

PROBABILISTIC COMBINATORICS AND THE RECENT WORK OF PETER KEEVASH

W. T. GOWERS

1. THE PROBABILISTIC METHOD

A *graph* is a collection of points, or *vertices*, some of which are joined together by *edges*. Graphs can be used to model a wide range of phenomena: whenever you have a set of objects such that any two may or may not be related in a certain way, then you have a graph. For example, the vertices could represent towns and the edges could represent roads from one town to another, or the vertices could represent people, with an edge joining two people if they know each other, or the vertices could represent countries, with an edge joining two countries if they border each other, or the vertices could represent websites, with edges representing links from one site to another (in this last case the edges have directions—a link is *from* one website *to* another—so the resulting mathematical structure is called a *directed graph*).

Mathematically, a graph is a very simple object. It can be defined formally as simply a set V of vertices and a set E of unordered pairs of vertices from V . For example, if we take $V = \{1, 2, 3, 4\}$ and $E = \{12, 23, 34, 14\}$ (using ab as shorthand for $\{a, b\}$), then we obtain a graph known as the *4-cycle*.

Despite the simplicity and apparent lack of structure of a general graph, there turn out to be all sorts of interesting questions one can ask about them. Here is a classic example. To state it, we shall need a small amount of standard graph-theory terminology. A *triangle* in a graph is, as one would expect, a triple of vertices x, y, z such that all of xy, yz , and xz are edges. A graph is called *triangle free* if it contains no triangles. A *proper colouring* of a graph is a way of assigning colours to its vertices such that no two vertices of the same colour are joined by an edge. The *chromatic number* of a graph is the smallest number of colours you need in a proper colouring. (One of the most famous results of graph theory, the four-colour theorem, asserts that any graph that you can draw in the plane without any of its edges crossing each other has chromatic number at most 4. This can be interpreted as saying that if you want to colour the countries in a map in such a way that no two adjacent countries have the same colour, then four colours will suffice.)

And now the question.

Question 1.1. Does there exist a triangle-free graph with chromatic number 2016?

How might one answer Question 1.1? The obvious thing would be to sit down and try to construct one. But this seems to be rather hard. How can we ensure

Received by the editors August 3, 2016.
2010 *Mathematics Subject Classification*. Primary 05-02.

that plenty of colours are needed? The only method that is immediately apparent is to have a large set of vertices that are all joined to each other by edges, since then all those vertices have to have different colours. But of course, if we do that, then we massively violate the other main constraint, that the graph should contain no triangles.

That suggests that perhaps one can always properly colour a triangle-free graph with a small number of colours. So how might we do that? Perhaps we could list the vertices in some order and colour them in that order, using new colours only when forced to do so. We could colour the first vertex red, then the second blue (if it is joined to the first). The third would be joined to at most one of the first two, so could be coloured either red or blue. And so on.

If one keeps going like this, one soon discovers that the number of colours one needs can grow without limit. Unfortunately, this does not solve the problem: it merely shows that choosing the colours in a greedy way does not work.

Although it is not the only way of solving the problem, there is an extraordinarily simple and powerful idea that does the job. It is to take an appropriate *random* graph. (This hint does not entirely spoil the problem, since it is still a very nice challenge to come up with an explicit construction of a graph that works.)

Here is how the argument works. Let p be a probability that we will choose later. Now let G be a random graph where we join a vertex x to a vertex y with probability p , making all choices independently.

If there are n vertices, then the expected number of triangles in the graph is at most $p^3 \binom{n}{3}$. This means that it is possible to make the graph triangle free by deleting at most $p^3 \binom{n}{3}$ edges.

We now make two simple observations. Recall that an *independent set* in a graph is a set of vertices no two of which are joined by an edge. The very simple observation is that in a proper colouring, the vertices of any one colour must form an independent set, so if a graph has n vertices and chromatic number k , then it must have an independent set of size at least n/k . Therefore, to prove that the chromatic number of a graph is large, it is sufficient to find a small upper bound on the *independence number* of the graph—that is, the size of its largest independent set.

The second observation is that if every set of m vertices contains more than $p^3 \binom{n}{3}$ edges, then we can remove an arbitrary set of $p^3 \binom{n}{3}$ edges and will be left with a graph that has independence number greater than m . So our task is reduced to proving that every reasonably large set of vertices includes a reasonably large number of edges.

