see that if you look at a static page — but instead a grey blur where the text appears, and white where there is no text (or other material). What is significant is that the white appears at the same place on every page: above the headline, below the footline, between the head/footlines and the body of the text, and to left and right of the text, in the margins of the page. If the book is set in multiple columns (usually two, but rarely more, except for rather specialised works), then a further block (or blocks) of white space will appear, separating the columns from each other.

And in many senses, this white space is the most important of the graphic elements which will go to make up each page. It provides the framework or matrix within which the 'dark matter' — the text, graphics, etc., which make up the *information* content of the page — is set. But probably because it does not of itself appear to carry any information, it is frequently afforded less respect than it deserves, particularly by those undertaking design without any formal background. And yet, although it does not *appear* to carry any information, in fact it carries a great deal: without it, we would not know where the headline stopped and the page body began; where the page body stopped, and the footline began; where the left column stopped and the right began, and so on… In fact, it is *vital* to our comprehension of the contents of the page, and is therefore *at least* as important as every other element on the page, if not more so.

Because the white space and the dark matter are inextricably interlinked — one starts wherever the other stops, until the physical limits of the page are reached — any discussion of the uniformity of white space must equally be interlinked with a discussion of the uniformity of the dark matter of the page. But there is a third element to this uniformity which is even more dependent upon the inextricable interlinking of white space and dark matter, and that is the sense of 'greyness' of every page. The human eye is remarkably sensitive to small variations in grey level, and if the apparent greyness varies either within a page or between pages (particularly between facing pages, which form a *spread*), the effect can be quite discomforting. Such variations in apparent greyness can result from a variety of causes, of which the most common are: (a) use of letterspacing for justification; (b) inconsistent leading between two or more blocks of text in the same font; (c) inappropriate changes of leading or font (or both) when deliberately setting a block of text in a different font (in a multi-line quotation, for example). The cures for each of these ills are fairly straightforward: *never* use letterspacing to achieve justification, unless its use is so subtle that the eye cannot perceive the variation in inter-letter spacing; *never* allow the typesetting system to vary the leading in order to achieve vertical justification (and never set two blocks of text in the same font but with a different leading without being aware of the effect which will be achieved); and be aware of the perceived grey-level (white-matter : dark-matter ratio) when setting blocks of text in different fonts.

In an ideal world, attention to the suggestions of the previous paragraph would do much to ensure that the apparent greyness of each page was uniform: but there is another problem which results from our less-than-perfect world which can also significantly affect perceived greyness, and this is the problem of 'print through'. Ideal paper presents a uniform opaque whiteness on which the black of the ink is superimposed; real paper, on the other hand, whilst uniformly white (at least, as far as high-quality printing papers are concerned) is rather less than opaque; if held up to a bright light, even the best paper will allow some light to shine through, and poorer papers are so translucent that printed material can be read almost as easily from the back as from the front (albeit as a mirror image). This in itself would be no problem were it not for the fact that the two sides of each sheet are logically independent entities: not only are they printed in separate operations, but the material appearing on one side bears little or no correlation with the material appearing on the other. However, in *designing* such pages, the effect of their back-to-back nature must be borne in mind, and a good design will attempt to ensure that each line of text on the obverse is matched by another line of text on the reverse. Of course, in practice this is not achievable; sections break up the flow of the text, as do illustrations and other graphics: but it must be the *intention* of the designer to achieve this line-for-line equivalence, and on this philosophy is predicated the whole concept of the *grid*.

The grid represents an abstract model of each page; special pages (e.g. chapter openings) may be afforded a special grid of their own, but normal 'running' pages will each use the same grid, onto which is mapped the various elements of the page. The grid can be perceived in a hierarchical manner: at its most superficial, it will have lines for the physical limits of the page, for the left and right limits of the text (or of the columns, if a multi-column work), for the upper and lower limits of the page body, and for the headline and footline. At the next level of refinement, the page body will be divided into lines of text (which is why most traditional specifications for books express the dimensions of the page body in terms of lines of text, rather than so many picas or so many inches or centimetres). Superimposed back-to-back, two of these grids will intermesh perfectly, each line of text on one side corresponding to a line of text on the other; the left edge of the text on the obverse will exactly align with the right edge of the text on the reverse, and so on (which has implications for the margins, as we shall see).

