Cycles and expansion in graphs

Richard Montgomery

Cycles are fundamental objects in graph theory, where their inherent simplicity belies the depth of even some simply stated questions. In this article, I will discuss three problems on cycles in graphs and recent progress on them. In each case, the progress has been made by new and different tools involving graph expansion, itself an important topic in extremal graph theory.

1 Eulerian graphs and the Erdős–Gallai problem

The advent of graph theory is often pinned to the Königsberg bridge problem from the 18th century. At the time, Königsberg had seven bridges connecting either side of the Pregel River and the two islands within it (see Figure 1). Was it possible to walk through the city while crossing each bridge exactly once? In 1735, this problem reached Euler, who comprehensively solved it in full generality. Representing each connected land mass by a vertex and each bridge by an edge between the two vertices it connects, we get a graph. Euler showed that there is a walk in a graph passing through every edge exactly once if and only if it is connected[1] and at most two vertices have odd degree[2] (if two such odd-degree vertices exist they must be the start and end vertices of the walk). Thus, there is no solution to the Königsberg bridge problem, as the corresponding graph has four vertices with odd degree.

A slightly neater equivalent formulation of this general problem is to ask when a graph has an *Eulerian tour*, which is a walk through the edges of the graph, covering each edge exactly once, and *arriving back at the start*. As Euler showed, there is an Eulerian tour if and only if the graph is connected and every vertex has even degree. Note that the general solution to the original bridge-crossing problem can be deduced by applying this, after first adding a fictitious edge between two odd-degree vertices when they exist.

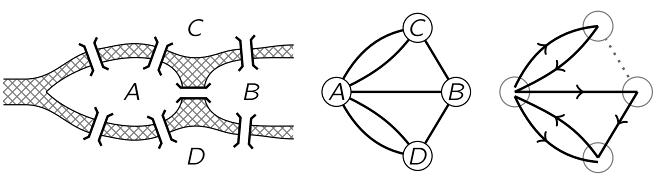

Figure 1. The Königsberg bridge problem, its representation as a (multi)graph and a walk crossing all but one edge/bridge.

[1] Any vertex can be reached from any other by a path in the edges.

[2] The number of edges containing that vertex.

Behind Euler's work is a simple result: if all the vertex degrees of a graph G are even, then its edges can be decomposed into cycles, i.e., its edges can be exactly partitioned into cycles (see Figure 2 for a similar decomposition). As removing the edges of a cycle maintains the parity of each vertex degree, this can be proved easily by induction, as every graph with even vertex degrees and at least one edge has at least one cycle (and on the other hand any graph decomposable into cycles has even vertex degrees). To reach Euler's result, then, first decompose any connected graph with even vertex degrees into cycles, $C_1, \dots, C_k$. Then, start off walking around the edges of C_1 but follow the rule that whenever we first encounter a vertex on a new cycle C_i, we break off and start walking around that cycle, giving an iterative process that can be seen to produce an Eulerian tour. On the other hand, the existence of an Eulerian tour in a graph easily shows that it must be connected and have even vertex degrees.

Above, in the space of a few lines, we have characterised with proof exactly those graphs which can be decomposed into cycles. However, with a slight change, we can reach a much deeper and more challenging problem. First, let us note the following: so far we have quietly been dealing with *multigraphs*, where a pair of vertices can have more than one edge between. From now on, every graph we consider will be assumed to be a *simple* graph, with at most one edge between each pair of vertices. Now, what if we ask for a decomposition of the graph into *few* cycles? In 1966, Erdős and Gallai [17] conjectured that every n-vertex Eulerian graph should have a decomposition into $O(n)$ cycles, in the following equivalent form.

Conjecture 1.1 (Erdős and Gallai). *Every n-vertex graph has a decomposition into $O(n)$ cycles and edges.*

Of course, if true, then this conjecture is tight up to the implicit constant, as demonstrated, for example, by the n-vertex graph with

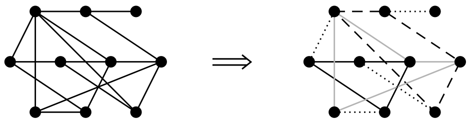

Figure 2. A graph on the left decomposed into three cycles and four (dotted) edges on the right.

all possible edges: when n is even any such decomposition needs at least $n/2$ edges and at least $(n-2)/2$ cycles. Erdős observed that a construction of Gallai can be improved to show that at least $(3/2 - o(1))n$ cycles and edges may be needed. Interestingly, the number of cycles and *paths* required to decompose any graph is well understood, thanks to an old result of Lovász [32], who unimprovably showed that any n-vertex graph can be decomposed into at most $\lceil n/2 \rceil$ paths and cycles.

