Discrete mathematics in industry

David Allwright, Smith Institute

There are so many discrete mathematical problems arising from all industrial sectors that it is difficult to know how to begin to classify them. Naturally they often involve some continuous aspect as well, arising from time or space or the physical variables in the real-world problem. For instance, consider the error-detecting and error-correcting codes that are used in all the digital communications we make each day. The construction of these codes is a purely discrete problem, yet the reason why such codes are needed arises from an underlying continuous problem of signal transmission over a radio link or optical fibre, and the probability distribution of errors in such transmissions. To give another example, a timetabling problem or a resource allocation problem will generally take a purely discrete form, yet it has arisen from real-world continuous constraints such as how long it takes a data packet to get from A to B, or a train, or a class of schoolchildren. So, when I describe some discrete industrial mathematics problems here, some of them also involve continuous variables: but the characteristic feature in each case is that the mathematically challenging part of the problem is essentially discrete. The first two problems are industrial applications of graph theory and illustrate, incidentally, that not all industrial graph theory problems are the travelling salesman problem. The third will be from combinatorial auctions of the kind used by Ofcom for spectrum licences, and the fourth is a regulatory problem to do with aircraft noise.

The two graph theory problems arise from channel assignment. So we are thinking of a network of radio transmitters and receivers communicating with mobile devices in the surrounding area. They might be the base stations of a phone network, communicating with mobile phones in their area, or they might be wireless access points for the WLAN in an office building, communicating with the laptops, printers and so on in the nearby rooms. Each transmitter needs to operate on one of several available "channels", which may, for instance, be different versions of a coding system. Neighbouring transmitters need to use *different* channels so that a device that is perhaps midway between them is able to lock unambiguously onto one of the incoming signals despite that fact that they are incident with similar intensity. There are various other constraints in practice, depending on the particular application, which we may think of as "non-interference" constraints. Finding a good channel assignment, that gives a good quality of service to the users of the system, is an interesting discrete problem in industrial mathematics that has been widely studied in recent years. It is very closely related to graph colouring, for we can think of the transmitters as the nodes of a graph, and join 2 nodes when their cells are geographical neighbours. Then if we think of the available channels as colours, of which we have to assign one to each node, the simple constraint mentioned earlier is that joined vertices must be different colours, *i.e.* we have to produce a vertex-colouring of the graph. However, in 2005 Motorola brought to the UK Maths-in-Industry Study Group this *recolouring* problem: suppose one channel assignment is in operation and we wish to change to another, by changing one node at a time [1]. During this change, when some transmitters have gone to their new channel and others are still on the old, all the non-interference constraints must be satisfied. How do we find an order of carrying out the changes that achieves this ? To give some idea of the scale of this problem, there might be 2000 nodes in the area in question, so the number of possible orders to consider would be 2000! and it is way beyond the wildest dreams of any computer to consider them one by one. Nevertheless, the problem turns out to be tractable by suitable combinations of ingenuity and heuristic methods. One of the important considerations, as often in channel assignment problems, is that the graph is not like a general random graph, because of the underlying geography. If A and C are each neighbours of B then it is much more likely than average that A and C are neighbours. The theory of random *geometric* graphs has been developed to study such situations [3] and there will undoubtedly be further interesting discrete industrial mathematics problems in this area.

The second graph theory problem is the one mentioned by John Ockendon in his IMA Gold Medal Lecture printed in this issue, and concerns the case of a WLAN network in a modern office complex, where many different companies can be working in the same building, each with its own "territory". There are various ways of reducing the risk of loss of security of company-confidential data, one of which is to arrange that walls and ceilings between *different* territories have a special treatment, "stealthy wallpaper", that provides reasonable radio attenuation at the relevant frequencies. However, it is still necessary that companies whose territories are in *face-to-face* contact on opposite sides of just one wall or floor-ceiling should be on *different* channels. Territories that are separated by two walls or floors, or that touch only along an edge or at a corner can share a channel. In general of course, arbitrarily many channels may be needed, but what if each territory is a cuboid aligned with a fixed set of Cartesian axes? It is not difficult to make an example that needs 6 channels, but Bruce Reed was able to show that in fact arbitrarily many channels may still be needed in this case, and the construction goes as follows [2]. Given any integer k we arrange blocks such that if they can be coloured with k colours then all of the colours are needed. We begin with a *staircase* of a_k blocks, called