Of course, the grid is an ideal, but conformity to the grid must at times be allowed to be violated; if this were not the case, there would only be a finite number of positions at which a heading (for example) could appear above the text which it introduces: one line, two lines, three lines, etc. But such granularity is far too coarse for the æsthetic demands of real book design, and headings therefore need to be treated as special cases, allowed to float away from their 'natural' grid line whilst the paragraphs of text above and below the heading remain bound to the grid. Illustrations and graphics, too, must be treated as special cases, and float within the white space equivalent to an integral number of lines of text, thereby themselves being independent of the grid whilst leaving their surrounding paragraphs locked firmly in place.

But sometimes the requirements of page makeup will dictate that a particular page be underfull: a paragraph, for example, may finish in such a way that there is insufficient room (e.g., only a single blank line) for a new paragraph to start; or there may be room for a heading but not for a heading plus post-heading vertical white space plus at least two lines of text. How, in those circumstances, can the contradictory requirements of uniformity and page makeup be reconciled? In the limiting case, there is no general solution which will always work, and practical (real-life) books may occasionally have to violate one or other constraint; but equally often there is a solution which is both elegant and æsthetically pleasing: violate, by the same amount, the constraint of uniformity for two facing pages (i.e., for a *spread*). If, for example, the verso (left-hand) page runs one line short, then *force* the recto (right-hand) to run one line short also; if the verso page would ideally run one line long, then allow it so to do, but require the matching recto page to run one line long also.

And in this concept of balancing the *spread*, as opposed to achieving uniformity between *all* pages, lies, I believe the essence of good design. For when the book is held open in the hand, or laid open on the desk or lectern, it is not a single page that is seen at all, but a double-page spread; and if the two facing pages of the spread appear uniform (uniform in greyness or visual density; uniform in placement of headline and footline; uniform in size of margins — outer margins the same size as each other, inner margins also the same size but not necessarily the same apparent[4] size as the outer; and uniform in terms of grid-lock, in that verso lines appear in perfect vertical alignment with their recto counterparts) and balanced (with both verso and recto page bodies starting at the same height from the bottom of the page, and extending for the same depth), then much will have been accomplished; and if this same uniformity and balance can be carried through every spread of the book (thereby avoiding problems of print-through and so on), then much of the framework of good design will already be in place.

But there are many practical problems associated with the concepts of gridlock and balanced spreads; some of these are particularly true when using TeX, whilst others are more general. Those that are particular to TeX will be addressed in the sequel to this paper[5], whilst those that are more general are discussed below.

Considering first the problems of balanced spreads: it was suggested above that if the natural height for a verso page was one line short, or one line long, then it should be set to its natural size and its counterpart recto page *forced* to the same size. But what if the verso page naturally sets at the target size of the page, whilst the recto page runs one line light or one line over? If page makeup is performed on a page-by-page basis, then it is already too late to re-set the verso page, and either the recto page will have to be set to a non-natural size (if

---

4: I use the term *apparent* here quite intentionally, for as we shall see, the apparent size of the inner margins is always less than their actual size, by an amount which is a function both of the thickness of the book and of the binding technique(s) used.

5: Book Design for TeX Users; Part 2: Practice, to appear in Zeszyt 6.

there is sufficient flexibility in the page makeup to allow this), or the balance constraint for the spread violated. And therefore we must postulate that any typesetting system intended for the production of well-designed books *must* be capable, at the very least, of setting a *spread* as an entity, rather than a page. Of course, this does not let us off the hook completely: for example, if the verso page naturally runs one line over, but the recto page finishes a paragraph at the natural height for the page, then it may not be possible to graft an additional line onto the recto page without violating some other (tacit) constraint; in these circumstances it may be necessary to backtrack even further, and to start asking questions such as "what if I were to set the preceding spread one line light, or one line over", and so on; in the final analysis, the more decisions about page makeup that can be deferred, the better the final volume is likely to be. As computer memory becomes cheaper and cheaper, it is by no means unreasonable to think about optimising a complete chapter at a time.