As observed by Erdős and Gallai, $O(n \log n)$ cycles and edges can easily be seen to suffice for decomposing any n-vertex graph. Indeed, given an n-vertex graph G, you can iteratively remove a longest cycle until no cycles (and thus at most $n-1$ edges) remain. Any graph with average degree $d \geq 2$ has a subgraph with minimum degree at least $d/2$ and thus a cycle with length at least $d/2$ (seen, for example, by considering a longest path and the neighbouring vertices of one of its endpoints, which must lie within that path). Therefore, if we iteratively remove longest cycles from a graph with average degree d, after removing $O(n)$ cycles the average degree will be below $d/2$. Tracing the decrease in the average degree of G as longest cycles are removed, we therefore will remove $O(n \log n)$ cycles before only edges remain.

It took around 50 years for this simple bound to be improved, despite Erdős highlighting the problem in many of his problem collections. Finally, in 2014, it was shown by Conlon, Fox and Sudakov [9] that $O(n \log \log n)$ cycles and edges suffice to decompose any n-vertex graph. As discussed below, a critical concept behind this breakthrough was *expansion* in graphs. More recently, Bucić and I [7] were able to use a much more delicate form of this, known as *sublinear expansion*, to push this bound lower. Specifically, this allowed us to improve the $\log \log n$ term to the iterated logarithm function $\log^* n$, defined as the least k such that the k-fold logarithm of n, $\log(\log(\dots \log(n)))$, is at most 1. That is, we showed the following.

Theorem 1.2. *Any n-vertex graph decomposes into $O(n \log^* n)$ cycles and edges.*

Thanks to the result of Lovász quoted above, we know that any n-vertex graph decomposes into $O(n)$ cycles and *paths*. A potential strategy – rooted in a variety of combinatorial techniques from the last 50 years – would be to start by setting aside a small, perhaps randomly chosen, selection of the edges. Then, decomposing the remaining edges into a collection of few paths and cycles using Lovász's result, we could try to use the edges set aside to join the paths in the collection into cycles. Hopefully, then, we would have few cycles, and few edges remaining as well if we did not set many edges at the start.

This strategy is hard to achieve, but essentially was done by Conlon, Fox and Sudakov using a form of *graph expansion*. Expanders – graphs satisfying some type of graph expansion condition or other – are an important topic in their own right, both in combinatorics and in their application to computer science. Here we will concentrate on the use of graph expansion as a practical tool in extremal graph theory, and in particular for the study of cycles in graphs. For details on expanders more generally, we recommend the survey of Krivelevich [26] and its references.

In the simplest formulation, a graph expansion condition in a graph G might be that, for every set U of vertices which is not too large, there are plenty of vertices not in U which have at least one neighbouring edge to U (see Figure 3). More specifically, we might have, for some parameters m and α, that $|N_G(U)| \geq \alpha |U|$ for every set U of at most m vertices, where $N_G(U)$ is the set of vertices of G not in U which have a neighbouring edge to U (the *neighbourhood* of U).

The possible utility of expansion here is that it can allow us to find paths connecting any pair of vertices. For example, if G has n vertices, $m = n/3$ and $\alpha = 3/2$, we can consider iterative neighbourhoods $N_G(N_G(\dots N_G(x)))$ and $N_G(N_G(\dots N_G(y)))$ around any vertices x and y, respectively, and use the expansion condition to show that they grow in size exponentially until they each have more than $n/2$ vertices and thus overlap (see Figure 3). Thus, there must be a path from x to y, and moreover one which has length $O(\log n)$. This is only one path, however, and in the above very light sketch we would like to be able to connect many pairs of vertices (from the endpoints of the paths) by paths simultaneously while not using any edge more than once. Therefore, this is only some indication of how we might start to use expansion conditions to connect up paths into cycles, but the general principle is that if the expansion conditions are strong enough, then it might be possible even to find these connections iteratively in this fashion.

However, Conjecture 1.1 applies to all graphs, not just all graphs which conveniently satisfy some expansion condition! Conlon, Fox and Sudakov's first step, then, was to partition a graph into sub-

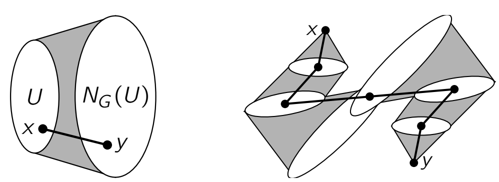

Figure 3. On the left, the neighbourhood $N_G(U)$ of a set U, the set of vertices y not in U for which there is some x in U such that xy is an edge. On the right, iteratively expanding about vertices x and y until a path (here of length six) between them is found.

graphs which satisfied some expansion conditions, before carrying out the outline above. This allowed them to show that, for any d, if an n-vertex graph has average degree d (and thus $dn/2$ edges in total), then $O(n)$ edge-disjoint cycles can be removed from it to leave a graph with average degree at most $d^{1-\varepsilon}$, for some small fixed $\varepsilon > 0$. Iterating this in $O(\log\log n)$ steps then gave a decomposition into at most $O(n\log\log n)$ cycles and edges.