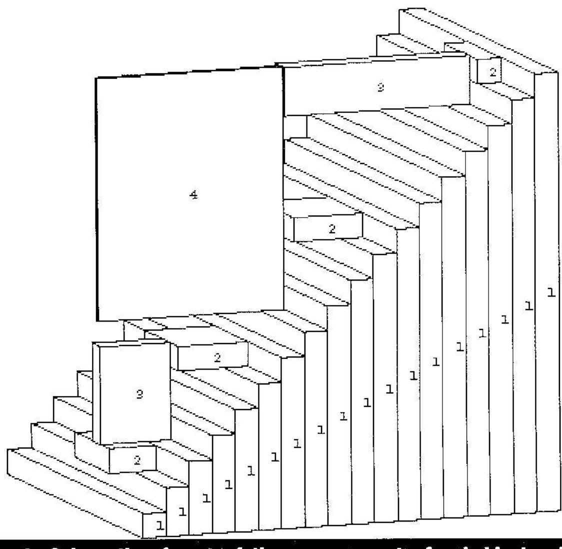

Figure 1: Schematic of part of the arrangement of cuboids to give an adjacency graph needing at least k colours

the 1-blocks, like the blocks labelled 1 in Figure 1. They all have width 1 in the x-direction, length l_1 in the y-direction, and heights $1, 2, \ldots, a_k$ where we shall define $a_r = 2^r r!$ (for each $r = 1, \ldots, k$). We let $b_k = a_k / k$, and a set of b_k of the 1-blocks will be called a 1-set, so there are $n_k = \binom{a_k}{b_k}$ such 1-sets. Our "2-blocks" are going to be of length $l_2 = l_1 / n_k$. For each 1-set, we allocate a section of length l_2 in the y-direction, and arrange $b_k / 2 = a_{k-1}$ 2-blocks as illustrated in the Figure. So, if the 1-set starts with the 1-blocks whose ends are shaded, we take them in order, temporarily label them A,B,A,B,. . . , and for each consecutive pair A,B, we place a 2-block with its right face in contact with B and its left face vertically above the left face of A. The total number of 2-blocks added is therefore $n_k a_{k-1}$, since there are a_{k-1} in each of the n_k sections. The arrangements of the 3-blocks, 4-blocks *etc.* are going to be obtained by a recursive iteration of this construction. So, consider any particular set of a_{k-1} 2-blocks that have been constructed in this way from a particular 1-set. A set of $b_{k-1} = a_{k-1} / (k-1)$ of these 2-blocks will be called a 2-set, so there are $n_{k-1} = \binom{a_{k-1}}{b_{k-1}}$ such 2-sets. We let $l_3 = l_2 / n_{k-1}$ and for each 2-set we allocate a section of length l_3 and place $b_{k-1} / 2 = a_{k-2}$ blocks of Type 3 that are arranged as shown in the Figure. The total number of 3-blocks is therefore $n_k n_{k-1} a_{k-2}$ since for each of the n_k 1-sets there are n_{k-1} 2-sets, and each of those has a_{k-2} 3-blocks placed in its section. This whole recursive refinement process is repeated up to the k-blocks, so the result looks totally unlike any real building, and the blocks become *extremely* thin in the y-direction!

Now we have to describe the consequences of this arrangement for the colouring problem. First, since there are a_k 1-blocks, if they are coloured in k colours, there must be some set of $b_k = a_k / k$ that are all the same colour, which we may call colour 1. Somewhere along the y-direction is the section of length l_2 where those blocks were the 1-set, and therefore the 2-blocks are arranged as shown in the Figure. So the 2-blocks in that section must all avoid colour 1, and are therefore coloured using the $k-1$ colours $2, \ldots, k$. However, since there are a_{k-1} of these 2-blocks, there must be some subset of $b_{k-1} = a_{k-1} / (k-1)$ of them that are all the same colour, which we may call colour 2. Somewhere along our section of length l_2 in the y-direction is the arrangements of 3-blocks that corresponds to this particular monochromatic 2-set. So within that section of length l_3 the 3-blocks are arranged as shown in the Figure. So the 3-blocks in that section must all avoid colours 1 and 2, and are therefore coloured using the $k-2$ colours $3, \ldots, k$. Repeating this argument through each stage of the construction we see that in fact all k colours must be used, as we wished to show.

The number of blocks in this construction is

$$X_k = a_k + n_k a_{k-1} + n_k n_{k-1} a_{k-2} + \ldots = a_k + n_k X_{k-1}, \qquad (1)$$

with $X_1 = a_1 = 2$. One can then show that for $k \geq 6$, $X_k < k^{k^k}$ and so the number of WLAN channels needed for n cuboid territories can grow at least as "fast" as $\log \log n / \log \log \log n$.