And what of uniformity: what if a multiline quotation, set in a smaller font with correspondingly reduced leading, *must* appear as an entity on a page, whilst there is no matching quotation on the other (physical) side of the same page? Then print-through will undoubtedly occur for the duration of the quotation, and in the worst case there will be an interference effect as the lines of the quotation drift into and out of synchronism with the lines of paragraph text on the other side. Here no matter how much material we accumulate can a makeup solution be postulated; and in the end we are dependent more on the skills of the paper maker in achieving near opacity than we are on our own skills in using and programming our typesetting system.

But there is much more to uniformity than simply gridlock and balanced spreads. Uniformity is a concept which percolates every element of good book design. Consider, for example, the treatment of chapter headings, section headings, paragraphs, quotations and so on: in what sense can they, too, be made 'uniform'? Clearly each must be unique, in order for the reader to immediately identify at what sort of entity he or she is looking; yet if they are not only unique but are also afforded wildly disparate typographic treatment, then any sense of coherence is lost and the book starts to take on the appearance of a mismatched hotchpotch of design ideas.

We might start by positing that there should only be a small number of different fonts used — 'the fewer the better' is hard to equal as an axiom for the selection of fonts! — whilst equally there should only be a small number of placements. For example, if paragraphs are fully justified (as would usually be the case for a book, although exceptions to this rule will be discussed elsewhere) and if section headings are set ranged left, then the book as a whole should probably restrict itself to these two styles of setting: it would normally be inappropriate to have centered headings in a book that otherwise has a fully-justified or ranged-left theme running through it. But if section headings are set ranged left (perhaps in conjunction with semantic line breaks[6]), whilst normal paragraphs are set fully justified, then quotations could either be set fully justified (like paragraphs) or ranged left (like section headings), but should probably not be set ranged right without good reason.

And what of indentation? Here two different schools of thought obtain. One would argue that the requirement of uniformity encompasses indentation, and that the indentation, once chosen, should apply to the whole book: thus, for example, lists would be indented by the same amount as paragraphs; quotations might be set with an additional left margin equal to this indentation; and the bibliography might be set with reverse indentation also equal to this same amount. The other would say that the requirements of clarity and lack of ambiguity dictate that a *different* indentation should be used wherever different entities occur, thereby giving the reader maximum indication of the nature of the entity being indented even on the most superficial glance at the page. I have sympathy with both points of view, but my inherently conservative background renders the former more appealing than the latter; I do not think I have yet seen an example in which the reader could have been mislead had a uniform indentation been adopted. But this whole area transcends the boundary between *uniformity* (which suggests a uniform indentation), and *information* (which suggests

6: A concept whereby a ragged-right setting is used in conjunction with 'strongly recommended' line breaks, thereby ensuring that complete ideas (phrases, clauses, sentences, etc.) are not unnecessarily split over two lines.

different indentations for different purposes), and brings us naturally to the next section.

Information

The primary function of any book is to convey information; yet the preceding discussion has concentrated almost entirely on the æsthetics of book design, rather than on its rôle as a medium for the communication of information. However, provided that the two ideas do not come into conflict, a uniform and æsthetically pleasing appearance does much to assist the book in its communication rôle, for it allows the reader to concentrate on the text (i.e., the *information content* of the book) whilst not being distracted by its design (a fact which is sadly ignored by many of today's more *avant garde* designers). But there comes a point at which excessive adherence to the precept of uniformity would start to detract from the book's primary rôle as information source, and it is therefore to this area that we must now turn our attention.

Consider first of all the rôle of section headers: those single (or occasionally multiple) lines of text which serve to introduce the reader to the ideas which follow. This paper, for example, makes use of only a single level of section header, the author preferring to lapse into straight prose within each section; other authors, particularly those with a strong scientific background, feel happier if they can classify their ideas in a strongly hierarchical manner, and frequently have recourse not only to *A-heads* (as in this paper), but B-heads, C-heads, D-heads and even E-heads on rare occasions. The first requirement for such headers is that they shall, *unambiguously*, refer to the text which follows: it should not be possible, in a well-designed book, to mentally attach them to the preceding text. The means by which this is accomplished is simplicity itself, yet is so often violated in amateur-designed books and other documents that one wonders whether the idea has ever occurred to their designers at all: the section header shall be physically closer to the text which it introduces than to the text which precedes it. Note that this is strictly a 'less than' relationship, not a 'less than or equals' one: the header must *never* be equi-spaced between the preceding and following texts. This rule has some interesting knock-on effects: for example, a header must *never* appear in isolation at the bottom of a page, for were it so to do, it would by definition be nearer to the preceding text than from the text which follows.