As so often, there is a playoff: the stronger the expansion conditions we use, the easier it will be to connect up the paths into cycles, but we will need more edges remaining at each iteration to guarantee such strong conditions. Conlon, Fox and Sudakov used a strong expansion condition, where the ingenuity of their proof overall allowed them to work in subgraphs where, when they had t vertices remaining, say, the expansion condition corresponds very roughly to the above condition with $\alpha = t^{1-\mu}$ (though the precise condition used involves the number of edges leaving the set rather than the size of its neighbourhood).

In order to make much of an improvement to this, we need to use expansion conditions that can be found in an n-vertex graph with only $\log^{O(1)} n$ edges. To do this we must allow α to be much smaller. As α will need to be substantially less than 1, this is known as *sublinear expansion*, which was introduced in the 1990s by Komlós and Szemerédi [24, 25]. For illustration, a basic example of the type of expansion we might have here in an n-vertex graph G is that $|N_G(U)| \geq \frac{1}{\log^2 n}|U|$ for each vertex set U with at most $n/2$ vertices.

As this condition is so much weaker, whether or not we can work with it to some desired end often rests delicately on the details around the exact sublinear expansion conditions that can be achieved, but for simplicity here we will avoid these technicalities (for more details, and many applications and methods, see the recent comprehensive survey by Letzter [30]). Instead, let us emphasise a key point. These conditions may be weak, but in compensation essentially every graph has some subgraph satisfying some variation of these conditions. Indeed, as Komlós and Szemerédi showed, any graph contains within it a subgraph which has average degree not much smaller but which has some sublinear expansion properties.

For Theorem 1.2, Bucić and I followed the broad outline of Conlon, Fox and Sudakov [9] sketched above but using sublinear expansion conditions in place of the much stronger expansion conditions. To try this is a simple idea, the challenge is how hard it is to work with these much weaker conditions and this is what required considerable novelties and further work. It is particularly hard to show that, in very sparse sublinear expanders, a randomly chosen vertex set is likely to retain some expansion properties. However, we were able to show that if an n-vertex graph has average degree d, then we can remove $O(n)$ cycles to leave a graph with average degree $\log^{O(1)} d$. Thus, iterating this produces a graph with $O(n)$ edges in only $O(\log^* n)$ rounds, and therefore uses only $O(n\log^* n)$ cycles and edges in total.

While the iterated logarithm function grows extremely slowly with n, so that Theorem 1.2 comes within a hair's breadth of proving Conjecture 1.1, of course we should not be satisfied. The fundamental issue we currently have in our techniques is the iteration used, and more ideas appear needed, perhaps to decompose the graph more efficiently when this iteration is required by retrospectively lengthening the cycles found in previous rounds. Conjecture 1.1 is very simple to state, but its truth or falsity appears to reflect something deep about the structure of graphs, and its resolution may need to take into account different structural extremes that we do not yet understand.

2 Cycle lengths and the Erdős–Hajnal odd cycle problem

Any n-vertex graph with at least n edges has at least one cycle, but what can we say about its cycles? For any given k and n, how many edges in an n-vertex graph are sufficient to guarantee a cycle with length k, that is, with k edges? This (in wider generality) is the fundamental *Turán question* and a central part of extremal graph theory.

The case for odd cycles can be answered well without much trouble. An n-vertex graph can have a great many edges – $\lfloor n^2/4\rfloor$ – yet contain no odd cycle (see Figure 6), and a single edge more immediately gives any odd cycle with length less than around $n/2$. Far fewer edges are needed to guarantee any specific even cycle. Though, in this sense, it is easier to find an even cycle, getting a satisfactory answer to the Turán question is much harder. It has been known since the 1970s, due to Bondy and Simonovits [5], that, for each fixed k, a bound of the form $O(n^{1+1/k})$ on the number of edges can be sufficient to guarantee a cycle with length $2k$ in an n-vertex graph. However, while this is widely expected to be tight up to the implicit constant multiple, this is only known for $k = 2, 3$ and 5.

In n-vertex graphs with $n^{1+o(1)}$ edges, we cannot guarantee any particular cycle length k. Could we instead guarantee a graph has some cycle whose length lies within some sequence $k_1, k_2, k_3, \ldots$? Answering a question of Erdős, in 2005 Verstraëte [40] showed that there is some such increasing sequence with limiting density 0 for which there is some C such that, for any graph G with average degree at least C, G contains a cycle of length k_i for some $i \geq 1$. This proof was non-constructive and thus did not determine any particular sequence of lengths $k_1, k_2, \ldots$ which has this property. Erdős [15] asked in particular whether the powers of 2 might satisfy this property. In 2008, Sudakov and Verstraëte [37] were able to show that any n-vertex graph with no cycle whose length is a power of 2 must have average degree at most $e^{O(\log^* n)}$, where, again, $\log^* n$ is the iterated logarithm function. The powers of 2 in this result is only an example sequence: the proof works for any sequence $k_1, k_2, \ldots$ of even numbers in which each term is at most C times the previous term, for any fixed $C > 0$ (i.e., $k_{i+1} \leq Ck_i$ for each $i \geq 1$).