At this point I will not give full details. I will just say that the probability $q(m, r)$ that some given set of m vertices spans fewer than r edges is given by the binomial distribution $B(\binom{m}{2}, p)$, for which we have standard estimates. So to ensure that, with probability greater than $1/2$, every set of m vertices spans at least r edges, one needs to choose the parameters in such a way that $\binom{n}{m} q(m, r) < 1/2$. This is not hard to do, and it leads to a proof that there exists a triangle-free graph with chromatic number proportional to $\sqrt{n}/\log n$. Finally, choosing n large enough, we obtain a triangle-free graph with chromatic number at least 2016. (If we want, we can remove edges from it until it has chromatic number exactly 2016, thereby answering Question 1.1 exactly as it was asked.)

This general approach—showing that graphs with certain properties exist by proving that appropriately distributed random graphs have the properties with non-zero probability—was pioneered by Erdős in a famous paper that gave an exponential lower bound for the Ramsey number $R(k, k)$ [3]. The proof just sketched is the special case for girth 3 of an argument, also by Erdős, that there can be graphs with arbitrarily high girth (the girth of a graph is the length of the shortest cycle it contains) and arbitrarily high chromatic number [4]. The probabilistic method is now a huge part of combinatorics, and it has had an impact on many other subjects as well. For an excellent introduction to it, see a classic text of Alon and Spencer [2].

2. WHEN IS THE PROBABILISTIC METHOD APPROPRIATE?

This is not an easy question to answer with complete precision, but there are some useful rough guidelines that one can give. Take the proof I have just sketched. After the two simple observations, our aim became to find a triangle-free graph with the property that for any set of m vertices there would be at least r edges between those vertices. This task has the following four properties.

- Graphs are very structureless objects, in the sense that a general graph has very few constraints that it needs to satisfy.
- It seems to be hard to define such a graph explicitly.
- We want to scatter the edges around very evenly.
- We want to do that very efficiently (so that we do not end up choosing so many edges that they bunch together to form triangles).

In such a situation, choosing edges at random makes very good sense: somehow, it allows us to do a lot of tasks in parallel, each with a very high chance of success.

Contrast that with a task such as solving the famous Burnside problem: does there exist a prime p and an infinite but finitely generated group such that every element has order p ? It would be ludicrous to suggest picking a random group as a possible solution. That is partly because for an object to be a group, it has to satisfy a set of axioms that impose significantly greater constraints than those of a graph, where one simply chooses an arbitrary set of edges. Indeed, it is far from obvious what a good model for a random group should be. (There are in fact some very interesting models introduced by Gromov [5, 6] where one chooses a set of generators and a random set of relations of a certain length, but this model of random groups is of no help for the Burnside problem.) It is also because for the Burnside problem there is an obvious construction: choose the free group $G(k, p)$ on k generators subject to the relations that every element has order p . This is a well-defined group, and if such groups can be infinite, then this one must be. So the problem boils down to showing that some $G(k, p)$ is infinite.

There are also some intermediate problems—that is, problems that seem to have too many constraints for probabilistic methods to be appropriate, but too few for there to be obvious explicit constructions. An example of such a problem is the following beautiful question. Recall that a *Hamilton cycle* in a graph is a cycle that visits every vertex. That is, it is a sequence v_1, v_2, \dots, v_n of vertices, where each vertex is included exactly once, such that $v_1v_2, v_2v_3, \dots, v_{n-1}v_n$ and v_nv_1 are all edges.

Question 2.1. Let $n = 2m + 1$ be an odd integer greater than or equal to 3, and define a graph G_n as follows. Its vertices are all subsets of $\{1, 2, \dots, n\}$ of size m or $m + 1$, and the edges are all pairs AB such that $|A| = m$, $|B| = m + 1$, and $A \subset B$. Does every G_n contain a Hamilton cycle?

The graph G_n is called the *middle layers graph* in the *discrete cube*.

If one starts trying to build a Hamilton cycle in G_n , one runs into the problem of having too much choice and no obvious way of making it. (A natural thing to try to do is find some sort of inductive construction, but a lot of people have tried very hard to do this, with no success—a natural pattern just doesn't seem to emerge after the first few small cases.)

So there are not enough constraints to force one's hand and, in that way, lead one to a solution. At first this sounds like just the kind of situation for which the probabilistic method was designed: why not start at a set and then keep randomly choosing neighbouring vertices that have not yet been visited? Of course, that cannot be the whole story, since there is no guarantee that one will not get stuck at some point, but perhaps one can combine this basic idea with a little bit of backtracking here and there and end up proving the existence of a Hamilton cycle.