But in a strongly hierarchical book or paper, it is just as important that the different levels of header (A-head, B-head, etc.) shall be capable of being differentiated at a glance. How should this hierarchy of headers best be conveyed to the reader? We have available several options: (1) Higher-level headers may be separated from their preceding text by greater amounts of vertical white space than lower-level; (2) Higher-level headers may be separated from their qualifying (following) text by greater amounts of vertical white space than lower order; (3) Higher-level headers may be set in a larger font than lower-level; (4) Higher-level headers may be set in a bolder font than lower-level; (5) Some other typographic differentiation (e.g. the use of a *sans serif* font in a book or document otherwise set in a *serif* font) may be used for one or more levels of header; (6) Run-in headers may be used for the lowest level of header. Indeed, these are only some of the available options: for example, in some works a new page is taken for each new top-level section, even where that section is only one of many similar sections in a chapter.

Clearly the range of options is vast, and it is not possible in a paper of this brevity to give more than a few typical conventions, but one requirement is tantamount: if two or more conventions are adopted within a single document, then no combination of those conventions must lead to ambiguity. For example, if A-heads are set in 16 point roman, B-heads must not be set in 14 point bold, for the boldness of the B-head would counteract the effect of the smaller font and lead to ambiguity in the mind of the reader. Even if a bold font is not explicitly used, it is possible (by, for example, selecting an ill-matched *sans serif* font for B-level headers in an otherwise *serif* document) to accidentally specify an *apparently* bolder font for a subsidiary-level header. Such ambiguities must be avoided.

In what other ways can the book designer ensure that information is most clearly conveyed? Perhaps most important of all by ensuring that the book can be *read*! This goes without saying, you may say, but there are sadly only too many counter-examples already published for this particular requirement to be omitted from any reasonably critical analysis. Perhaps we need to start by defining what we mean by “to read”; I suggest that if reading is to be conducted efficiently and pleasurably, then it must (for the normally-abled adult) be an almost unconscious activity. If I pick up a book hoping to gain information from it, then the *last* thing that I want is to have the designer’s personality forced down my throat (unless it is a book on book design, in which case I may be able to judge from the book’s design whether or not to bother to read it!); the design must therefore be very ‘quiet’ and unintrusive, allowing the content to flow naturally forth through the medium of the form, rather than having the form leap out from the page and distract the reader from the content. Naturally there are exceptions to this rule, and books on design clearly come into that category, being inherently self-referential, but generally speaking the reader wants to know as little about the designer and as much about the content as possible.

Furthermore, reading must be able to proceed in a linear and uninterrupted manner; it is well known that any infelicity on the part of the author which results in ambiguity in the reader’s mind will cause the latter to back-track through the work, hoping to gain further clues and thereby disambiguate the text on a second or subsequent reading. Classic authors on grammar (Fowler, Weseen, Partridge, Onions, Gowers, Quiller-Couch, Sweet) pay much attention to this. But there are many typographic pitfalls which can also cause a reader to have to backtrack, and it is as important for the designer to avoid these as it is for the author to avoid the grammatical infelicities.

For example, during the 1930s, there was a great vogue for *sans serif* faces: they were modern, *avant garde*, stylish, modish — use whatever term you will. And particularly in North America, and to a lesser extent in Europe, such was the pressure to use these typefaces that their *raison d’être* — to provide a simple, minimalist, style for short sections of text which would not draw attention away from the main theme (frequently an accompanying graphic) — were forgotten, and they were advocated (and used) as *the* typefaces for every conceivable purpose. These purposes were not restricted to their classic use in headings, captions, posters, etc., but were instead extended to encompass even the running text of books; every page was set in *sans serif* text, with little feeling for the comfort and convenience of the reader. The effect on the reader was all too predictable (with hindsight): readers found it difficult to concentrate on such books for any period of time, finding it tiring and even distressing; and the reason was very simple, although not well understood at the time: even though the *serifs* which characterise most of our classic typefaces today are in reality no more than artifacts dating back to the original letterforms of stonecutters (particularly in the case of upper-case letterforms), and later of typecutters, these *serifs* perform a very important function when the letterform occurs in running text: they serve to draw the eye naturally along the line of text, very much reducing the risk of the eye vacillating between two adjacent lines of text, and also help to minimise the amount of backtracking within a single line. And so, with the benefit of hindsight and of psychological and physiological research, it has now been established that the typeface of choice for passages of running text (as opposed to captions, etc, which extend for at most a few lines) is almost invariably a *serif* face. Sadly this fact is still occasionally ignored.