In the last few years, and using a new and fundamentally different approach, Liu and I [31] were able to improve this to answer Erdős's question, as follows.

Theorem 2.1. *There is some $d > 0$ such that every graph with average degree at least d has a cycle whose length is a power of 2.*

My methods with Liu apply to an even wider selection of sequences than those of Verstraëte and Sudakov, requiring only the sequence $k_1, k_2, \ldots$ consist of even numbers such that $k_{i+1} \leq \exp(k_i^{1/10})$ for each $i \geq 1$. The wide applicability of this result is perhaps its strongest quality, as for the powers of 2 the result is likely to hold for a much smaller constant than could be deduced from our methods. Indeed, Gyárfás and Erdős [16] conjectured that any graph with minimum degree at least 3 has a cycle whose length is a power of 2, as follows.

Conjecture 2.2. *Any graph with minimum degree at least 3 has a cycle whose length is a power of 2.*

Rather than only one cycle length from a sequence, we might also be able to say something about the set of cycles more generally. If a graph has average degree d and n vertices, then it may have no cycle shorter than $\Omega_d(\log n)$. However, if it has no such short cycles, then perhaps it correspondingly has many long cycles. Erdős and Hajnal suggested the harmonic sum of the cycle lengths as a measure of the density of the cycle lengths of a graph. Specifically, in 1966, they asked whether imposing a condition on the chromatic number[3] of a graph G is sufficient to force $\sum_{\ell \in C(G)} \frac{1}{\ell}$ to be large, where $C(G)$ is the set of integers ℓ for which there is a cycle of length ℓ in G.

Erdős later wrote that they felt a much weaker condition, one only on the average degree, should actually be sufficient. In 1984, this was confirmed by Gyárfas, Komlós, and Szemerédi [22], who moreover showed that any graph G with average degree d satisfies $\sum_{\ell \in C(G)} \frac{1}{\ell} \geq c \log d$ for some small constant $c > 0$. This is tight up to the value of c, as shown by the example of the complete bipartite graph with d vertices in each class (see Figure 6 on the left for a similar graph). This showed that c cannot be taken to be larger than $\frac{1}{2}$ here, and led Erdős [12] to suggest in 1975 that this should be the best possible asymptotically. Using our methods behind Theorem 2.1, Liu and I confirmed that this is the right bound and c can be taken to be arbitrarily close to $\frac{1}{2}$ for sufficiently large d. That is, the following is true.

Theorem 2.3. *Every graph G with average degree d satisfies $\sum_{\ell \in C(G)} \frac{1}{\ell} \geq (\frac{1}{2} - o_d(1)) \log d$.*

[3] The minimum number of colours required to colour the vertices so that no edge lies between two vertices of the same colour.

As noted above, there is a sharp distinction between odd and even cycles with regard to the Turán question, and graphs may have a very high average degree indeed and yet no odd cycle. Average degree in a graph G is therefore not the right parameter to determine the appearance of cycle lengths from a sequence of odd numbers. Here the original suggestion by Erdős and Hajnal from 1966 to consider the chromatic number $\chi(G)$, as mentioned above, is more promising. Let $C_{\text{odd}}(G)$ be the set of odd numbers appearing in $C(G)$. In 1981, Erdős and Hajnal [14] asked whether $\sum_{\ell \in C_{\text{odd}}(G)} \frac{1}{\ell} \to \infty$ if $\chi(G) \to \infty$. This is a more difficult problem than the corresponding question for average degree and all cycle lengths and as such was widely open, though in 2011 Sudakov and Verstraëte [38] showed that it is true under an additional condition imposed on the 'independence ratio' of G. Using additional ideas on top of the methods behind Theorems 2.1 and 2.3, Liu and I were able to build on our techniques to answer this, giving the following asymptotically-tight lower bound.

Theorem 2.4. *Every graph G with chromatic number at least k satisfies $\sum_{\ell \in C_{\text{odd}}(G)} \frac{1}{\ell} \geq (\frac{1}{2} - o(1)) \log k$.*

If G is a graph with k vertices and every possible edge, then G has chromatic number k and $\sum_{\ell \in C_{\text{odd}}(G)} \frac{1}{\ell}$ is the sum of the odd numbers in the interval $[3, k]$, and, thus, is equal to $(\frac{1}{2} - o(1)) \log k$. Therefore, the constant $\frac{1}{2}$ in Theorem 2.4 cannot be improved.

Unlike all the other results mentioned on and towards these problems, the progress made for Theorems 2.1, 2.3 and 2.4 uses sublinear expansion. Building on our previous work – both together and apart – Liu and I showed that, for any sublinear expander H with at least a large constant average degree, there is a long interval in which $C(H)$ contains every even number. As every graph G with at least a large constant average degree contains a sublinear expander with almost the same average degree, $C(G)$ thus also contains every even number from a long interval. These intervals are long enough (relative to their start) that they will always catch, for example, some power of 2. Moreover, summing the reciprocals of the even numbers in any such interval will lead to a proof of Theorem 2.3.

