Chapter 10. Section 2

Set Systems and Bipartite Graphs

We start with a slight modification to our example from the previous section:

Example $E = \{a, b, c, d, e\}$, $\mathcal{A} = \{\{a, c\}, \{a, d\}, \{a, b\}, \{a\}, \{d, e\}\}$.

We saw that previously $W = \{3, 4\}$ failed Hall's Theorem because sets 3 and 4 of $\mathcal{A}$ were both $\{a\}$ and did not collectively have at least $|W|$ elements. However, with $\mathcal{A}$ as above, $|A_3 \cup A_4| = |\{a, b\} \cup \{a\}| = 2$: the necessary and sufficient condition of Hall's Theorem is now satisfied. But *only for this choice of* W: we have to try *every* subset W of E to confirm that $\mathcal{A}$ has a transversal. In this case there are $2^{|E|} = 32$ subsets to try which is not too difficult; but for large E the task suffers from **combinatorial explosion**: 2^5 is easy; 2^{50} requires more than 35 years, testing one subset per microsecond; 2^{500} is impossible with current technology.

Nevertheless a good characterisation such as Hall's Theorem (or Kuratowski's theorem) usually suggests that the problem of checking existence is accessible to a fast algorithm. We shall see that this is the case for complete transversals by exploiting ideas from graph theory.

Take a set system $\mathcal{A}$ on ground set E. Turn it into a bipartite graph as follows: the vertex bipartition X, Y is defined by making X the ground set E and by making Y the sets of $\mathcal{A}$. Now an edge joins $e \in X$ to $A_i \in Y$ if and only if $e \in A_i$. This is called the **incidence relation** on E and $\mathcal{A}$ and the resulting bipartite graph is the **incidence graph** of the relation.

Example For the above set system the bipartite graph is shown below at (a) (the $\mathcal{A}$-sets are written without the brackets for simplicity).

(a) (b)

Now transversals correspond to disjoint sets of edges in the bipartite graph, i.e. subsets of the edges no two sharing an endpoint: each such edge pairs a different ground set representative with an $\mathcal{A}$-set to which it belongs. This is illustrated above at (b): a represents $\{a\}$, b represents $\{a, b\}$, c represents $\{a, c\}$, d represents $\{a, d\}$, and e represents $\{d, e\}$.

A subset of disjoint edges is called a **matching**. A spanning matching is called **perfect**, that is when all vertices of the graph are included in some edge of the matching. So if the ground set has the same size as $\mathcal{A}$ then complete transversals correspond to perfect matchings, and this is what happens at (b) in the above illustration.

This correspondence between transversals and matchings is not a 1–1 correspondence: a transversal is a subset of the ground set; a matching shows *one way* for this subset to represent the sets of $\mathcal{A}$. There may be others, and this leads to a combinatorial question: *how many* perfect matchings are there? This counting problem appears to be very hard. There are fast algorithms for finding perfect matchings in bipartite graphs, and therefore for checking whether the necessary and sufficient condition of Hall's Theorem is satisfied. But these algorithms stop when they have found a maximum-size matching — they are not designed to list all the possible matchings.

We can be a bit more explicit about why counting perfect matchings is hard. We represent our bipartite graph using a matrix whose rows represent the sets of $\mathcal{A}$ and whose columns represent the ground set E. Now record a '1' in row i, column j if A_i contains element j of E; all other entries are zero. Thus each '1' in the matrix corresponds to an edge in the incidence graph. This matrix is called the **incidence matrix** of the set system.

The incidence matrix for our example set system is given on the right. Zero entries are omitted for clarity; and the entries corresponding to the perfect matching given at (b) above have been circled. Now two edges are disjoint in the incidence graph of $\mathcal{A}$ if and only if the corresponding '1's in the incidence matrix occupy different rows and columns. And a subset of edges is spanning if and only if every row and column contains one of the corresponding '1's.

	a	b	c	d	e
ac	1		①		
ad	1			①	
ab	1	①			
a	①				
de					①

So we have

perfect matchings in the incidence graph = rook placements in the incidence matrix = permutations.

The above perfect matching corresponds to the permutation $1 \to 3, 2 \to 4, 3 \to 2, 4 \to 1, 5 \to 5$ which, in disjoint cycle notation, is given more concisely as (1 3 2 4).

Now recall the basic definition of the **determinant** function: for an $n \times n$ matrix X, $\det X$ is the sum, over all permutations π of $[n]$, of $\mathrm{sgn}(\pi) \times$the product of the entries of X corresponding to π, where $\mathrm{sgn}(\pi)$ is the 'sign' of π: $+1$ if π is an 'even' permutation and -1 if π is 'odd' (don't worry if you have forgotten the details — we are about to dump the sgn anyway!)

More formally,

$$\det X = \sum_{\pi \in \mathcal{S}_n} \mathrm{sgn}(\pi) \prod_{i \in [n]} x_{i,\pi(i)}.$$

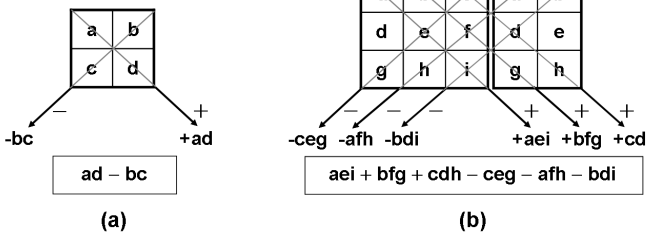

where $\mathcal{S}_n$ is the group of all permutations of $[n]$. We repeat on the right the illustration for $n = 2$ and $n = 3$ which we gave in Chapter 7, Section 3 which lists the permutations and their signs explicitly. Although these signs complicate the definition of the determinant function they are precisely the reason why Gaussian elimination works and why calculating the value of $\det X$ can be done rapidly (in the sense that it takes about n^3 multiplications instead of about $n!$: compare $10^3 = 1000$ with $10! = 3628800$ for a 10×10 matrix).

Now to count perfect matchings in an incidence graph G we want to find every permutation in the incidence matrix M. This is the same as taking *all* permutations and multiplying together the corresponding entries in M, since this will record a matching exactly when every permutation entry is 1. And this is precisely the determinant function *but ignoring the sgn*(π) *because we want every perfect matching to contribute* $+1$ *to our count*.

The function which is the same as the determinant but without the $\mathrm{sgn}(\pi)$ is called the **permanent**:

$$\mathrm{per}X = \sum_{\pi \in \mathcal{S}_n} \prod_{i \in [n]} x_{i,\pi(i)}.$$

Example For the matrix $X = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$ we have $\det X = -0 - 0 - 1 + 1 + 1 + 0 = 1$ and $\mathrm{per}X = 0 + 1 + 1 + 1 + 1 + 0 = 3$, using the 'Rule of Sarrus' scheme above right.

Unlike the determinant function, there is no known way to sum the n! terms of the permanent other than explicitly: Gaussian elimination does not work. Indeed calculating the permanent belongs to the class of so-called **#P-complete** counting problems: solving any one would give a solution to every other; since this includes a large number of different problems which have been studied using many different techniques for many years, it seems unlikely that any can be solved. Counting perfect matchings is one example, counting colourings is another: the only way to evaluate $C(G; \lambda)$ appears to involve something like constructing an exponentially large deletion/contraction tree.