COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **22** (1) (2002) 1–18

RANDOM GRAPH COVERINGS I: GENERAL THEORY AND GRAPH CONNECTIVITY

ALON AMIT, NATHAN LINIAL*

Received June 21, 1999
Revised November 16, 2000

In this paper we describe a simple model for random graphs that have an n-fold covering map onto a fixed finite base graph. Roughly, given a base graph G and an integer n, we form a random graph by replacing each vertex of G by a set of n vertices, and joining these sets by random matchings whenever the corresponding vertices are adjacent in G. The resulting graph covers the original graph in the sense that the two are locally isomorphic. We suggest possible applications of the model, such as constructing graphs with extremal properties in a more controlled fashion than offered by the standard random models, and also "randomizing" given graphs. The main specific result that we prove here (Theorem 1) is that if $\delta \geq 3$ is the smallest vertex degree in G, then almost all n-covers of G are δ-connected. In subsequent papers we will address other graph properties, such as girth, expansion and chromatic number.

1. Introduction

The notion of covering maps between graphs is essentially a restriction to the case of graphs (as, say, one dimensional simplicial complexes) of the general topological notion of covering map. It is described in purely combinatorial terms as a mapping of graphs that maps the neighbours of a vertex one-to-one onto the neighbours of its image vertex (We will later refine this slightly to properly account for multiple edges and loops).

Covering maps have received considerable attention from several different points of view. For example, Leighton's remarkable theorem on common

Mathematics Subject Classification (2000): 05C80, 05C10, 05C40

* Work supported in part by grants from the Israel Academy of Aciences and the Binational Israel-US Science Foundation.

finite coverings [8] was inspired by Angluin's application of coverings to problems in multi-processor networks [2], and in turn was found to have connections with deep algebraic questions related to groups acting on trees (Bass and Kulkarni [4]). Graph coverings are naturally related to subgroups of free groups; see Stillwell [12] and Stallings [11], for example. Finally, as purely graph theoretic objects, they were studied by Gross and Tucker [6], [7], Negami [10] and Archdeacon and Richter [3], among others.

The main theme of this work is the introduction of a probabilistic structure on the set of graphs that cover a fixed base graph. This enables one to apply the powerful probabilistic method to questions concerning coverings. We may investigate what a *typical* covering looks like, and for instance prove the existence of coverings having properties that may be hard to construct explicitly. Some results of this kind appear in our subsequent [1].

Moreover, we can "forget" that the random graphs we generate cover some base graph, and view our model as some sort of random model for (a special class of) graphs. Inasmuch as the probabilistic method is used to prove existence theorems of graphs with extremal properties, a new model for random graphs (complete with "knobs" we can fiddle with) may be a useful addition to our arsenal.

In a different direction, we may hope that the model can be used to "randomize" a given graph, in the following sense. Suppose we wish to calculate, or approximate, some invariant of a graph G. It is often the case that the invariant is hard to determine efficiently, but an efficient algorithm is known to work on a *random* graph. We can now run the algorithm on a random covering $\tilde{G}$ of G, hoping that it is "sufficiently random" so that the algorithm runs reliably, but still "sufficiently similar" to G itself so the result is close to the original invariant. We hope to pursue this idea in subsequent papers.

In the following subsections we describe the model and its basic properties. Section 2 contains our first application of the model, showing that almost every covering of a graph with minimal degree δ is δ-connected. In Section 3 we present variations and generalizations of the model, some of which we will utilize in forthcoming papers.

1.1. Coverings

Let G and $\tilde{G}$ be graphs (possibly with multiple edges and loops), and let $\pi : \tilde{G} \to G$ be a graph homomorphism, namely a pair of maps $V(\tilde{G}) \to V(G)$ and $E(\tilde{G}) \to E(G)$ that preserve adjacency. π is a *covering* if the star of each vertex $\tilde{v} \in \tilde{G}$ is mapped bijectively to the star of its image $\pi\tilde{v}$, where a star

is the collection of edges emanating from a vertex. A loop is considered as *two* edges in the star.

We call G the *base graph* and the inverse image $\pi^{-1}(v)$ is called a *fiber*, denoted $\tilde{G}_v$. We use suggestive terms such as "$\tilde{v}$ lies above v" when $\pi\tilde{v}=v$ etc. Indeed it is best to visualize $\tilde{G}_v$ as a vertical stack of vertices above v. It is easy to see that the inverse image of an edge $[u,v]$ sets a bijective correspondence between $\tilde{G}_u$ and $\tilde{G}_v$, since each $\tilde{u}$ in $\tilde{G}_u$ is connected to precisely one $\tilde{v}$ in $\tilde{G}_v$.

In particular all the fibers have the same cardinality (when the base graph is connected). This common cardinality is called the *degree* of the covering; if this is finite and equal to n we call π an n-covering. Every path in G is covered by n disjoint paths in $\tilde{G}$, called the *liftings* of the path. This is the well-known *unique path-lifting* property of coverings.

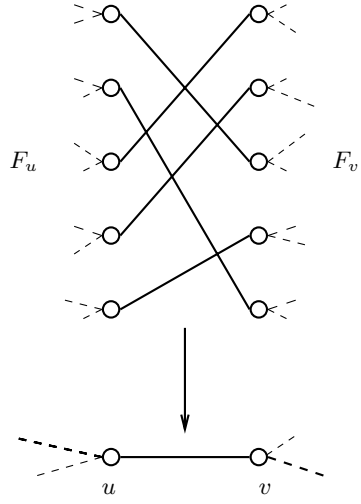

Fig. 1. Typical adjacent fibers

Figure 1. shows part of G and the part of a typical $\tilde{G}$ covering it.