Unfortunately, this does not seem to work either. After a while, the constraints start to bite, and one gets sufficiently stuck (at least potentially) that even being allowed a little local tinkering does not allow one to proceed.

So here we have a situation that is difficult because it is somehow “intermediate”: it falls between being highly structured, where one has few options, so constructions, if they exist, are in a certain sense easier to find, and being highly unstructured, where one has so many options that making random choices does not violate the constraints.

This well-known problem appears to have been recently solved by Torsten Mütze [9], but I do not know whether his argument has yet been thoroughly checked.

3. THE EXISTENCE OF DESIGNS

A *Steiner triple system* is a finite set X and a collection T of triples of elements of X with the property that every pair x, y of elements of X is contained in exactly one triple xyz from T . One can think of it as a set of triangles that partition the edges of the complete graph with vertex set X .

Do Steiner triple systems exist? Well, an obvious constraint is that the number of edges in the complete graph with vertex set X should be a multiple of 3. That is, if $|X| = n$, we need that $3 \mid \binom{n}{2}$. This holds if and only if $n \equiv 0$ or $1 \pmod{3}$. A slightly less obvious constraint (until it is pointed out) is that n should be odd. That is because the triangles xyz that contain a given vertex x will partition the $n - 1$ edges that meet x into sets of size 2. Thus, n must, for very simple reasons, be congruent to 1 or 3 mod 6.

When $n = 1$, the empty set is a trivial Steiner system. When $n = 3$, a single triangle does the job. The next case, when $n = 7$, is more interesting: there is a famous example known as the *Fano plane*. If we number its vertices from 1 to 7, then the triples can be given as follows: 123, 246, 257, 167, 347, 356, 145. A less mysterious definition is to take as vertices the set of all triples of 0's and 1's apart from the triple 000, and then to take as our triple system (which must consist of triples of these triples!) the set of all xyz such that $x + y + z = 000$, where the

addition is mod 2 in each coordinate. For example, one of the elements of this triple system is the triple xyz with $x = 100, y = 101,$ and $x = 001$.

Why does this work? Well it is clear that an edge xy can be contained in at most one triple from the system, since the only possibility is the unique triple xyz , such that $x + y + z = 000$, which forces z to equal $-(x + y)$, which is the same as $x + y$ since addition is mod 2. The only thing that could conceivably go wrong is if z turned out to equal x or y , but if $x + y$ is equal to x or y , then one of x and y is 000, which is not allowed.

Note that this construction gives an infinite family of Steiner triple systems, since the same proof works if we take non-zero 01-sequences of any fixed length k . So there are Steiner triple systems for every n of the form $2^k - 1$. It turns out that there are constructions similar in spirit to this one (making use of objects called quasigroups) that prove that Steiner triple systems exist for every n that satisfies the obviously necessary divisibility conditions—that is, for every n congruent to 1 or 3 mod 6.

A *design* is a generalization of a Steiner triple system, where instead of aiming to include every set of size 2 in exactly one set (from a carefully chosen collection) of size 3, one wishes to include every set of size r in exactly one set of size s . People also consider a further generalization where one wishes for every set of size r to be contained in exactly λ sets of size s . When $\lambda = 1$, the case I shall concentrate on here, designs are called *Steiner systems*.

No sooner is this definition presented, than several obvious questions immediately arise. Here are five, in decreasing order of optimism. Let us say that a *Steiner (n, r, s) -system* is a collection Σ of subsets size r of a set X of size n such that every subset of X of size s is contained in exactly one set from Σ .

Question 3.1. If obvious necessary divisibility conditions are satisfied, does it follow that a Steiner (n, r, s) -system exists?

Question 3.2. If obvious necessary divisibility conditions are satisfied and n is sufficiently large, does it follow that a Steiner (n, r, s) -system exists?

Question 3.3. Is it the case that for all $s < r$ there are infinitely many n such that a Steiner (n, r, s) -system exists?

Question 3.4. Do Steiner (n, r, s) -systems exist with s arbitrarily large?

Question 3.5. Do Steiner (n, r, s) -systems exist with r arbitrarily large?

For the rest of this note I want to discuss a remarkable recent result of Peter Keevash [8], though I will only be able to scratch the surface of what he did. For more detail, see an article of Gil Kalai [7], though this is still far from a complete description of Keevash's ideas. To set the scene, let me give some idea of what was known before his work.

