Acyclic orientations

Peter J. Cameron

University of St Andrews
Queen Mary University of London

Probability and graphs
Eindhoven, 6 January 2014

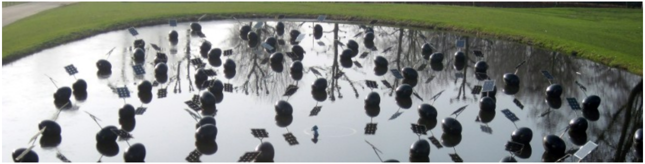

Acknowledgements

This is joint work with Celia Glass and Robert Schumacher (City University, London), and some of it will be in Robert's PhD thesis (currently being written).

Acknowledgements

This is joint work with Celia Glass and Robert Schumacher (City University, London), and some of it will be in Robert's PhD thesis (currently being written).
I am currently a part-time professor at the University of St Andrews and emeritus at Queen Mary.

Counting and probability

"I count a lot of things that there's no need to count," Cameron said. "Just because that's the way I am. But I count all the things that need to be counted."

Richard Brautigan, The Hawkline Monster: A Gothic Western

Counting and probability

> *"I count a lot of things that there's no need to count," Cameron said. "Just because that's the way I am. But I count all the things that need to be counted."*
>
> *Richard Brautigan, The Hawkline Monster: A Gothic Western*

Most of this talk is about counting, not probability. But, as we know, there is a close connection between being able to count a set and being able to choose one of its elements at random; so this is a preliminary to studying typical properties of elements of the set.

Acyclic orientations

In this talk, G will be a graph with n vertices and m edges; the vertex and edge sets are V and E respectively. All graphs are simple.

Acyclic orientations

In this talk, G will be a graph with n vertices and m edges; the vertex and edge sets are V and E respectively. All graphs are simple.

An acyclic orientation of a graph is an orientation of the edges of the graph such that there are no directed cycles. I will often abbreviate "acyclic orientation" to a.o.

Acyclic orientations

In this talk, G will be a graph with n vertices and m edges; the vertex and edge sets are V and E respectively. All graphs are simple.

An acyclic orientation of a graph is an orientation of the edges of the graph such that there are no directed cycles. I will often abbreviate "acyclic orientation" to a.o.

Any acyclic orientation can be obtained by ordering the vertices and then orienting all edges from the smaller to the greater vertex.

Acyclic orientations and colourings

Any acyclic orientation gives rise to a colouring. Colour the sources with the first colour; then delete them and repeat.

Acyclic orientations and colourings

Any acyclic orientation gives rise to a colouring. Colour the sources with the first colour; then delete them and repeat. The colour given to a vertex is the length of the longest directed path ending at the vertex. So the total number of colours used is equal to the longest directed path.

Acyclic orientations and colourings

Any acyclic orientation gives rise to a colouring. Colour the sources with the first colour; then delete them and repeat.
The colour given to a vertex is the length of the longest directed path ending at the vertex. So the total number of colours used is equal to the longest directed path.
Depending on the a.o. chosen, this can take any value from the chromatic number of the graph up to the length of the longest path.

Acyclic orientations and colourings

Any acyclic orientation gives rise to a colouring. Colour the sources with the first colour; then delete them and repeat.
The colour given to a vertex is the length of the longest directed path ending at the vertex. So the total number of colours used is equal to the longest directed path.
Depending on the a.o. chosen, this can take any value from the chromatic number of the graph up to the length of the longest path.
It would be interesting to know something about the distribution of the number of colours over all a.o.s of a given graph.

Moving between acyclic orientations

Define the distance between two a.o.s of G to be the number of edges oriented differently in the two a.o.s.

Moving between acyclic orientations

Define the distance between two a.o.s of G to be the number of edges oriented differently in the two a.o.s.

Proposition

- *Any a.o. of G has an edge whose orientation can be flipped while preserving the acyclic property.*

Moving between acyclic orientations

Define the distance between two a.o.s of G to be the number of edges oriented differently in the two a.o.s.