1.2. The Model

Given the base graph G and the degree n, we can construct a covering graph $\tilde{G}$ by putting n vertices $v_1,\ldots,v_n$ above every vertex v of G. For every edge $e=[u,v]$, we need to decide which u_i is connected to which v_j. This is determined by a single permutation on n elements: Given a permutation

$\sigma = \sigma(e)$ we connect u_i to $v_{\sigma(i)}$, and this is done independently for every edge. For example, if we number the vertices in Figure 1. from the bottom up, the corresponding permutation has cycle presentation $(124)(35)$.

Notice that we need to assign an orientation to e so that we know which end is u and which is v. This choice, however, is arbitrary and has no real effect on the possible outcomes, i.e. reversing the edge and inverting the permutation yield the same covering. Having made this choice, we see that a covering is defined by attaching permutations to the edges of G. All n-coverings of G are obtained in this manner.

It is now fairly obvious how to define a *random* n-covering: simply choose a permutation $\sigma(e) \in S_n$ uniformly and independently for every edge e in G, and form the covering $\tilde{G}$ as above. This is our model for a random *labeled* n-covering of G (see 1.3). Already we see several ways we can tweak it to suit our needs: we may choose the permutations non-uniformly in S_n, and possibly not independently. Unless otherwise stated, however, we will work with this standard model. We summarize it in the following definition:

Definition 1. Given a graph G, a *random labeled n-covering* of G is obtained by arbitrarily orienting the edges of G, choosing a permutation σ_e in S_n for each edge e uniformly and independently, and constructing the graph $\tilde{G}$ with n vertices $u_1, \ldots, u_n$ for each vertex u of G and edges $e_i = (u_i, v_{\sigma_e(i)})$ whenever $e = (u, v)$ is an oriented edge. A covering $\pi : \tilde{G} \to G$ is defined by $\pi(u_i) = u$ and $\pi(e_i) = e$.

The random model is well-defined in the context of pseudographs as well. Different permutations can be assigned to parallel edges, and a single permutation is attached to each loop.

1.3. The Role of Labels

We note that the model actually gives a little more than what we required, since the resulting graph $\tilde{G}$ is equipped with a labelling $\{1, \ldots, n\}$ of the vertices in each fiber. A covering $\pi : \tilde{G} \to G$ with $\tilde{G}$ so labeled is called a *labeled* covering, and it is these objects for which we have a random model. The situation is analogous to that of random graphs, where the standard model is defined for labeled graphs instead of (isomorphism classes of) abstract graphs.

Let $L_n(G)$ be the set of labeled n-coverings of G. The same notation is used also for the probability space we have just defined. Clearly, the probability distribution is uniform. Let $C_n(G)$ be the set of isomorphism classes of unlabeled n-coverings of G. Obviously $|L_n(G)| = (n!)^e$ where $e = |E(G)|$ is

the number of edges in G. By considering the natural action of S_n^v on $L_n(G)$ (here $v = |V(G)|$, and S_n^v acts by relabeling the vertices in each fiber) one can show that

$$|C_n(G)| = \sum_{\mathcal{C}} s(\mathcal{C})^{e-v} \tag{1}$$

where the sum is taken over conjugacy classes in S_n, and $s(\mathcal{C})$ is the size of the centralizer of any element in $\mathcal{C}$. Estimating this sum we find that, when $e - v > 0$,

$$|C_n(G)| = (1 + o(1))(n!)^{e-v}.$$

The details of this argument will be published elsewhere. It follows that almost every covering has a trivial automorphism group, and more importantly, that properties of coverings have the same asymptotic distribution in the labeled and unlabeled models. (This is similar to the situtation with random graphs $G(n,p)$. Compare with [5], chapter IX). We may therefore adopt the standard policy of proving results in the labeled model, and viewing them as valid statements in the unlabeled case as well.

Finally, it is worth noting that imposing certain restrictions on the random permutations yields an equivalent model. Specifically, it is not difficult to prove the following: If E is a set of edges that does not contain a cycle, then the probability of any graphical property of the covering is unchanged if we condition on all the permutations assigned to edges in E being the identity.

2. Connectivity

For a fixed base graph G, what is the typical degree of connectivity of a random (labeled) covering of G? If $\delta = \delta(G)$ is the minimal degree of G, it is also the minimal degree in every covering of G. Therefore no covering can have connectivity higher than δ. We prove that a random covering is indeed almost surely δ-connected when $\delta \geq 3$.

Theorem 1. *Let G be a connected simple graph with minimal degree $\delta \geq 3$. Then with probability $1 - o(1)$, a random labeled n-covering of G is δ-connected.*

The $o(1)$ term is taken with respect to n, of course. Here, and throughout, we use the term "almost surely" instead of "with probability $1 - o(1)$ as n tends to infinity", so the theorem says that a random labeled covering is almost surely δ-connected. As mentioned above, the same statement also applies to random unlabeled coverings.

We point out that for some base graphs G, it is not even true that a random covering is likely to be connected. A trivial example is that of a base graph that is a tree T: *Every* covering of T is disconnected, except for the 1-covering T itself. Another case is when G is unicyclic: a tree T with one extra edge (i.e., a graph with zero Euler characteristic). In this case, a random covering is essentially determined by a single permutation σ (see the discussion of flattening in Section 1.3), and it is connected iff σ is cyclic, which happens only with probability $1/n$. There is no contradiction, though, since both these examples fail to satisfy $\delta \geq 3$. See also the discussion of $K_2(2)$ and $K_2(3)$ in Section 2.2.