Again, the weak properties of sublinear expanders make them very difficult to work with, and the methods introduced by Liu and myself here required several key ideas that subsequently led to other work (see, for example, [19], and again the survey by Letzter [30]). While Conjecture 2.2 appears currently far beyond our current techniques, progress continues to be made. Indeed, in forthcoming work, Milojević, Pokrovskiy, Sudakov and I have been able to answer, for large d, a question of Erdős [13] by giving an exact extremal result corresponding to Theorem 2.3. That is, there is some d_0 such that for every $d \geq d_0$, among the n-vertex graphs G with at least $d(n-d)$ edges, $\sum_{\ell \in C(G)} \frac{1}{\ell}$ is minimised exactly by the n-vertex graph with every possible edge between a set of d vertices and a set of $n - d$ vertices, and no other edges.

3 Hamilton cycles in expanders

A *Hamilton cycle* in a graph is a cycle which passes through every vertex exactly once. It is named for the Irish astronomer and mathematician William Rowan Hamilton, who considered the following question in 1857. Can you walk along some of the edges of a dodecahedron, pass through all the corners exactly once, and finish where you started? Equivalently, does the graph formed from the corners and edges of the dodecahedron contain a Hamilton cycle (see Figure 4)?

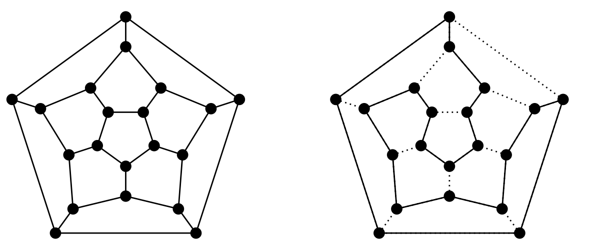

Figure 4. The graph of a dodecahedron, and a Hamilton cycle in it.

Hamilton was sufficiently enamoured with this question to sell the idea as a physical puzzle, which was sold in the UK and more widely in Europe as the 'Icosian game' and challenged the player to find such a cycle on a dodecahedron. Determining whether a graph contains a Hamilton cycle is a computationally difficult task. Indeed, this is one of Karp's original examples of an NP-complete problem. It is perhaps then ironic that Hamilton's Icosian game was a commercial failure in large part because it was too easy to solve! Hamilton may wish to be remembered more, then, for his invention of quaternions and his long period as the Royal Astronomer of Ireland.

Hamilton cycles in polyhedra had been considered a little earlier by the clergyman and amateur mathematician, T. P. Kirkman. Much earlier still, Hamilton cycles arose in the knight's tour problem, whose history dates back at least to the 9th century in India. In this problem, the knight is to make a sequence of legal moves around the chessboard so that it occupies every square exactly once (see Figure 5). This gives a path through every vertex (a *Hamilton path*) in the graph corresponding to its possible moves, and, if this can be done so that it can return immediately to its starting square, then it gives a Hamilton cycle. The knight's tour problem perhaps led to the first examples of Hamilton cycles in modern mathematics, with solutions for example given by Euler [18]. In the 20th century, Hamilton cycles have been an increasingly important object of study, with relevance for example to the travelling salesman problem.

The general difficulty of determining whether a graph has a Hamilton cycle has led to a lot of attention on proving simple conditions that imply a graph is Hamiltonian. For example, a stalwart of many a first course in graph theory is Dirac's theorem from 1953 that any graph with $n \geq 3$ vertices and minimum degree at least $n/2$ contains a Hamilton cycle. That this is tight is seen by two different *extremal examples* – an unbalanced complete bipartite graph and the disjoint union of two large complete graphs (see Figure 6). It is easy to determine whether Dirac's *minimum degree* condition for Hamiltonicity holds, or not, but it will only be satisfied by graphs with very many edges. This condition has been generalised to others concerning the degrees of the graph (for example, Ore's theorem), but all such conditions for Hamiltonicity require many edges if they are to be satisfied.

A famous condition for Hamiltonicity that can apply to sparser graphs is the Chvátal–Erdős condition from 1972. The corresponding result states that if the connectivity[4] of a graph G is at least as large as the independence number[5] of G, then it is Hamiltonian. Any n-vertex graph with average degree d is known to have an independent set with at least $n/(d+1)$ vertices, and the connectivity can easily be seen to be at most the average degree d. Thus, the Chvátal–Erdős condition can only hold in n-vertex graphs with average degree $\Omega(\sqrt{n})$.

Looking for conditions implying Hamiltonicity that apply to sparser graphs still, it is natural to consider two difficult and widely open conjectures from the 1970s. The first of these is by Chvátal and suggests a link between the *toughness* of a graph and its Hamiltonicity. A graph G is said to be *t-tough* if, for any s, the deletion of any set of s vertices from G gives either a connected graph, or one with at most s/t connected components. Any Hamiltonian graph, then, can be seen to be 1-tough. In 1973, Chvátal suggested that any sufficiently tough graph should be Hamiltonian, as follows.