Proposition

- *Any a.o. of G has an edge whose orientation can be flipped while preserving the acyclic property.*
- *Any a.o. can be transformed into any other by flipping (in some order) just those edges which are oriented differently in the two orientations; so the distance is equal to the minimum number of edge flips required to transform one into the other.*

A Markov chain

The edge flips define a Markov chain on the set of a.o.s of G: choose a random edge and flip it if possible. (This chain is irreducible and aperiodic. But its limiting distribution is not uniform.)

A Markov chain

The edge flips define a Markov chain on the set of a.o.s of G: choose a random edge and flip it if possible. (This chain is irreducible and aperiodic. But its limiting distribution is not uniform.)
It would be nice to modify this Markov chain so as to bias it towards a.o.s which give rise to colourings with few colours, but we have not succeeded in doing this.

Counting acyclic orientations

The number of a.o.s of a given graph is an interesting graph parameter.

Counting acyclic orientations

The number of a.o.s of a given graph is an interesting graph parameter.

Theorem (Stanley)

The number of acyclic orientations of G is equal to $(-1)^n P_G(-1)$, where P_G is the chromatic polynomial of G.

Counting acyclic orientations

The number of a.o.s of a given graph is an interesting graph parameter.

Theorem (Stanley)

The number of acyclic orientations of G is equal to $(-1)^n P_G(-1)$, *where* P_G *is the chromatic polynomial of G.*

So it is an evaluation of the Tutte polynomial.

Counting acyclic orientations

The number of a.o.s of a given graph is an interesting graph parameter.

Theorem (Stanley)

The number of acyclic orientations of G is equal to $(-1)^n P_G(-1)$, where P_G is the chromatic polynomial of G.

So it is an evaluation of the Tutte polynomial.
It is hard to compute exactly, and it is currently unknown whether there is an efficient approximation procedure (fpras) for it.

Distribution: average

Given n and m, what is the distribution of the number of a.o.s of graphs with n vertices and m edges?

Distribution: average

Given n and m, what is the distribution of the number of a.o.s of graphs with n vertices and m edges?
Let $a(n,m)$ be the number of acyclic digraphs with n labelled vertices and m arcs, and $A_n(x)$ the generating polynomial

$$\sum_{m=0}^{n(n-1)/2} a(n,m)x^m.$$

Distribution: average

Given n and m, what is the distribution of the number of a.o.s of graphs with n vertices and m edges?
Let $a(n,m)$ be the number of acyclic digraphs with n labelled vertices and m arcs, and $A_n(x)$ the generating polynomial

$$\sum_{m=0}^{n(n-1)/2} a(n,m)x^m.$$

Theorem (Bender, Richmond, Robinson, Wormald)

$$A_n(x) = \sum_{i=1}^{n}(-1)^{i+1}\binom{n}{i}(1+x)^{i(n-i)}A_{n-i}(x).$$

Distribution: average

Given n and m, what is the distribution of the number of a.o.s of graphs with n vertices and m edges?
Let $a(n, m)$ be the number of acyclic digraphs with n labelled vertices and m arcs, and $A_n(x)$ the generating polynomial
$\sum_{m=0}^{n(n-1)/2} a(n, m)x^m$.

Theorem (Bender, Richmond, Robinson, Wormald)

$$A_n(x) = \sum_{i=1}^{n} (-1)^{i+1} \binom{n}{i} (1+x)^{i(n-i)} A_{n-i}(x).$$

Now $a(n, m) \Big/ \binom{n(n-1)/2}{m}$ is the average number of a.o.s. and is easily computed.

Distribution: minimum

Theorem (Linial)

The graph with n vertices and m edges having the smallest number of acyclic orientations is constructed as follows: Let $r(r-1)/2$ be the largest triangular number not exceeding m; take a complete graph on r vertices, and add a new vertex joined to $q = m - r(r-1)/2$ vertices in the clique. The number of a.o.s is $q!(r+1)$.