Let us briefly indicate the structure of the proof. After some preliminaries we begin in Section 2.2 by considering connectivity properties of coverings of $K_2(\alpha)$, the graph with two vertices and α edges. Usually, such coverings are highly connected: every set of vertices has many neighbours. This is essentially just a property of random independent permutations. The "usually" here refers to the fact that for $\alpha = 1, 2$ we don't get high connectivity at all, and even when $\alpha = 3$, very small sets may fail to have enough neighbours. It is this difficulty with seemingly minor cases that makes the proof a bit messy.

Given a general base graph G, we consider a set X of vertices in the covering $\tilde{G}$, and try to show that it has a large vertex boundary. We look for a topological $K_2(\alpha)$ inside G over which X has enough vertices to apply the previous analysis. This works well for most sets X, namely those that have at least two vertices in some fiber. The special cases $\alpha = 1, 2$, however, require some additional effort. This is handled in Section 2.3. Section 2.4 deals with "thin" sets, that have at most one vertex in each fiber.

2.1. Preliminaries

To show that a covering $\tilde{G}$ is δ-connected, we need to show that for every set X of vertices with $|X| < |\tilde{G}|/2$, $|\partial X| \geq \delta$ where ∂X is the set of vertices outside of X that are adjacent to some vertex in X. An important characteristic of sets X in $\tilde{G}$ is the way they are distributed across the fibers. Let $X_v = X \cap \tilde{G}_v$ be the set of vertices of X that lie above a vertex $v \in V(G)$, and $x_v = |X_v|$ its size. We begin with a simple observation concerning these numbers.

Suppose that for some u, v in G, $x_u - x_v \geq \delta$. Since G is connected, there is a u-v path P in G. Consider the liftings of P to paths in $\tilde{G}$ that start at a vertex in X_u. They are disjoint, so at least δ of them end in a vertex *outside* X_v. Therefore each such path contains a vertex in ∂X (possibly more than

one), so X satisfies $|\partial X| \geq \delta$ automatically, regardless of the actual covering. In Figure 2., $|\partial X|$ is even larger than $x_u - x_v$, since the third (from the top) lifted path contributes more than one vertex to ∂X.

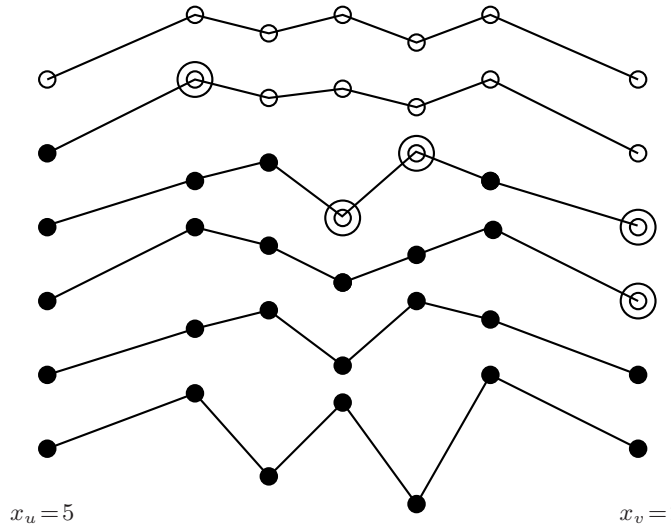

Fig. 2. An Example with $x_u - x_v = 3$. The black vertices are in X, the white ones outside X, and the circled white are in ∂X.

Thus, in proving δ-connectivity, we may restrict to sets X that are rather evenly distributed across the fibers, in the sense that

$$|x_u - x_v| \leq \delta - 1 \tag{2}$$

for every $u, v \in G$.

Let $m = \max\{x_v \,|\, v \in G\}$. A set X for which $m = 1$ is called *thin*. Thus a thin set is a set of vertices of $\tilde{G}$ which contains at most one vertex from each fiber. X is called *thick* if it is not thin.

For distinct vertices a, b in a graph G, define their local connectivity as the maximal number of vertex disjoint paths in G between them. This is denoted $\kappa(a,b) = \kappa(a,b;G)$. The maximum local connectivity, i.e. the maximum of $\kappa(a,b;G)$ over all pairs of vertices of G, is denoted $\bar{\kappa}(G)$. Note that the analogous *minimum* is just the vertex connectivity of G.

We need the following result of Mader:

Theorem 2 (Mader [9]). *In every finite graph G there is an edge $[a,b]$ such that $\kappa(a,b) = \min(\deg(a), \deg(b))$.*

2.2. Coverings of $K_2(\alpha)$

Lemma 1. *Let $\tilde{G}$ be a random n-covering of $G = K_2(\alpha)$, the graph with two vertices and α edges, where $\alpha \geq 3$. Then a.s. every subset $X \subset V(\tilde{G})$ such that $3 \leq |X| < 2|\tilde{G}|/3$ satisfies $|\partial X| \geq \alpha$.*

It is interesting to see why the restrictions in the proposition are necessary. For $\alpha = 1$, $K_2(1) = K_2$ is a tree with one edge, and every covering is just a matching, so there are many sets X with $|\partial X| = 0$.

For $\alpha = 2$, a covering is determined by two permutations σ_1 and σ_2, given that an orientation has been assigned for the two edges, say from a to b where a, b are the vertices of $K_2(2)$. Suppose that i is a fixed point of $\sigma_2^{-1}\sigma_1$, so that $\sigma_1(i) = \sigma_2(i) = j$, say. Then the set $X = \{a_i, b_j\}$ spans a $K_2(2)$ in $\tilde{G}$, and is disconnected from the rest of the graph. For every constant k the probability that there are k such fixed points is bounded away from zero, and so there are likely to be several such small disconnected sets.