Conjecture 3.1 (Chvátal). *There is some t such that any t-tough graph is Hamiltonian.*

In 2000, Bauer, Broersma and Veldman [3] showed that the stronger conjecture of Chvátal [8] that 2-toughness implies Hamiltonicity

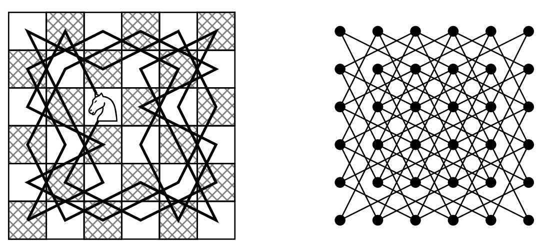

Figure 5. On the left, a closed knight's tour on a 6×6 chess board, which corresponds to a Hamilton cycle in the graph of legal moves on the right.

[4] The minimum number of vertices whose removal disconnects G.

[5] The maximum number of vertices with no edges between them.

is false, and indeed if Conjecture 3.1 is true, then t must be at least $9/4$. Conjecture 3.1 remains wide open.

Perhaps even more difficult is an elegant conjecture due to Lovász [33] from the 1960s, and a variant later given by Thomassen (see, e.g., [21]). Lovász conjectured that any vertex-transitive graph[6] which is connected contains a path that goes through every vertex exactly once. Thomassen suggested that only finitely many such graphs lack a cycle that goes through every vertex exactly once, as follows.

Conjecture 3.2 (Thomassen). *All but finitely many connected vertex-transitive graphs have a Hamilton cycle.*

There are only five known connected vertex-transitive non-Hamiltonian graphs with more than two vertices, all of which do not have many vertices. Note that the class of graphs which Conjecture 3.2 applies to contains the graph of the dodecahedron shown in Figure 4. Thus, in his very wide generalisation of the 'Icosian game,' Thomassen asks so much more of us than Hamilton ever did!

That Conjecture 3.2 is very challenging can be seen by our progress on the following question. Given any n-vertex vertex-transitive connected graph, how long a path or cycle can we find? (Hoping to eventually find an n-vertex path or cycle.) The best bound on this question for more than 40 years was that of Babai [2], who showed that in such graphs a cycle with length $\Omega(\sqrt{n})$ always exists. Only in 2023 was this improved, to $\Omega(n^{3/5})$ in a breakthrough by DeVos [10]. The current state of the art is that any n-vertex vertex-transitive connected graph has a cycle with length $\Omega(n^{13/21})$, due to Groenland, Longbrake, Steiner, Turcotte, and Yepremyan [21], and a path with length $\Omega(n^{9/14})$, due to Norin, Steiner, Thomassé, and Wollan [35].

Our discussion so far reflects the difficulty of finding sufficient conditions for Hamiltonicity applicable to sparse graphs. The sublinear expansion conditions we considered previously are not strong enough. In particular, if a graph G is Hamiltonian, then any vertex set A in G which contains no edges must satisfy $|N(A)| > |A|$. That is to say, we could have a local problem at any scale that looks like the first of the extremal examples in Figure 6 and prevents the existence of any Hamilton cycle. It is natural, then, to consider expansion conditions which might imply Hamiltonicity, as has been done since the pioneering work of Pósa [36] on Hamilton cycles in random graphs.

Building on Pósa's work, whether a Hamilton cycle is likely to appear or not in the most studied random graph models is very well understood, due to work by Bollobás [4] and by Ajtai, Komlós, and Szemerédi [1]. In all the corresponding methods, expansion conditions (as are likely to occur in the random graphs) are used in conjunction with random techniques. It is desirable then to find simple properties likely to hold in random graphs which will imply Hamiltonicity and avoid studying the random graph model directly. That is, which *pseudorandom* conditions in sparse graphs are sufficient to imply Hamiltonicity?

The study of pseudorandom graphs was begun by Thomason [39] in the 1980s, creating an active and influential area of research (for more on which, see, for example, the survey of Krivelevich and Sudakov [28]). A major class of graphs known to exhibit pseudorandom properties are (n, d, λ)-graphs: n-vertex d-regular[7] graphs satisfying a certain condition (governed by λ) on the eigenvalues of their adjacency matrices. In their foundational study of (n, d, λ)-graphs in 2003, Krivelevich and Sudakov [27] conjectured that if $d/\lambda \geq C$ (for some universal constant $C > 0$), then any (n, d, λ)-graph is Hamiltonian. As evidence, they showed that, if $d/\lambda \geq \log n$ (for large n), then any (n, d, λ)-graph is Hamiltonian. For more details on this, and (n, d, λ)-graphs, see [27]. For our purposes now, however, we will consider the following conjecture (appearing, for example, in [6]), which implies the conjecture of Krivelevich and Sudakov on the Hamiltonicity of (n, d, λ)-graphs, and suggests pseudorandom expansion conditions for Hamiltonicity.