Distribution: minimum

Theorem (Linial)

The graph with n vertices and m edges having the smallest number of acyclic orientations is constructed as follows: Let $r(r-1)/2$ be the largest triangular number not exceeding m; take a complete graph on r vertices, and add a new vertex joined to $q = m - r(r-1)/2$ vertices in the clique. The number of a.o.s is $q!(r+1)$.

Note that the minimum is linear between consecutive triangular numbers.

Distribution: minimum

Theorem (Linial)

The graph with n vertices and m edges having the smallest number of acyclic orientations is constructed as follows: Let $r(r-1)/2$ be the largest triangular number not exceeding m; take a complete graph on r vertices, and add a new vertex joined to $q = m - r(r-1)/2$ vertices in the clique. The number of a.o.s is $q!(r+1)$.

Note that the minimum is linear between consecutive triangular numbers.
We have a new proof of this theorem which shows that the same graphs are also extremal for several further graph parameters.

Distribution: maximum

Finding the maximum number of a.o.s of a graph with n vertices and m edges is considerably more challenging.

Distribution: maximum

Finding the maximum number of a.o.s of a graph with n vertices and m edges is considerably more challenging.
A Túran graph is a complete multipartite graph, the sizes of whose parts differ by at most 1. Turán graphs famously arise as the graphs with the largest number of edges subject to forbidding complete graphs of given size.

Distribution: maximum

Finding the maximum number of a.o.s of a graph with n vertices and m edges is considerably more challenging.
A Túran graph is a complete multipartite graph, the sizes of whose parts differ by at most 1. Turán graphs famously arise as the graphs with the largest number of edges subject to forbidding complete graphs of given size.

Conjecture

If n and m are such that a Turán graph with n vertices and m edges exists, then the Turán graph has the maximum number of acyclic orientations.

Distribution: maximum

Finding the maximum number of a.o.s of a graph with n vertices and m edges is considerably more challenging.
A Túran graph is a complete multipartite graph, the sizes of whose parts differ by at most 1. Turán graphs famously arise as the graphs with the largest number of edges subject to forbidding complete graphs of given size.

Conjecture

If n and m are such that a Turán graph with n vertices and m edges exists, then the Turán graph has the maximum number of acyclic orientations.

This is proved only in the case where m is only slightly less than $n(n-1)/2$: if we only omit a few edges from the complete graph, they should be disjoint.

Some numbers

vertices	2	4	6	8	10	12
edges	1	4	9	16	25	36
Min # a.o.s	2	12	96	1440	25200	362880
Ave # a.o.s	2	12	167	3851	156636	9312017
Max # a.o.s?	2	14	230	6902	329462	22934774

Table: Numbers of acyclic orientations of graphs with $m = n^2/4$

Some numbers

vertices	2	4	6	8	10	12
edges	1	4	9	16	25	36
Min # a.o.s	2	12	96	1440	25200	362880
Ave # a.o.s	2	12	167	3851	156636	9312017
Max # a.o.s?	2	14	230	6902	329462	22934774

Table: Numbers of acyclic orientations of graphs with $m = n^2/4$

The values given for the maximum are those suggested by the conjecture. Robert Schumacher computed some of these, and looked them up in the On-Line Encyclopedia of Integer Sequences. The references led to the following discussion.

Complete bipartite graphs

There is a formula for the number of acyclic orientations of a complete bipartite graph.

Complete bipartite graphs

There is a formula for the number of acyclic orientations of a complete bipartite graph.

Theorem

The number of acyclic orientations of K_{n_1,n_2} is

$$\sum_{k=1}^{\min(n_1,n_2)+1} ((k-1)!)^2 S(n_1+1,k) S(n_2+1,k).$$

Complete bipartite graphs

There is a formula for the number of acyclic orientations of a complete bipartite graph.