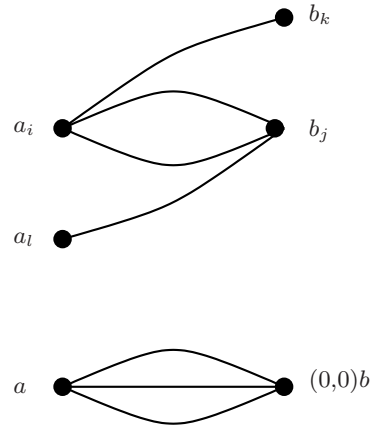

Fig. 3. A bad set in a covering of $K_2(3)$

When $\alpha = 3$, a typical covering is *nearly* 3-connected, but not quite. The covering is determined by three permutations $\sigma_1, \sigma_2, \sigma_3$. Consider again an index i for which $\sigma_1(i) = \sigma_2(i) = j$. Then a_i is connected to b_j by two parallel edges, and is further connected to b_k where $k = \sigma_3(i)$. Also b_j is connected to some a_l where $l = \sigma_3^{-1}(j)$. This yields a set $X = \{a_i, b_j\}$ which only has 2 neighbours, as in Figure 3. This shows why $2 < |X|$ is necessary in Proposition 1. Incidentally, this also demonstrates why simplicity is required in Theorem 1.

Proof (of the Lemma). Let $\{a,b\}$ be the vertices of $G=K_2(\alpha)$. A random covering of G is determined by α random permutations $\sigma_1,\ldots,\sigma_\alpha$ in S_n, by connecting each vertex a_i to the α (not necessarily distinct) vertices $a_{\sigma_p(i)}$, $1\le p\le\alpha$.

For a subset $A\subset\{1,\ldots,n\}$, let $A_p=\sigma_p(A)$ and $A_*=\cup_p A_p$. The letter m always denotes the size of A. Call a set A *bad* if $2\le m\le 3n/4$ and $|A_*|<m+\alpha-\epsilon$ where ϵ is 0 unless $m=2$ and $\alpha=3$, in which case it is 1. (The role of ϵ will soon be revealed). We claim that, if the σ_ps are chosen uniformly and independently in S_n, then a.s. there are no bad sets.

Proof of claim. We start with small sets A, those for which $m\le 10\alpha$, say. Given A, the probability that $|A_*|\le t$ for some t is not greater than the probability that $s=\alpha m$ balls that are thrown into n cells occupy no more than t cells, which is bounded by

$$\binom{n}{t}\left(\frac{t}{n}\right)^s.$$

Since there are $\binom{n}{m}$ sets A with $|A|=m$, to show that no set of size m is bad a.s. it suffices to verify that

$$\binom{n}{m}\binom{n}{t}\left(\frac{t}{n}\right)^s\to 0$$

as $n\to\infty$, where $s=\alpha m$ and $t=\alpha+m-\epsilon-1$. By our assumption that $m\le 10\alpha$, s and t are bounded by constants so the expression is $\Theta(n^{m+t-s})$. Now $m+t-s=(2-\alpha)m+\alpha-\epsilon-1$ which is always negative. Note that, when $m=2,\alpha=3$, it is not zero only thanks to the ϵ. It is indeed true that given three permutations there is likely to exist a set of size two whose image contains just $4<2+3$ elements.

For large sets A, $10\alpha<m<3n/4$. $|A_*|$ is the union of α randomly chosen sets of size m, and we need to bound the probability that this union has size $m+\alpha$ or less. A union bound over the possible sets of size $m+\alpha$ gives (again recalling that we have $\binom{n}{m}$ possible candidates for A itself)

$$B(m)=\binom{n}{m}\binom{n}{m+\alpha}\left(\binom{m+\alpha}{m}/\binom{n}{m}\right)^\alpha$$

and we must show that this tends to 0 as $n\to\infty$. The easy inequality

$$\frac{\binom{n}{m+\alpha}}{\binom{n}{m}}\le n^\alpha$$

implies

$$B(m) \leq n^{\alpha} \binom{n}{m}^{2-\alpha} \binom{m+\alpha}{\alpha}^{\alpha}.$$

We rewrite the upper bound as the product of two terms

$$B(m) \leq n^{\alpha} \binom{n}{m}^{2-3\alpha/4} \cdot \left(\frac{\binom{m+\alpha}{\alpha}}{\binom{n}{m}^{1/4}} \right)^{\alpha}$$

and we turn to show that each of these two terms tend to 0. In the first one, the exponent is negative since $\alpha \geq 3$, and so it is maximized when m is minimized, namely $m = 10\alpha$. For this m it is bounded by

$$(ne)^{\alpha + 10\alpha(2-3\alpha/4)} \to 0$$

as required. The second expression is certainly negligible in the range $3n/4 > m \geq n/10$ where the numerator is polynomial in n whereas the denominator is exponential in n. For $m \leq n/10$ one easily sees that the maximum is obtained again when $m = 10\alpha$, where the numerator is bounded by a constant and the denominator grows like a polynomial. This proves our claim.

We now translate what all this means for the covering $\tilde{G}$ of G. Every set X of vertices in $\tilde{G}$ is the union of two disjoint sets X_a and X_b lying above a and b respectively, with $|X_a| = x_a$ and $|X_b| = x_b$. Assume w.l.o.g that $x_a \geq x_b$. If x_a and x_b differ by more than α, then ∂X automatically contains α vertices, as we have seen. This, and our assumptions on $|X|$, implies that $2 \leq x_a \leq 3n/4$.