First, it was known that Question 3.1 was too optimistic: there are triples (n, r, s) for which no Steiner (n, r, s) -system exists even though the existence of such a system is not ruled out on simple divisibility grounds. Second, the answer to Question 3.5 was known to be positive: in a famous sequence of papers in the 1970s, Richard Wilson proved that for any fixed r , a Steiner $(n, r, 2)$ -system exists provided that n is sufficiently large and satisfies the obviously necessary divisibility conditions [11–13]. (His proof also gives a similar result for all higher values of λ .)

However, the answers to the questions in between were not known. There have been many ingenious constructions of designs for specific values of n , r , and s (and λ if one makes that extra generalization), and a few small pairs (r, s) for which Question 3.2 has a positive answer. But to give an idea of our level of ignorance, no Steiner (n, r, s) -systems at all were known with $s > 5$. Thus, even an answer to Question 3.4 would have been a remarkable achievement. But Keevash went much further than this: he proved that the answer to Question 3.2 is yes! This is about as non-incremental a result as one could imagine: going from a situation where it was a huge struggle to prove the existence of even one design when s was even slightly large, to proving that for each r and s the only obstacle to the existence of a Steiner (n, r, s) -system was the trivial divisibility condition, for all but at most a finite set of n .

4. ARE PROBABILISTIC METHODS APPROPRIATE FOR THE PROBLEM?

In the previous section we saw a simple algebraic construction of a family of Steiner triple systems. That might suggest that explicit constructions are the right way of tackling this problem, and indeed much research in design theory has concerned using algebra to create interesting examples of designs.

On the other hand, the constraints on a Steiner system are not all that strong. Even for a triple system, if one wants to include all pairs in exactly one triple, the constraints are very slack: there is no sense at all in which one's moves are forced, at least to start with.

In fact, the situation here is rather similar to the situation with the problem about Hamilton cycles in middle-layer graphs. Suppose one just chooses sets of size r more or less arbitrarily, but making sure that no two of them intersect in a set of size s or more. (This ensures that no set of size s is contained in more than two sets from our collection.) For a long time we will have no problem, but eventually we will start to find that there are sets of size s that we do not seem to be able to cover with a set of size r without that set overlapping too much with a set we have already chosen. So the problem is difficult for a similar reason: there is too much choice for algebraic constructions to be easy to discover, and too little choice for simple probabilistic arguments to work.

Interestingly, however, it turns out that more sophisticated probabilistic arguments enable one to prove the existence of collections of sets that are *almost* Steiner systems. (This can mean one of two things: either one asks for no set of size s to be covered more than once and almost all sets of size s to be covered, or one asks for all sets of size s to be covered and almost no sets of size s to be covered more than once. It turns out that if you can achieve one, then you can achieve the other.) Even more interestingly, Keevash's proof uses an intriguing *mixture* of probabilistic and algebraic methods, thus having the best of both worlds, and reflecting the "intermediate" nature of the problem.

5. THE RÖDL NIBBLE

Let us return to the simplest non-trivial case of our general problem, that of Steiner triple systems. I shall give some indication of how probabilistic methods can be used to create almost Steiner triple systems. The basic technique was invented by Vojta Rödl [10], though an important precursor to it appeared in a paper of Ajtai, Komlós, and Szemerédi [1].

For this example, Rödl’s technique, which was dubbed the *Rödl nibble*, works as follows. Let us discuss the version of the problem where we are trying to cover all edges at least once and almost no edges more than once.

Let G_0 be the complete graph on n vertices. We begin by taking a small “bite” out of G_0 by choosing triangles randomly with probability $\epsilon/(n-2)$ for some small ϵ and removing them. The expected number of triangles we choose that contain any given edge xy is then ϵ (since there are $n-2$ possible z and each xyz has a probability $\epsilon/(n-2)$ of being chosen). We can in fact say more than this: the distribution of the number of triangles we choose that contain an edge xy is Poisson with mean ϵ .

Once we have done this, we will typically have covered a fraction roughly equal to ϵ of all the edges of the graph. We now eliminate all the edges in all the triangles that we have chosen and call the resulting graph G_1 .

We then repeat the process for G_1 . That is, we choose each triangle in G_1 independently with probability p , for a suitably chosen p , and remove all the edges in all the triangles we have chosen, obtaining a new graph G_2 . Note that none of the triangles we choose in this second bite share an edge with any of the triangles we chose in the first bite, since they are all triangles in G_1 .