Conjecture 3.3. *There exists $C > 0$ such that any n-vertex graph satisfying the following two conditions is Hamiltonian.*
1. $|N(A)| \geq C|A|$ *for any vertex set A of at most $n/2C$ vertices.*
2. *For any disjoint vertex sets A, B of at least n/C vertices each, there is an edge between A and B in G.*

On a personal note, Conjecture 3.3 is a problem I first studied when writing my PhD thesis in 2015. There, to aid a long argument embedding certain graphs known as spanning trees into random graphs [34], an answer to this question would have been very convenient! In an example of how it is often easier to study random graphs rather than pseudorandom graphs, I was able to avoid the

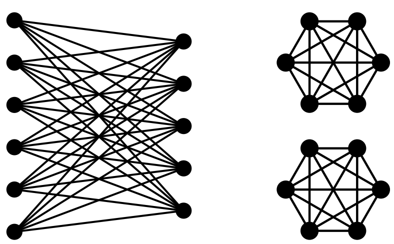

Figure 6. Extremal graphs with high minimum degree yet no Hamilton cycle. On the left, the graph with 11 vertices has no Hamilton cycle, as it cannot alternate between the right and left due to the imbalance in the number of vertices. On the right, the graph with 12 vertices has no Hamilton cycle, as it is not connected.

[6] I.e., a graph G in which, for any pair of vertices, there is an isomorphism of G mapping one to the other.

[7] I.e., every vertex has degree d.

difficulty of Conjecture 3.3 by working in random graphs directly using some methods of Sudakov and Lee [29].

In 2009, Hefetz, Krivelevich, and Szabó [23] made progress towards Conjecture 3.3, essentially showing this is true if the constant expansion coefficient C is replaced by $\log n$. In 2024, Glock, Munhá Correia, and Sudakov [20] gave an improved bound on the conjecture of Krivelevich and Sudakov on the Hamiltonicity of (n, d, λ)-graphs, improving the condition $d/\lambda \geq \log n$ to $d/\lambda \geq (\log n)^{1/3}$. Very recently, I was able to confirm the conjecture in full with Draganić, Munhá Correia, Pokrovskiy, and Sudakov [11], which followed from our positive resolution of Conjecture 3.3. This confirms the most general, natural, conditions for Hamiltonicity which apply to even very sparse graphs, but the openness of Conjectures 3.1 and 3.2 shows that much remains to be done in this fascinating area.

References

[1] M. Ajtai, J. Komlós and E. Szemerédi, First occurrence of Hamilton cycles in random graphs. In *Cycles in graphs*, North-Holland Math. Stud. 115, North-Holland, Amsterdam, 173–178 (1985)

[2] L. Babai, Long cycles in vertex-transitive graphs. *J. Graph Theory* **3**, 301–304 (1979)

[3] D. Bauer, H. J. Broersma and H. J. Veldman, Not every 2-tough graph is Hamiltonian. *Discrete Appl. Math.* **99**, 317–321 (2000)

[4] B. Bollobás, The evolution of sparse graphs. In *Graph theory and combinatorics (Cambridge, 1983)*, Academic Press, London, 35–57 (1984)

[5] J. A. Bondy and M. Simonovits, Cycles of even length in graphs. *J. Combinatorial Theory Ser. B* **16**, 97–105 (1974)

[6] S. Brandt, H. Broersma, R. Diestel and M. Kriesell, Global connectivity and expansion: long cycles and factors in f-connected graphs. *Combinatorica* **26**, 17–36 (2006)

[7] M. Bucić and R. Montgomery, Towards the Erdős–Gallai cycle decomposition conjecture. *Adv. Math.* **437**, article no. 109434 (2024)

[8] V. Chvátal, Tough graphs and Hamiltonian circuits. *Discrete Math.* **5**, 215–228 (1973)

[9] D. Conlon, J. Fox and B. Sudakov, Cycle packing. *Random Structures Algorithms* **45**, 608–626 (2014)

[10] M. DeVos, Longer cycles in vertex transitive graphs. arXiv: 2302.04255v1 (2023)

[11] N. Draganić, R. Montgomery, D. M. Correia, A. Pokrovskiy and B. Sudakov, Hamiltonicity of expanders: optimal bounds and applications. arXiv:2402.06603v2 (2024)

[12] P. Erdős, Some recent progress on extremal problems in graph theory. In *Proceedings of the Sixth Southeastern Conference on Combinatorics, Graph Theory and Computing (Florida Atlantic Univ., Boca Raton, FL, 1975)*, Congress. Numer. 14, Utilitas Math., Winnipeg, 3–14 (1975)

[13] P. Erdős, On the combinatorial problems which I would most like to see solved. *Combinatorica* **1**, 25–42 (1981)

[14] P. Erdős, Problems and results in graph theory. In *The theory and applications of graphs (Kalamazoo, MI, 1980)*, Wiley, New York, 331–341 (1981)

[15] P. Erdős, Some new and old problems on chromatic graphs. In *Combinatorics and applications (Calcutta, 1982)*, Indian Statist. Inst., Calcutta, 118–126 (1984)