We may think of X_a as a subset A of $\{1, \ldots, n\}$, since the vertices of $\tilde{G}_a$ are indexed by this set. By the construction of the covering, A_* corresponds to ∂X_a, as a subset of $\tilde{G}_b$. Thus we may assume that $|\partial X_a| \geq x_a + \alpha - \epsilon_1$. If $x_b \geq 2$ then similarly $|\partial X_b| \geq x_b + \alpha - \epsilon_2$. Since $\partial X = \partial X_a \cup \partial X_b \setminus (X_a \cup X_b)$, its size is at least

$$|\partial X| \geq x_a + x_b + 2\alpha - (\epsilon_1 + \epsilon_2) - (x_a + x_b) = 2\alpha - (\epsilon_1 + \epsilon_2) \geq \alpha$$

or, if $x_b \leq 1$, we just ignore ∂X_b and

$$|\partial X| \geq x_a + \alpha - \epsilon_1 - x_b \geq \alpha$$

as required. ∎

This result can be naturally extended to topological $K_2(\alpha)$s as well. Let G be a topological $K_2(\alpha)$, namely a graph with two vertices a, b and α vertex-disjoint paths between them.

Definition 2. An n-covering $\tilde{G}$ of $G = K_2(\alpha)$ is called *good* if every set $X \subset V(\tilde{G})$ with $3 \le x_a + x_b < 2 \cdot 2n/3$ has $|\partial X| \ge \alpha$.

Notice that the numbers x_v for $v \neq a,b$ are not important. The upper bound $2 \cdot 2n/3$ is actually the same as in Lemma 1, since $|\tilde{G}_a| + |\tilde{G}_b| = 2n$. An immediate corollary of Lemma 1 is

Corollary 1. *A random covering of a topological $K_2(\alpha)$ with $\alpha \ge 3$ is a.s. good.*

This follows at once since a covering of a topological $K_2(\alpha)$ is a subdivision of a covering of $K_2(\alpha)$, so there is a one-to-one correspondence between coverings of $K_2(\alpha)$ and coverings of a topological $K_2(\alpha)$ which preserves the probability distribution and the graphical properties. ∎

2.3. Thick Sets

Consider a finite base graph G and a subgraph H which is a topological $K_2(\alpha)$ for $\alpha \ge 3$, namely H consists of α internally disjoint paths between two vertices a and b of G. Given a covering $\tilde{G}$ of G, we can look at its restriction $\tilde{H}$ covering H, i.e., the subgraph of $\tilde{G}$ comprising of vertices and edges that are mapped to H by the covering projection. It is clear that if $\tilde{G}$ is a uniformly random covering, then $\tilde{H}$ is a uniformly random covering of H. For every H, this covering is a.s. good, and since there are only a finite number of possible such subgraphs H, a.s. *every* restricted covering $\tilde{H} \to H$ is good.

In this section we prove that for such coverings, thick sets (namely, sets that have more than one vertex in at least one fiber) have no less than δ neighbours. Throughout the proof, "disjoint paths" are vertex-disjoint, except possibly at their ends.

Proposition 1. *Let G be a finite graph with $\delta = \delta(G) \ge 3$, and let $\tilde{G} \to G$ be a covering such that the restriction $\tilde{H} \to H$ is good for every H that is a topological $K_2(\alpha)$ with $\alpha \ge 3$. Then for every thick set $X \subset V(\tilde{G})$, $|\partial X| \ge \delta$.*

We fix the base graph G and an n-covering $\tilde{G}$ satisfying the condition of the Proposition. X denotes a thick set of vertices in $\tilde{G}$ with $|X| < \tilde{G}/2$, and as usual, $X_v = X \cap \tilde{G}_v$, $x_v = |X_v|$ and $m = \max x_v$ so $m \ge 2$. Notice that (2) implies that $m < 2n/3$. We assume that X is such that $|\partial X| < \delta$, and arrive at a contradiction.

Notice that if $a,b \in G$ are such that $x_a = x_b = m$, and there are $\alpha \ge 3$ disjoint paths between a and b, then ∂X contains at least α vertices which

all lie above these paths (that is, they belong to the fibers of vertices in the paths). This is an immediate consequence of the definition of good covering and the bounds on m.

Let M' be the subgraph of G spanned by the vertices v for which $x_v = m$, and let M be any connected component of M'. For a vertex $a \in M$, set $d(a) = \deg(a; M)$ and $e(a) = \deg(a; G) - d(a)$, so $e(a)$ is the number of vertices outside M that are adjacent to a. Let $E(a)$ denote this set of vertices.

We will make frequent use of the following observation: If a vertex $v \in M$ is adjacent to a vertex $z \notin M$, then ∂X contains at least one vertex in $\tilde{G}_z$. (This is because the edges covering $[v, z]$ form a matching between $\tilde{G}_v$ and $\tilde{G}_z$, and $x_z < x_v$ so some vertex in X_v is matched to one outside X_z). It follows that every $a \in M$ implies the existence of at least $e(a)$ vertices in ∂X, which lie above $E(a)$. We also see that $|\partial M| < \delta$, and in particular M cannot be a single vertex.

Let $d = \delta(M)$ be the minimal degree of M, and $k = \bar{\kappa}(M)$ be its maximum local connectivity. We will gradually restrict the structure of M, until all possibilities are exhausted.

First, assume that $d \geq 3$. We use Mader's Theorem 2 to find vertices a, b in M with $d(a) = \alpha \leq \beta = d(b)$ having α disjoint paths between them. Since $2 \leq m \leq 2n/3$ and $\alpha \geq d \geq 3$, we find at least α vertices of ∂X lying above M. But $e(a) \geq \delta - \alpha$, and so ∂X contains also $\delta - \alpha$ vertices that do not lie over M, giving a total of δ vertices in ∂X, a contradiction. Therefore $d \leq 2$. See Figure 4.A.