But what is a suitable probability p ? There would not be a clear answer to this question were it not for the fact that G_1 has an extremely useful property that is crucial to Rödl’s argument: with very high probability it is *quasirandom*. I will not say exactly what a quasirandom graph is here but, roughly speaking, a quasirandom graph is one that behaves in many ways like a random graph of the same density. It is a remarkable and important fact that this can be defined in a precise and useful way—in fact, it can be defined in several ways that turn out to be equivalent, not always for trivial reasons.

Once we know that G_1 is quasirandom, we know that almost all edges will be contained in approximately the same number of triangles in G_1 . If this number is t , then we can choose the probability to be ϵ/t .

The Rödl nibble continues this process, showing at each stage that the graph G_r , that results after r bites have been taken, is quasirandom. Eventually, one loses control of the quasirandomness of G_r , but by that time there are so few edges that one can simply include each remaining edge in a triangle and one will not have covered too many edges more than once.

Rödl used a generalization of this argument to prove the existence of almost Steiner (n, r, s) -systems for all sufficiently large n . (If you do not require an exact Steiner system, then the divisibility constraints no longer apply.)

6. KEEVASH’S CONTRIBUTION

Before Peter Keevash’s work, the received wisdom was that with probabilistic methods, results like that of Rödl were the best one could hope for. Indeed, one can almost prove it: the probabilistic methods are insensitive to the divisibility constraints, and we know that Steiner systems do not exist when the divisibility constraints are not satisfied. Thus, it is very hard to see how a probabilistic argument could be made to work without also proving a false result. Of course this does not completely rule out using probabilistic methods—it just means that it is very hard to see how one could use them.

A natural if hopeless-looking preliminary idea is that one might try to use probabilistic methods to get almost all the way and then use some kind of clever local backtracking and adjustments right at the end to get from an almost Steiner system to an exact one. To oversimplify a lot, this is roughly what Keevash does. His proof is long and complicated, and there is no hope of giving a proper explanation in an article like this, so I will content myself with trying to explain roughly how algebra enters the picture.

Recall the example we saw earlier of a family of Steiner triple systems. More or less the same example (suitably generalized for the more general Steiner systems) plays an important role in Keevash's proof, but he uses the finite field \mathbb{F}^{2^k} rather than just the Abelian group \mathbb{F}_2^k .

More precisely, he chooses k such that $2n \leq 2^k < 4n$, takes a random map $\phi : V(G) \rightarrow \mathbb{F}^{2^k}$, and lets T be the system of triangles xyz such that $\phi(x) + \phi(y) + \phi(z) = 0$. As before, no two triangles in T share an edge. He then defines G^* to be the graph obtained by taking all the edges in all the triangles of T . This graph has positive density, but there are many edges not included.

Next, he uses the Rödl nibble to cover the complement of G^* as well as he can, but for the reasons discussed above, the best he can hope for is to cover *most* of the edges outside G^* in this way. Let H be the union of all these edges. So now he has an almost Steiner triple system, one part of which covers G^* and one part of which covers H .

Why is he in an any-better situation than one would be after simply applying the Rödl nibble? This is (as it must be) the key to his argument. Whereas, after applying the Rödl nibble, there is almost nothing useful one can say about the small graph that is not yet covered, we are now in a situation where we have the union of a small graph (about which we can still say almost nothing) and the graph G^* , about which it turns out that we can say rather a lot, thanks to its algebraic definition.

What kind of thing would one like to say about G^* ? To answer this question, let us think about what we plan to do next. Let E be the small graph that consists of the edges that we have not managed to cover. Obviously, we would now like to cover E , but there is no guarantee that E contains any triangles, so we are going to have to be prepared to modify the system of triangles we have already chosen. We do not touch any of the triangles used to cover H , because we do not know enough about it, but we are happy to touch the triangles used to cover G^* , because those are described very explicitly. As Keevash puts it in a key phrase, G^* “carries a rich structure of possible local modifications”. Roughly speaking, given an edge of E , there are rather a lot of ways in which one can try to contain it in a triangle with the other two edges in G^* , and once one does contain it, there are many ways in which one can “repair the damage” by removing the triangles used to cover those other two edges and then, in a clever way, organize for the triangles covering G^* to be adjusted slightly so that the edges that have just been uncovered get covered up again.

One might ask why it is necessary for G^* to be defined algebraically. Would it not be possible just to take G^* as a random union of edge-disjoint triangles? In fact, why bother with G^* at all? Why not just do the local modifications in H , which will look pretty random? Would the randomness not make it highly likely that local adjustments of the kind one wants to make are possible?