But if the choice of a *serif* face is almost mandatory to avoid vacillation between lines of text and backtracking within a single line, what other psychological or physiological factors can also affect the readability of the text? Perhaps the most important of all, and one for which plain TeX sadly gives most inappropriate guidance, is the size of font with respect to the *measure* (i.e., the width) of the text. Plain TeX is predicated on the use of 10 point fonts on a measure of 6.5 inches (39 picas), which simply gives far too many characters per line. Psychologists have shewn that the optimal number of characters per line for normally sighted people lies in the range 40–70, and peaks somewhere near the upper bound of that range; below it, people become frustrated: they are forced to take in too little information per glance; and above it, they tend to lose their place, and either backtrack within the line, or on re-scanning to the start of the next line, lose their vertical place and re-scan to the start of the wrong line. Even LaTeX, which generally gives better guidance than plain TeX in matters of typographic design, allows the user complete freedom to select between 10 point, 11 point and 12 point fonts, regardless of the style chosen and therefore of the measure of the text. For Europeans readers, accustomed to the DIN series of paper sizes, the best guidance I can give is as follows: if you are setting on a sheet of A4 paper (which is unlikely for a book but quite possible for a report or other similar document), with 'normal' margins (circa 1 inch), then a 12 point font is called for; you can get away with 11 point, but 10 point is out of the question. The same goes for North American readers with 1 inch margins on a sheet of American 'letter' paper, $8.5'' \times 11''$. And for a book? Well, 'how big is a book' is a question to which I will return in the sequel to this paper, but generally speaking books *are* set in 10 point typefaces; however, as the width of the paper increases, two columns become obligatory or pathologically large margins become required.[7] In unusually small books, 9 point fonts may be used, but anything less than this poses problems of legibility for normally sighted people.

In the preceding paragraph, I have spoken of a "10 point font" as if it were some sort of ISO standard; but sadly it is anything but. Fonts vary enormously both in their actual size (as measured), and in their perceived size, and the quoted size is at best an approximation and at worst a d@mned lie! For what it is worth, the notional size of a font is that distance which may separate consecutive lines of text in a paragraph set in that font without the descenders of one line overlapping the ascenders of the line below; it is also approximately the height + depth of a parenthesis glyph. But in practice one designer's 10 point font may well be another's 11 point; and if you are using two or more fonts in a single document, then it is your responsibility as designer to ensure that the size at which they are used renders them visually conformable, even if this means loading one at 10 point and another at 11 point (or even at 10.6347 point, if that represents the true ratio between their perceived sizes).

And for the leading: some authorities will suggest "1.2 times the design size of the font"; others will suggest "2 points more than the design size of the font"; and others will suggest yet further formulæ. The answer is, of course, that no one formula will be right for every font, or for every size, and until experience has given you the insight to look at a font sample and *know* the appropriate leading for the target font size, then you will have to use the most powerful tool available to you: your eyes. In other words, you will have to print samples of the text at various leadings (probably of the order of magnitude suggested by the formulæ above), and adjust until it looks right to you. But when you print these samples, you will come up against another, and very subtle, psychological quirk: assume you do as most people do, and print your proofs on a laser printer; then your output will appear either on a sheet of A4, or on a sheet of 'letter' size paper, and most unusually on anything else. And try as you might, you will not be able to judge the size of the font and the size of the leading as they will appear in the final book form, even if you draw a box around your sample text to represent the dimensions of the final trimmed page; your eye/mind will refuse to believe that the white paper which lies outside that line is not attached to the text, and will judge the size of the text and the size of the leading in terms of the untrimmed sheet of A4 or 'letter' paper. The solution, of course, is to guillotine the paper to the final trimmed size, and then to paste two such trimmed sheets together (or to print a double page spread in the first place) and to look at a full-size replica of the final spread of the book; and then, and only then, will you be able properly to judge the size of type and the size of the leading in terms of the printed page.