Now suppose that $k \geq 3$. Let a be a vertex of M with $d(a) = \alpha \leq 2$, and let H be a topological $K_2(3)$ in M. $e(a)$ is at least $\delta - 2$, and there are at least 3 vertices in ∂X above H. This gives at least $\delta + 1$ vertices in ∂X, a contradiction. Note that it makes no difference if a belongs to H or not. This case is depicted in Figure 4.B.

Now we know that $k \leq 2$. Suppose that K is a block of M which is not an edge. Then K is 2-connected, and we claim that it is a cycle. Otherwise, let C be a cycle in K and $[c, x]$ an edge outside $E(C)$, with $c \in C$. If $x \in C$ then $\kappa(x, c; M) \geq 3$ which is impossible. Otherwise, since c is not a cutvertex of K, there is a path which avoids c from x to some vertex z on C. Thus $\kappa(c, z; M) \geq 3$ which is again impossible. We obtain that every block of M is either an edge or a cycle.

If M has at least two leaves, we proceed as follows. Let a and b be leaves. Since $e(a) \geq \delta - 1$, every vertex in $E(b) \setminus E(a)$ would give an additional vertex in ∂X which cannot happen. Therefore, $E(b) = E(a)$, and in particular there are at least δ disjoint a-b paths in G (one in M and $\delta - 1$ paths of length

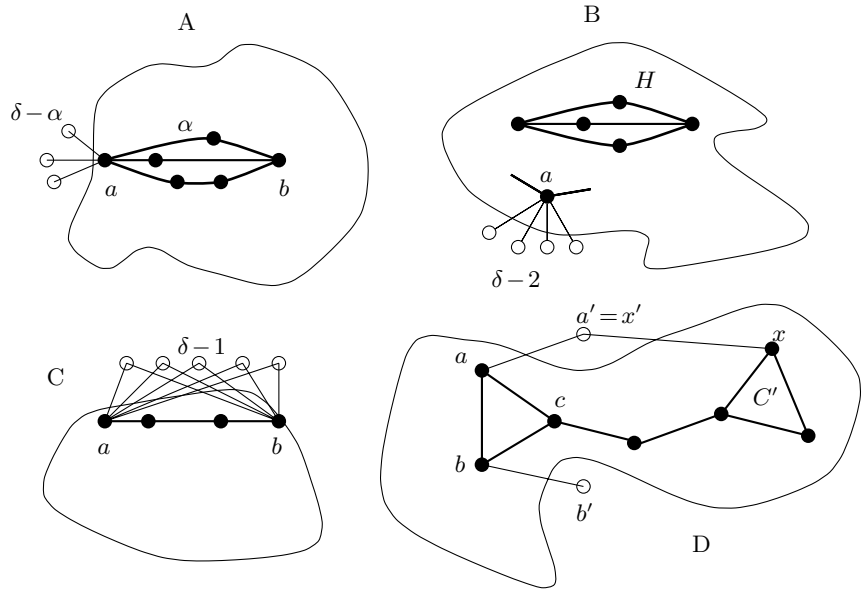

Fig. 4. Various cases for M.

2; See Figure 4.C). But this is also impossible, since it implies δ vertices of ∂X above these paths.

Since M is not a single vertex and contains at most one leaf, it has an endblock C which is a cycle. Let a and b be two vertices in C with $d(a)=2$. We claim that $E(a)$ is disjoint from $E(b)$. Otherwise, there are 3 disjoint a-b paths in G, two of which are in M and another one outside it. $E(a)$ still contains at least $\delta-3$ vertices disjoint from these paths, which together yield δ elements of ∂X, as usual.

It follows that C cannot contain 3 vertices a,b,c with $d(a)=d(b)=d(c)=2$, since then $E(a)$, $E(b)$ and $E(c)$ are three pairwise disjoint sets of size $\geq\delta-2$ each, and $3(\delta-2)\geq\delta$. In particular M is not a single cycle and must contain more than one block. Since C is an endblock, it must contain at least two vertices of degree 2, so the only remaining possibility is the one that appears in Figure 4.D: $C=\{a,b,c\}$ is a cycle of length 3 and one of its vertices, say c, belongs to another block. $E(a)$ and $E(b)$ together contain $2(\delta-2)$ vertices, so $\delta=3$ and $e(a)=e(b)=1$. Set $E(a)=\{a'\}$ and $E(b)=\{b'\}$.

Let C' be another endblock of M. There is a path from c to C' that avoids C. C' can be an edge or a cycle, but in either case it contains a vertex x with $d(x)\leq 2$, so $e(x)\geq 1$. Let $x'\in E(x)$ be some vertex. Then $x'=a'$ or $x'=b'$, since otherwise we get 3 distinct vertices in ∂M. Say $x'=a'$. There are now

3 disjoint paths from a to c, two in M (through C) and one outside, through a' and x. Once more this implies that $|\partial X| \geq 3$, which contradicts $\delta = 3$.

This final contradiction establishes the Proposition. ∎

2.4. Thin Sets

In this section we show that a.s. every *thin* set in $\tilde{G}$ has more than δ neighbours. Unlike our analysis of thick sets, here we work with certain subgraphs of $\tilde{G}$, the covering graph.

Definition 3. A subgraph H of a covering $\tilde{G}$ is called *edge-thin* if it does not contain a pair of *parallel edges*, namely edges covering the same edge of the base graph G.

An edge-thin subgraph contains at most one edge in each edge fiber, in contrast with a thin set which contains at most one vertex in each vertex fiber. Notice that a thin set spans an edge-thin subgraph, but the vertex set of an edge-thin subgraph is not necessarily thin. Also, a pair of parallel edges cannot be incident with the same vertex, by the definition of a covering.