[16] P. Erdős, Some problems in number theory, combinatorics and combinatorial geometry. *Math. Pannon.* **5**, 261–269 (1994)

[17] P. Erdős, A. W. Goodman and L. Pósa, The representation of a graph by set intersections. *Canadian J. Math.* **18**, 106–112 (1966)

[18] L. Euler, Solution d'une question curieuse que ne paroit soumise à aucune analyse. In *Mémoires de l'académie des sciences de Berlin*, Vol. 15, 310–337 (1766)

[19] I. G. Fernández, J. Kim, Y. Kim and H. Liu, Nested cycles with no geometric crossings. *Proc. Amer. Math. Soc. Ser. B* **9**, 22–32 (2022)

[20] S. Glock, D. Munhá Correia and B. Sudakov, Hamilton cycles in pseudorandom graphs. *Adv. Math.* **458**, article no. 109984 (2024)

[21] C. Groenland, S. Longbrake, R. Steiner, J. Turcotte and L. Yepremyan, Longest cycles in vertex-transitive and highly connected graphs. *Bull. Lond. Math. Soc.* **57**, 2975–2990 (2025)

[22] A. Gyárfás, J. Komlós and E. Szemerédi, On the distribution of cycle lengths in graphs. *J. Graph Theory* **8**, 441–462 (1984)

[23] D. Hefetz, M. Krivelevich and T. Szabó, Hamilton cycles in highly connected and expanding graphs. *Combinatorica* **29**, 547–568 (2009)

[24] J. Komlós and E. Szemerédi, Topological cliques in graphs. *Combin. Probab. Comput.* **3**, 247–256 (1994)

[25] J. Komlós and E. Szemerédi, Topological cliques in graphs. II. *Combin. Probab. Comput.* **5**, 79–90 (1996)

[26] M. Krivelevich, Expanders – how to find them, and what to find in them. In *Surveys in combinatorics 2019*, London Math. Soc. Lecture Note Ser. 456, Cambridge Univ. Press, Cambridge, 115–142 (2019)

[27] M. Krivelevich and B. Sudakov, Sparse pseudo-random graphs are Hamiltonian. *J. Graph Theory* **42**, 17–33 (2003)

[28] M. Krivelevich and B. Sudakov, Pseudo-random graphs. In *More sets, graphs and numbers: A salute to Vera Sós and András Hajnal*, Bolyai Soc. Math. Stud. 15, Springer, Berlin, and Janos Bolyai Mathematical Society, Budapest, 199–262 (2006)

[29] C. Lee and B. Sudakov, Dirac's theorem for random graphs. *Random Structures Algorithms* **41**, 293–305 (2012)

[30] S. Letzter, Sublinear expanders and their applications. In *Surveys in combinatorics 2024*, London Math. Soc. Lecture Note Ser. 493, Cambridge Univ. Press, Cambridge, 89–130 (2024)

[31] H. Liu and R. Montgomery, A solution to Erdős and Hajnal's odd cycle problem. *J. Amer. Math. Soc.* **36**, 1191–1234 (2023)

[32] L. Lovász, On covering of graphs. In *Theory of Graphs (Proc. Colloq., Tihany, 1966)*, Academic Press, New York, London, 231–236 (1968)

[33] L. Lovász, The factorization of graphs. In *Combinatorial structures and their applications, Proc. Calgary Internat. Conf., Calgary, Alberta, 1969*, Gordon and Breach, New York, 243–246 (1970)

[34] R. Montgomery, Spanning trees in random graphs. *Adv. Math.* **356**, article no. 106793 (2019)

[35] S. Norin, R. Steiner, S. Thomassé and P. Wollan, Small hitting sets for longest paths and cycles. arXiv:2505.08634v2 (2025)

[36] L. Pósa, Hamiltonian circuits in random graphs. *Discrete Math.* **14**, 359–364 (1976)

[37] B. Sudakov and J. Verstraëte, Cycle lengths in sparse graphs. *Combinatorica* **28**, 357–372 (2008)

[38] B. Sudakov and J. Verstraëte, Cycles in graphs with large independence ratio. *J. Comb.* **2**, 83–102 (2011)

[39] A. Thomason, Pseudo-random graphs. In *Random graphs '85 (Poznań, 1985)*, North-Holland Math. Stud. 144, North-Holland, Amsterdam, 307–331 (1987)

[40] J. Verstraete, Unavoidable cycle lengths in graphs. *J. Graph Theory* **49**, 151–167 (2005)

Richard Montgomery is a professor of mathematics at the University of Warwick, working in extremal and probabilistic combinatorics. He received his PhD in 2015 at the University of Cambridge under the supervision of Andrew Thomason, after which he spent time at Trinity College, Cambridge, and then the University of Birmingham. Recognition of his work includes a 2024 EMS Prize and a 2025 LMS Whitehead Prize.

richard.montgomery@warwick.ac.uk