Theorem

The number of acyclic orientations of K_{n_1,n_2} is

$$\sum_{k=1}^{\min(n_1,n_2)+1} ((k-1)!)^2 S(n_1+1,k) S(n_2+1,k).$$

Here $S(a,b)$ is the Stirling number of the second kind, the number of partitions of a set of size a into b parts.

Proof

Let the bipartition of the vertex set of the graph have parts A and B, which we think of as coloured amber and blue.

Proof

Let the bipartition of the vertex set of the graph have parts A and B, which we think of as coloured amber and blue.

Any acyclic orientation is defined by a sequence $(A_1, B_1, \ldots, A_k, B_k)$, where the A_i form a partition of A and the B_i of B, and all parts are non-empty except possibly A_1 and B_k.

Proof

Let the bipartition of the vertex set of the graph have parts A and B, which we think of as coloured amber and blue.
Any acyclic orientation is defined by a sequence $(A_1, B_1, \ldots, A_k, B_k)$, where the A_i form a partition of A and the B_i of B, and all parts are non-empty except possibly A_1 and B_k.
To get round having to consider four cases, we add a new vertex a_0 to A_1 and a new vertex b_0 to B_k. So we get an acyclic orientation by first choosing partitions (which can be done in $S(n_1 + 1, k)S(n_2 + 1, k)$ ways), ordering them with A_1 first and B_k last (in $((k-1)!)^2$ ways), and then removing the added vertices a_0 and b_0.

Proof

Let the bipartition of the vertex set of the graph have parts A and B, which we think of as coloured amber and blue.
Any acyclic orientation is defined by a sequence $(A_1, B_1, \ldots, A_k, B_k)$, where the A_i form a partition of A and the B_i of B, and all parts are non-empty except possibly A_1 and B_k.
To get round having to consider four cases, we add a new vertex a_0 to A_1 and a new vertex b_0 to B_k. So we get an acyclic orientation by first choosing partitions (which can be done in $S(n_1+1,k)S(n_2+1,k)$ ways), ordering them with A_1 first and B_k last (in $((k-1)!)^2$ ways), and then removing the added vertices a_0 and b_0.
Finally, sum over k. The result is as in the theorem.

Poly-Bernoulli numbers

This is only a very brief introduction to these numbers, which were introucd by Masanobu Kaneko in 1997.

Poly-Bernoulli numbers

This is only a very brief introduction to these numbers, which were introucd by Masanobu Kaneko in 1997.
Kaneko gave the following definitions. Let

$$\mathrm{Li}_k(z) = \sum_{m=1}^{\infty} \frac{z^m}{m^k},$$

and let

$$\frac{\mathrm{Li}_k(1 - \mathrm{e}^{-x})}{1 - \mathrm{e}^{-x}} = \sum_{n=0}^{\infty} B_n^{(k)} \frac{x^n}{n!}.$$

The numbers $B_n^{(k)}$ are the poly-Bernoulli numbers of order k.

Poly-Bernoulli numbers

This is only a very brief introduction to these numbers, which were introucd by Masanobu Kaneko in 1997.
Kaneko gave the following definitions. Let

$$\mathrm{Li}_k(z) = \sum_{m=1}^{\infty} \frac{z^m}{m^k},$$

and let

$$\frac{\mathrm{Li}_k(1-\mathrm{e}^{-x})}{1-\mathrm{e}^{-x}} = \sum_{n=0}^{\infty} B_n^{(k)} \frac{x^n}{n!}.$$

The numbers $B_n^{(k)}$ are the poly-Bernoulli numbers of order k.
He gave a couple of nice formulae for the poly-Bernoulli numbers of negative order, of which the one on the next slide is relevant here:

Kaneko's Theorem

Theorem (Kaneko)

$$B_n^{(-k)} = \sum_{j=0}^{\min(n,k)} (j!)^2 S(n+1, j+1) S(k+1, j+1).$$

Kaneko's Theorem

Theorem (Kaneko)

$$B_n^{(-k)} = \sum_{j=0}^{\min(n,k)} (j!)^2 S(n+1, j+1) S(k+1, j+1).$$

This formula has the (entirely non-obvious) corollary that these numbers have a symmetry property: $B_n^{(-k)} = B_k^{(-n)}$ for all non-negative integers n and k.