Structure

Finally we turn our attention to *structure*, and in particular to the means by which a well-designed book can be efficiently referenced (and cross-referenced) in a quasi-random, rather than sequential, manner. At the coarsest level of granularity, a book is divided into volumes (if huge), parts (if large) and chapters (almost all books). Access to volumes need not worry us unduly: each will contain the name and/or number of the volume on the spine and front cover, and only if two or more volumes are concurrently open in front of the reader will it be necessary to be able to differentiate between volumes by inspection of only the open spreads.

Parts are not uncommon, but many of the potential problems associated with the identification of parts can be eliminated by sequential numbering of chapters independent of the part in which they happen to fall; with sequential chapter numbering, the reader can always be referred to *Chapter n*,

7: I am advised by a North American student that it is the practice in North America for students to annotate their books; for this reason, they *expect* far wider margins than European readers, which may explain something about the default LaTeX styles.

without needing to qualify it as *Chapter n of Part m*.

But the most important division of the majority of books is into chapters, and here we must start our investigations into *structure* in earnest. Consider the classic case of a multi-chapter, single-volume, book, with a table of contents ('TOC') among the *front matter* (a.k.a. 'the prelims'). The reader wishing to access the book through the TOC consults the latter and sees, for each chapter, its number, its name (if the chapters are named), and the page on which it commences. Selecting a chapter from those listed, the reader flicks through the pages looking for the page on which the chapter starts. This is not a random search: the page numbers increase monotonically with period 1, and if the reader overshoots he or she is invariably sufficiently familiar with the general concept of a book to realise that it is necessary to backtrack.[8] But an interesting phenomenon occurs as the reader converges on the page of interest, at least in many less-than-optimal books: the page numbers (*folios*, as they are frequently termed) traditionally alternate between top-left and top-right, occupying the top-left placement on verso pages and top-right on recto; this placement is believed to make them maximally visible. But on opening chapter pages it is traditional to suppress the running head ('headline'), because the design of these pages (discussed in greater detail in the sequel to this paper) is such that a running head is generally considered æsthetically displeasing. And therefore the very page which (logically) bears the number sought is also the very page which (physically) has no page number on it; and the reader is forced to perform a narrow binary search to ensure that the page of interest has truly been located, by comparing the last physical page number which can be found (and which will, in the worst case, not even be visible from the page of interest, if the previous chapter happens to finish recto, since it is also traditional to start new chapters recto and a completely blank page will therefore form the verso half of the spread) and the next physical page number, which will invariably also be invisible from the page of interest. Of course, the name and/or number of the chapter will be visible on the sought page, and it will be clear from its design that it *is* an opening chapter page, but none the less the reader who until then has been searching for a specific page number is forced to modify his/her search algorithm.

The solution generally advocated for this problem is to present the page number on opening chapter pages as a *drop(ped) folio*: a centered page number occupying a part of the footline. The percipient reader soon becomes familiar with this convention, and modifies his/her gaze to take in the bottom of the page rather than the top outside edge when reaching an opening chapter page. But if dropped folios are acceptable on opening chapter pages, why not use them consistently throughout the book? This would have two beneficial effects: (1) the reader would be able to find *any* page in the book by studying the same part of every page, regardless of the nature of that page, and (2) additional space would be released in the running heads for additional (cross-)referencing material, space which as we shall see becomes of a premium as the complexity (in terms of explicit hierarchical structure) increases.

Once we have ensured that page numbers occur on *every* page (blank pages excepted, since by definition no possible well-formed (cross-)reference could require the reader to turn to such a page), we have at a stroke ensured that our tables of contents, indexes, etc., all of which generally yield a *page number* when 'dereferenced' (consulted), will invariably result in a hit rather than a miss. We must now turn our attention to other techniques for (cross)-referencing, and in particular methods for locating logical sub-divisions of the book (e.g. sections, sub-sections, etc.) by their *name*, and also by their *number* if such entities are numbered.

