TEXT ALIGNMENT

Many modern text processors, such as LaTeX(which this is written on) use a sophisticated dynamic programming algorithm to assure that the lines are well aligned on the right side of the page. The problem is where to break the text into lines. Let us $l_1, l_2, \ldots, l_n$ denote the lengths of the words of the text. If $l_i, \ldots, l_j$ are placed on a line we assume they take up space $l_i + \ldots + l_j + j - i$, the extra being the space between the words. (In the special case of a single word l_i the length is simply l_i.) Let L denote the total length of the line. (We'll assume: all $l_i \leq L$, we can never have more than space L on a line, and that words can never be cut.) A line with text $l_i, \ldots, l_j$ then has a *gap* $G = L - (l_i + \ldots + l_j + j - i)$ at the end of the line. We'd like, of course, for G to be zero but we can't always [1] get this. We are given a function $P(x)$ and a line with gap G incurs a penalty of $P(G)$. For the sake of argument we will specify $P(x) = x^3$. We'll say a line with gap G has *badness* $P(G) = G^3$. [2] The total badness of a text is the sum of the badnesses of the lines. The object is to split the text into lines so as to minimize the total badnesses.

A natural inclination is to use a *greedy algorithm*: if it fits, put it in. Below is an example (the | represents the end of the line) where this does not work. The words have lengths $3, 4, 1, 6$ and the line length $L = 10$. The greedy algorithm would have $3, 4, 1$ on the first line and 6 on the second with gaps $0, 4$ respectively so total badness $0^3 + 4^3 = 64$. If instead we split $3, 4$ and $1, 6$ the gaps are $2, 2$ and the total badness is $2^3 + 2^3 = 16$, much better[3]

```
   NOW SING  |2  8        NOW SING A|0  0
   A MELODY  |2  8        MELODY    |4 64
 total badness  16                     64
```

[1] You might notice that the text you are reading (and most things written in LaTeXor other modern word processors) seems to be nearly perfectly aligned. But if you look very closely you'll see that the space between letters is not perfectly uniform. When LaTeXhas, say, 3 extra spaces at the end of a line it spreads the space out along the whole line by putting extra space between letters. Exactly how it does that is itself an interesting question, but one we do not persue.

[2] So badnesses go $0, 1, 8, 27, 64, 125, \ldots$. A gap of five (badness 125) is then counted as equivalent to nearly five gaps of three so the algorithm will really try to avoid them. The choice of badness function is a subjective decision guided by aesthetic considerations. The algorithm is given a particular badness function and a text to split into sentences.

[3] In actual application space on the last line is not so bad as space in the middle but we ignore that wrinkle in our presentation.

Now for the algorithm. Set $m(i)$ equal to the total badness of the text $l_1, \ldots, l_i$. We shall find $m(i)$ for $i = 1, 2, \ldots, n$ in increasing order.

`Initialization.` Suppose i is such that $l_1, \ldots, l_i$ fit on one line, i.e., such that $l_1 + \ldots + l_i + i - 1 \leq L$. The best splitting of the text is then to have everything on the same line. This gives a gap $G = L - (l_1 + \ldots + l_i + i - 1)$ and so $m(i) = P(G)$.

`The recursion.` Now look at larger i. We take the i one by one, going up. We loop on i going up to n. For a given i let k range over $i, i-1, i-2, \ldots$ for as long as $l_k, \ldots, l_i$ can fit on one line. For each of those k calculate $m(k-1)$ plus the badness of the line $l_k, \ldots, l_i$. This gives the badness of the text $l_1, \ldots, l_i$ *if* the last line is $l_k, \ldots, l_i$. Pick the k that gives the smallest sum and set $m(i)$ equal to that sum. We can also keep an auxilliary array s setting $s(i) = k$. This has the meaning that in the optimal splitting of text $l_1, \ldots, l_i$ the last line starts with l_k. At the end $m(n)$ denotes the total badness of of the full text. To actually do the splitting we work backwards. The last line goes from word $s(n)$ to word n. Now set $n \to s(n) - 1$ and the penultimate line goes from word $s(n)$ to word n, etc.

How long does this take. There is a loop on i of length n. There is an inner loop on k. Certainly k takes on at most n values so that this is a $O(n^2)$ algorithm. But in many cases we can say more. Suppose u is the maximum number of words that can fit on a line. Then k takes on at most u values so that this is a $O(un)$ algorithm. If we think of the line size as fixed (rather a natural assumption) and the words as having a minimal length (also natural) then u is a fixed number and so this becomes a *linear* algorithm in the size of the text.