The answer to this turns out to be no. The reason is that algebraic constructions have a lot of small structures that one does not get with random constructions. For example, in both a random construction and an algebraic construction there will be a large number of triples of triangles of the form $xy'z', x'y'z', x'y'z$. But whereas in the random case the probability that xyz also belongs to the system of triangles is proportional to n^{-1} (because the number of triangles we choose is proportional to n^2 and the total number of triples of vertices is proportional to n^3 , and because there is almost no correlation between choosing the first three triangles and the fourth), in the algebraic case it is 1, since in \mathbb{F}^{2^k} the equations

$$x + y' + z' = x' + y + z' = x' + y' + z = 0$$

imply the equation $x + y + z = 0$ (as can be seen by adding together the three sums). Thus, an algebraically constructed triangle system is rich in configurations that look like four alternate faces of an octahedron, and in other small configurations of a similar kind, whereas a randomly constructed system is not.

Something like this idea works in general but becomes much more complicated—and it is already complicated even in the case of triangles and edges. Keevash uses an algebraic construction to define what he calls a “template”, proves that the template is rich in small configurations that can be used to make local adjustments, and then uses that richness to deal with the small collection of sets that is left over after the Rödl nibble has done what it can. Part of the reason that the proof is complicated is that it uses induction on s , so he has to consider simplicial complexes and not just collections of sets of a given size. Thus, the paper is a technical *tour de force* based on a beautiful underlying idea, and it solves one of the oldest problems in combinatorics.

ABOUT THE AUTHOR

Timothy Gowers is a Royal Society 2010 Anniversary Research Professor in Mathematics at the University of Cambridge and a fellow of Trinity College.

REFERENCES

- [1] M. Ajtai, J. Komlós, and E. Szemerédi, *A dense infinite Sidon sequence*, European J. Combin. **2** (1981), no. 1, 1–11, DOI 10.1016/S0195-6698(81)80014-5. MR611925
- [2] Noga Alon and Joel H. Spencer, *The probabilistic method*, fourth ed., Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons Inc., Hoboken, NJ, 2016, with an appendix on the life and work of Paul Erdős. MR2437651 (2009j:60004)
- [3] P. Erdős, *Some remarks on the theory of graphs*, Bull. Amer. Math. Soc. **53** (1947), 292–294. MR0019911
- [4] P. Erdős, *Graph theory and probability*, Canad. J. Math. **11** (1959), 34–38. MR0102081
- [5] M. Gromov, *Spaces and questions*, Geom. Funct. Anal. **Special Volume** (2000), 118–161. GAFA 2000 (Tel Aviv, 1999). MR1826251
- [6] Misha Gromov, *Random walks in random groups*, Geom. Funct. Anal. **13** (2003), 73–146.
- [7] G. Kalai, *Designs exist!* [after Peter Keevash], <http://www.bourbaki.ens.fr/TEXTES/1100.pdf>, Séminaire Bourbaki 67ème année, 2014–2015, no. 1100, 2015.
- [8] P. Keevash, *The existence of designs*, arxiv:1401.3665 [math.CO], 2014.
- [9] T. Mütze, *Proof of the middle levels conjecture*, Proc. Lond. Math. Soc. (3) **112** (2016), no. 4, 677–713, DOI 10.1112/plms/pdw004. MR3483129
- [10] V. Rödl, *On a packing and covering problem*, European J. Combin. **6** (1985), no. 1, 69–78, DOI 10.1016/S0195-6698(85)80023-8. MR793489
- [11] R. M. Wilson, *An existence theory for pairwise balanced designs. I. Composition theorems and morphisms*, J. Combinatorial Theory Ser. A **13** (1972), 220–245. MR0304203

- [12] R. M. Wilson, *An existence theory for pairwise balanced designs. II. The structure of PBD-closed sets and the existence conjectures*, J. Combinatorial Theory Ser. A **13** (1972), 246–273. MR0304204
- [13] R. M. Wilson, *An existence theory for pairwise balanced designs. III. Proof of the existence conjectures*, J. Combinatorial Theory Ser. A **18** (1975), 71–79. MR0366695

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CAMBRIDGE, 16 MILL LANE, CAMBRIDGE CB2 1SB, UNITED KINGDOM

E-mail address: `W.T.Gowers@dpms.cam.ac.uk`