Generally speaking, the names and numbers of logical sub-divisions are used for cross-referencing (i.e., referencing from within the text), rather than

---

8: It is interesting to realise that the scenario outlined is the converse of what usually happens in practice: because books are generally either laid on the desk/lectern or held in the right hand with the highest number page at the bottom, it is far more natural for the reader to make a *backwards* search through the pages until the desired page is found, or until overshoot occurs, than it is to make a *forwards* search. This is because it is far easier to raise a number of pages, frequently almost the entire page set, in one hand and allow them to fall back individually under the effect of gravity than it is to lift each page individually whilst seeking the page of interest.

for direct referencing (e.g. from a table of contents or an index); but regardless of the source of the reference, the reader will ultimately be required either to *see Section 2.1.4* or to *see also Lagopus hyperboreus* — in neither case will the reader explicitly be instructed to turn to a specific page. It is frequently possible to convert one of these *indirect* references into a *direct* reference to a page number, by consulting the appropriate table of contents or index, but this two-stage process is both frustrating and time-wasting: a more direct method is required.

The mechanism by which this direct access to named or numbered logical sub-divisions of a text is generally accomplished is through the medium of *running heads*; these have been referred to previously in the current paper without any formal definition being given of their nature or purpose. A running head is so called because it recurs on (almost) every page; opening chapter pages and blank pages are usually excluded from the set of pages on which a running head can occur, and if an entire page is given over to an illustration then that page too may be excluded; but special cases apart, running heads occur on every page. But of course the *content* of the running head varies from page to page: were it not so, there would be no purpose to the running head at all (which is also frequently the case when it is used to echo the title of the book on every page or every second page; the reader is normally aware of the title of the current work, although there are counter-examples, as when consulting many works at once; thus the echoing of the title is not necessarily evidence of poor design). In general, the content of the running head is adjusted to reflect the content of the page over which it appears; thus, for example, if *Section 2.1.4: metalinguistic notions* commenced on page 23, the running head of page 23 would almost certainly reflect that fact. But in a hierarchically structured work, there are potential conflicts; consider a book with chapters, sections and sub-sections: which of these entities should the running head reflect? A convention frequently adopted is to ascribe different semantics to the verso and recto heads: the verso carries 'more significant' information (e.g. the name/number of the chapter), whilst the recto head carries 'less significant' information (e.g. the section name/number). Yet this is not enough: where should the sub-section information appear? Ultimately there is no solution to this problem: if the book is sufficiently complex (i.e., possesses too deep a nesting), then no matter how complex an arrangement of headers is adopted there will be a level of nesting beyond which it is simply not possible to reflect lower-order entities in the header. The designer, then, must perform a trade-off, and decide which information is most beneficial to the reader. Omissions are possible at either or both ends of the spectrum: it may be that knowledge of the name of the current chapter is less important than knowledge of the current section / sub-section / sub-sub-section / sub-sub-sub-section; or it may be that knowledge of the chapter is deemed more important than knowledge of the current sub-sub-etc. The designer and author must work together on this problem.

But there is one additional mechanism which is considerably under-used, yet which allows twice as much information to be packed into each header: if folios are removed to the footline, thereby releasing the outer edge of each running head for other usage, then provided that the author can be encouraged to provide *short* names for each of his/her chapters/sections/etc., each running head can serve double duty. For example, verso heads can carry (left) chapter name/number, whilst carrying (right) section name/number; recto heads can then carry (left) sub-section and (right) sub-sub-section. Adequate space must clearly be left between the two elements to avoid potential ambiguity.

Finally, is it the *name* or the *number* of each logical entity which is to appear in the header? Above I have hedged my bets by consistently referring to name/number, yet at some point a decision must be made. If space allows, and if the author co-operates by providing short names, then there is no reason why *both* should not appear; with less space, or longer names, it may be necessary to omit the numbers in order to allow the names to appear; and if the author is unconscionably prolix in naming the various entities, then the designer may have little choice but to simply give the hierarchical name (e.g. *Chapter*, *Section*) followed by the relevant number. But this last serves the author rather than the reader, and pressure should be brought to bear on the author to provide suitable 'short forms' purely for use in the running heads. Of course, some works use *only* numbered entities; in such works, there is no choice: the hierarchical names (if appropriate) and numbers must be used.