Lemma 2. *Let G be a finite base graph and $\tilde{G}$ a random n-covering. Then a.s. in every edge-thin subgraph of $\tilde{G}$, every connected component is a tree or is unicyclic.*

Proof. Let $\Delta = \Delta(G)$ be the maximum degree of G. An edge-thin subgraph H of $\tilde{G}$ contains at most Δ vertices in each fiber, since if $\tilde{G}_v$ has $\Delta+1$ vertices in H, there is an edge in G incident with v that is covered twice by H. In particular, the number of vertices of an edge-thin subgraph is bounded by a constant $K = \Delta|V(G)|$.

It suffices to show that for each $1 \leq k \leq K$, a.s. no edge-thin subgraph H with k vertices has more than k edges. The number of possible vertex-sets $V(H)$ is $O(\binom{n|V(G)|}{k}) = O(n^k)$.

Given a vertex set V, $|V| = k$, we bound the probability that there is an edge-thin subgraph H with $V = V(H)$ and $|E(H)| = j > k$ as follows. First we choose a collection of j putative edges $[u,v]$, $u,v \in V$, no two of which cover the same edge of G. For each pair, the probability that $[u,v]$ is indeed an edge of $\tilde{G}$ is $1/n$, and these events are independent by the assumption that no edge of G is covered twice. Therefore the probability that all putative edges exist is n^{-j}, and the probability that *some* choice of putative edges succeeds is still $O(n^{-j})$ because the possible number of choices is bounded by, say, 2^{K^2}.

A union bound over all possible Vs gives $n^k n^{-j}$ which is $O(1/n)$ since $j > k$. Therefore, a.s. there are no edge-thin subgraphs with k vertices, as required. ∎

We now prove

Proposition 2. *Let G b a finite graph with $\delta = \delta(G) \geq 3$, and let $\tilde{G} \to G$ be a covering such that every edge-thin subgraph H of $\tilde{G}$ satisfies $|E(H)| \leq |V(H)|$. Then for every thin set $X \subset V(\tilde{G})$, $|\partial X| \geq \delta$.*

Fix a thin set X, $|X| = k$, and let H be the subgraph of $\tilde{G}$ spanned by X. Since H is edge-thin, it has at most k edges. It follows that H has a connected component which is a tree or a tree with an extra edge. Let H' be that connected component. We proceed by considering various possibilities for H'. See Figure 5.

Case 1. Suppose that H' has two leaves, say x,y. Namely, $\deg(x;H) = \deg(y;H) = 1$. Since $\deg(x;\tilde{G}) \geq \delta$, x is adjacent to at least $\delta - 1$ vertices $z_1, \ldots, z_{\delta-1}$ outside H, namely outside X. If y has a neighbour outside H which is different from the z_i's, we have δ vertices in ∂X and we are done. If not, then each z_i is adjacent to both x and y.

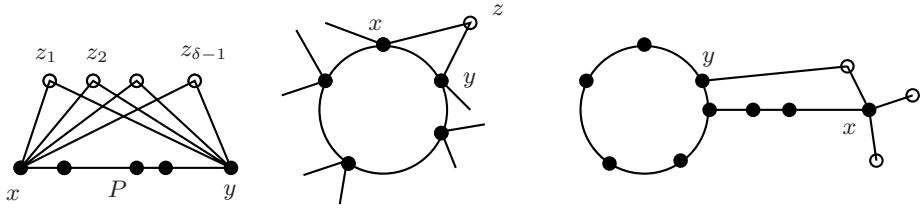

Case 1 Case 2 Case 3

Fig. 5. Cases for H'

We construct a certain subgraph Z of $\tilde{G}$: Take a simple x-y path P in H', along with the vertices z_i and all the edges $[x, z_i]$ and $[z_i, y]$. Clearly Z is edge-thin. Since $\delta \geq 3$, Z is a connected edge-thin subgraph with two distinct cycles, contrary to our assumption. Henceforth, H' has at most a single leaf.

Case 2. H' is a cycle. Since $\tilde{G}$ is simple, H' contains at least 3 vertices, each of which has at least $\delta - 2$ neighbours outside H'. Suppose that two vertices x, y of H' share a common neighbour z outside H'. Then the graph consisting of H', z and the edges $[x, z]$ and $[y, z]$ is edge-thin and has more edges than vertices, which is impossible by assumption. Therefore every vertex outside H' has at most a single neighbour in H'. Consequently $|\partial X| \geq 3(\delta - 2) \geq \delta$.

Case 3. The only remaining possibility is that H' is a lollypop: a cycle C with an attached path P. Let x be the endpoint of the path; it has $\delta-1$ neighbours outside H'. Let y be a vertex of degree 2 on C. This y has at least one neighbour outside H'. If it is not one of x's neighbours, then we are done. Otherwise the graph formed by adding this common neighbour to H' along with the two edges to x and y is edge-thin and has more edges than vertices, which is impossible. ∎

2.5. Conclusion

The proof of Theorem 1 is now complete. A random covering $\tilde{G}$ almost surely satisfies the conditions of both Proposition 1 and Proposition 2, as shown in the beginning of Section 2.3 and Lemma 2. Therefore both thick and thin sets have δ neighbours outside them, so $\tilde{G}$ is δ-connected. ∎

Let us make some further comments concerning the proof. It is not hard to extend Lemma 2 as follows: Fix a constant B. Then a.s. every subgraph H of $\tilde{G}$ with at most B parallel edges in each edge fiber satisfies $|E(H)| \leq |V(H)|$. It follows that topological $K_2(\delta)$s in $\tilde{G}$, which exist in abundance, must contain an unbounded amount of parallel edges, and in particular they are quite large. In a typical random covering, every two vertices are connected by δ disjoint paths, but these paths are necessarily quite long and utilize the lifts of some particular edges many times.