A third example of discrete mathematics in industry is in combinatorial auctions, and we can introduce this by contrasting first a Dutch auction and an English auction. Suppose a single item is being sold and there are n potential bidders, who are prepared to pay up to $a_1, a_2, \ldots, a_n$ for it, and we shall label them to make $a_1 < a_2 < \ldots < a_n$. In a Dutch auction, the asking price is gradually reduced until someone makes a bid, so the item will be sold to bidder n for price a_n. In an English auction, the asking price is gradually raised until there is only one bidder, so the item will still be sold to bidder n, but for price a_{n-1}. (In fact $a_{n-1} + \varepsilon$ but we ignore ε for this discussion along with the possibility of ties and the different psychology of the two auctions.) The English auction is thus in a certain sense economically fairer to the buyers: no non-winning bidder can say to the seller, "Look, I would have paid more for the item", but the price is as low as possible subject to this constraint. If it is run as a sealed-bid auction then it is called a second-price sealed-bid auction, or Vickrey auction. Now consider the corresponding problem for a combinatorial auction such as Ofcom is using for spectrum licences. Here the buyers submit sealed bids not just for one possible item at a time, but a whole table of bids for each combination of items they are interested in. The seller then has two combinatorial problems. The first is to find the combination of bids that maximizes the sum of the bid amounts subject to the available numbers of items of the different kinds. The second is to find the prices that are economically fair in the same sense as the English auction. In other words, no subset of the non-winning bidders should be able to say to the seller, "Look, if you had accepted these bids of ours for those items, and the prices that you are accepting from that subset of winning bidders for the remaining items, then your revenue would have been greater." But the prices should be as low as possible subject to this constraint. The need to verify this for different combinations of subsets of bidders can easily lead to a combinatorial explosion unless it is carried out carefully. These are not the only interesting discrete mathematical problems that arise in the course of combinatorial auctions, but they do illustrate the type.

Finally, an example illustrating a completely different regulatory problem with a discrete nature is that of aircraft noise. The regulations for aircraft noise (see for instance [4]) are based on Effective Perceived Noise Level (EPNL) and this is defined in terms of figures for the sound pressure levels SPL(i) in third-octave bands indexed by i running from 1 to 24. (These are in fact measured at 1 second intervals during take-off and landing but it is just the details within these figures for one second that we are concerned with here.) The figures are processed to produce a perceived noise level PNL, and a tone correction C, which are added to produce a tone-corrected perceived noise level PNLT. The PNL is produced by a detailed calculation from the SPL(i) using tables based on experiments that assessed how annoying people find the noise from different bands, and it is a monotonically increasing function of all the SPL(i). The tone correction C is intended to account for the fact that if the noise has a strong tonal component then it is more annoying. To find C, one first calculates a set of values SPL″(i) called the "Background sound pressure levels" that smooth out the SPL(i) after removing any pronounced spectral irregularities. Then the excesses F(i) = SPL(i) − SPL″(i) are calculated, converted to some tone corrections C(i) and the maximum of those is the final tone correction C. However, this has an unfortunate consequence. Suppose that the maximum of C(i) occurs for i_m. Then C = C(i_m) can be reduced by reducing SPL(i_m) − SPL″(i_m), and the background SPL″(i_m) is a smoothed average of the SPL(i) values for i near to i_m. So C(i_m) can be reduced by increasing SPL($i_m - 1$) and SPL($i_m + 1$), the sound pressure levels in the bands either side of the maximum. It turns out that this can decrease the penalty term C by *more* than the corresponding increase in the PNL term. In other words an aircraft could reduce its EPNL noise rating by judiciously increasing its noise output in the bands either side of its peak band. The problem has been introduced by the way the noise output is discretized into the bands and then processed in this way. It seems unfortunate, to say the least, that major international efforts on aircraft noise reduction such as the Silent Aircraft Initiative may end up being assessed by an internationally agreed EPNL rating that has this blemish in it. Let us hope that mathematicians can play a part in any future redrafting of these regulations, and others, to ensure that measures are monotonic where they should be.

Naturally I must apologise for this being such a small selection from the great host of interesting discrete problems that arise from industrial mathematics. I hope that others will present other examples in future issues of *Mathematics Today*.❑

REFERENCES

1 Frequency reassignment in cellular phone networks. Study Group report available from http://www.maths-in-industry.org/69/
2 Painting the office. Bruce Reed and David Allwright. To appear in Mathematics-in-Industry Case Studies.
3 Random Geometric Graphs. Mathew Penrose. OUP, 2003.
4 Federal Aviation Regulations. Part 36—Noise Standards: Aircraft Type and Airworthiness Certification.