Kaneko's Theorem

Theorem (Kaneko)

$$B_n^{(-k)} = \sum_{j=0}^{\min(n,k)} (j!)^2 S(n+1, j+1) S(k+1, j+1).$$

This formula has the (entirely non-obvious) corollary that these numbers have a symmetry property: $B_n^{(-k)} = B_k^{(-n)}$ for all non-negative integers n and k.
It also shows that the number of acyclic orientations of K_{n_1,n_2} is $B_{n_1}^{(-n_2)}$.

Lonesum matrices

This is closely related to recent work of Chet Brewbaker (2008).

Lonesum matrices

This is closely related to recent work of Chet Brewbaker (2008). A zero-one matrix is a lonesum matrix if it is uniquely determined by its row and column sum.

Lonesum matrices

This is closely related to recent work of Chet Brewbaker (2008). A zero-one matrix is a lonesum matrix if it is uniquely determined by its row and column sum.

Clearly a lonesum matrix cannot contain either $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ as a submatrix (in not necessarily consecutive rows or columns). (If one such submatrix occurred it could be flipped into the other without changing the row and column sums.)

Lonesum matrices

This is closely related to recent work of Chet Brewbaker (2008). A zero-one matrix is a lonesum matrix if it is uniquely determined by its row and column sum.

Clearly a lonesum matrix cannot contain either $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ as a submatrix (in not necessarily consecutive rows or columns). (If one such submatrix occurred it could be flipped into the other without changing the row and column sums.) Ryser showed that, conversely, a matrix containing neither of these is a lonesum matrix.

Lonesum matrices

This is closely related to recent work of Chet Brewbaker (2008).

A zero-one matrix is a lonesum matrix if it is uniquely determined by its row and column sum.

Clearly a lonesum matrix cannot contain either $\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$ or $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ as a submatrix (in not necessarily consecutive rows or columns). (If one such submatrix occurred it could be flipped into the other without changing the row and column sums.)

Ryser showed that, conversely, a matrix containing neither of these is a lonesum matrix.

Brewbaker showed that the number of $n_1 \times n_2$ lonesum matrices is given by our earlier formula.

The connection, 1

The number of $n_1 \times n_2$ lonesum matrices is equal to the number of a.o.s of K_{n_1,n_2}.

The connection, 1

The number of $n_1 \times n_2$ lonesum matrices is equal to the number of a.o.s of K_{n_1,n_2}.
Number the vertices in the bipartite blocks from 1 to n_1 (in A) and from 1 to n_2 (in B). Now given an orientation of the graph, we can describe it by a matrix whose (i, j) entry is 1 if the edge from vertex i of A to vertex j of B goes in the direction from A to B, and 0 otherwise. The two forbidden submatrices for lonesum matrices correspond to directed 4-cycles; so any a.o. gives us a lonesum matrix.

The connection, 2

Conversely, if an orientation of a complete bipartite graph contains no directed 4-cycles, then it contains no directed cycles at all.

The connection, 2

Conversely, if an orientation of a complete bipartite graph contains no directed 4-cycles, then it contains no directed cycles at all.
For suppose that there are no directed 4-cycles, but there is a directed cycle $(a_1, b_1, a_2, b_2, \ldots, a_k, b_k, a_1)$. Then the edge between a_2 and b_2 must be directed from a_1 to b_2, since otherwise there would be a 4-cycle $(a_1, b_1, a_2, b_2, a_1)$. But then we have a shorter directed cycle $(a_1, b_2, a_3, \ldots, b_k, a_1)$.
Continuing this shortening process, we would eventually arrive at a directed 4-cycle, a contradiction.