Theorem 1 is an asymptotic result. It will be interesting to estimate the probability that a random n-covering fails to be δ-connected in terms of n. We expect this probability to become small already for moderate values of n.

Following the proof of Theorem 1, we find that for most sets X, $|\partial X|$ is much larger than δ. However it is not true that a.s. only singletons of $\tilde{G}$ can be disconnected by removing δ vertices. For example, a random covering of K_4 is likely to contain a 3-cycle with only 3 neighbours. Yet the proof suggests that random coverings are good expanders, as large sets tend to have large boundaries. We will investigate expansion properties of random coverings in a forthcoming paper [1].

3. Variations

In this paper we considered the model of independent uniform permutations. Let us mention here some possible generalizations and variants of the model.

Given a group Γ acting on a set Ω, we can attach an element $g_e \in \Gamma$ to each edge e of G and view each fiber $\tilde{G}_u$ of as a copy of Ω, $\tilde{G}_u = \{u_\alpha\}_{\alpha\in\Omega}$.

The edges in $\tilde{G}$ now connect u_α to $v_{g_e\alpha}$, and $\tilde{G}$ is a covering of degree $|\Omega|$. If it is possible to select elements of Γ at random, we obtain random coverings. The standard model is the special case where $\Gamma = S_n$ and $\Omega = [n]$ with the usual action. Γ and Ω may be finite or infinite. (In fact, even G may be infinite as long as $\Gamma^{E(G)}$ is equipped with a probability measure).

A special case of this is that of the group Γ acting on itself by left multiplication. The resulting coverings of G will then be regular in the usual topological sense. If G is the bouquet B_k with one vertex and k loops, the covering will be the Cayley graph of Γ with respect to the set S of the elements attached to the loops and their inverses, assuming this set generates Γ.

A variation in a different direction is obtained as follows. Let $n = n_1 \cdot n_2 \cdots n_r$. Starting from G, we form a random n_1-covering, then a random n_2-covering *of the result*, and so on. Since a composition of covering maps is itself a covering, the resulting graph is an n-degree cover of G, distributed differently than one formed by taking a random n-covering directly. This kind of model may be more suited for some applications.

A far-reaching generalization of the above would be a model for random coverings of higher-dimensional simplicial complexes. At present this seems difficult to obtain even for 2-dimensional complexes. However, we mention here an interesting model for random branched coverings of a 2-manifold.

Let M be a 2-dimensional manifold and let G be a graph embedded in M in such a way that the regions (connected components of $M \setminus G$) are homeomorphic to open disks, and their boundaries form simple cycles on the graph (this is sometimes called a *strong embedding*). For example, we can let G be the 1-skeleton of a triangulation of M.

Consider a covering $\tilde{G}$ of G. Every cycle C in G that bounds a region in M is covered by the disjoint union of cycles $\tilde{C}_j$ in $\tilde{G}$, and every edge of $\tilde{G}$ belongs to exactly two such cycles. Let us attach a 2-cell to each $\tilde{C}_j$. We get a 2-complex $\tilde{M}$ and it is easy to verify that it is, in fact, a manifold. There is also a natural map $\tilde{M} \to M$ which maps the cell attached to C_j onto the region bounded by C, and restricts to $\tilde{G} \to G$ on the 1-skeleton. This map will not in general be a covering of manifolds: it may be branched at the centers of the attached cells.

If the original covering $\tilde{G} \to G$ is chosen randomly, this procedure yields a random branched covering $\tilde{M}$ of M, although we do not get all coverings of M that way. It will be interesting to see if anything useful can be done with this model.

References

[1] Alon Amit and Nathan Linial: Random lifts of graphs II: Edge-expansion, to appear.

[2] Dana Angluin: Local and global properties in networks of processors, In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, Los Angeles, California, 1980.

[3] Dan Archdeacon and R. Bruce Richter: On the parity of planar covers, *Journal of Graph Theory*, **14**(2) (1990).

[4] Hyman Bass and Ravi Kulkarni: Uniform tree lattices, *Journal of the Americn Mathematical Society*, **3**(4) (1990), 843–902.

[5] Béla Bollobás: *Random Graphs*, Academic Press, 1985.

[6] J. L Gross and T. W. Tucker: Generating all graph coverings by permutation voltage assignments, *Discrete Mathematics*, **18** (1977), 273–283.

[7] J. L. Gross and T. W. Tucker: *Topological Graph Theory*, Wiley-Interscience, 1987.

[8] Frank Thomson Leighton: Finite common coverings of graphs, *Journal of Combinatorial Theory Ser. B*, **33** (1982), 231–238.

[9] Wolfgang Mader: Grad und lokaler zusammenhang in endlichen graphen, *Math. Ann.*, **205** (1973), 9–11.

[10] Seiya Negami: The spherical genus and virtually planar graphs, *Discrete Mathematics*, **70** (1988), 159–168.

[11] John Stallings: Topology of finite graphs, *Inventiones Mathematicae*, **71** (1983), 551–565.

[12] John Stillwell: *Classical Topology and Combinatorial Group Theory*, Springer-Verlag, 1993.

Alon Amit

Institute of Mathematics,
Hebrew University,
Jerusalem 91904, Israel.
alona@math.huji.ac.il

Nathan Linial

Institute of Computer Science,
Hebrew University,
Jerusalem 91904, Israel.
nati@cs.huji.ac.il