The connection, 2

Conversely, if an orientation of a complete bipartite graph contains no directed 4-cycles, then it contains no directed cycles at all.

For suppose that there are no directed 4-cycles, but there is a directed cycle $(a_1, b_1, a_2, b_2, \ldots, a_k, b_k, a_1)$. Then the edge between a_2 and b_2 must be directed from a_1 to b_2, since otherwise there would be a 4-cycle $(a_1, b_1, a_2, b_2, a_1)$. But then we have a shorter directed cycle $(a_1, b_2, a_3, \ldots, b_k, a_1)$. Continuing this shortening process, we would eventually arrive at a directed 4-cycle, a contradiction.

(This simply says that the cycle space of the complete bipartite graph is generated by 4-cycles.)

Some questions

1. Is there are fully polynomial randomized approximation scheme (fpras) for the number of acyclic orientations of a graph?

Some questions

1. Is there are fully polynomial randomized approximation scheme (fpras) for the number of acyclic orientations of a graph?
2. What can be said about the length of the longest directed path in a random acyclic orientation (equivalently, the number of colours in the corresponding colouring)?

Some questions

1. Is there are fully polynomial randomized approximation scheme (fpras) for the number of acyclic orientations of a graph?
2. What can be said about the length of the longest directed path in a random acyclic orientation (equivalently, the number of colours in the corresponding colouring)?
3. Describe the graphs (with given numbers of vertices and edges) which have the maximum number of acyclic orientations.

Some questions

1. Is there are fully polynomial randomized approximation scheme (fpras) for the number of acyclic orientations of a graph?
2. What can be said about the length of the longest directed path in a random acyclic orientation (equivalently, the number of colours in the corresponding colouring)?
3. Describe the graphs (with given numbers of vertices and edges) which have the maximum number of acyclic orientations.
4. What can be said about the distribution of the number of acyclic orientations of graphs with n vertices and m edges? Can its variance be calculated?

Some questions

1. Is there are fully polynomial randomized approximation scheme (fpras) for the number of acyclic orientations of a graph?
2. What can be said about the length of the longest directed path in a random acyclic orientation (equivalently, the number of colours in the corresponding colouring)?
3. Describe the graphs (with given numbers of vertices and edges) which have the maximum number of acyclic orientations.
4. What can be said about the distribution of the number of acyclic orientations of graphs with n vertices and m edges? Can its variance be calculated?
5. Is there a formula for the number of acyclic orientations of a complete multipartite graph (in terms of the sizes of the parts)?

6. Are there analogues for acyclic orientations of correlation and anticorrelation results for spanning trees or forests? Any given edge has probability $\frac{1}{2}$ of having each direction in a random acyclic orientation, but are there any general results on correlation, perhaps in terms of the distance between the edges? Or any results that hold for particular graphs?

6. Are there analogues for acyclic orientations of correlation and anticorrelation results for spanning trees or forests? Any given edge has probability $\frac{1}{2}$ of having each direction in a random acyclic orientation, but are there any general results on correlation, perhaps in terms of the distance between the edges? Or any results that hold for particular graphs?

In the case of complete bipartite graphs, we are in a good position because all non-incident edge pairs are alike, and we have a formula for the number of a.o.s.

6. Are there analogues for acyclic orientations of correlation and anticorrelation results for spanning trees or forests? Any given edge has probability $\frac{1}{2}$ of having each direction in a random acyclic orientation, but are there any general results on correlation, perhaps in terms of the distance between the edges? Or any results that hold for particular graphs?

In the case of complete bipartite graphs, we are in a good position because all non-incident edge pairs are alike, and we have a formula for the number of a.o.s.

7. How is the number of binary matrices with given row and column sum distributed, over all feasible pairs of row/column sum tuples?