Epilogue: Sparsity in the infinite

Compilation date: January 22, 2018

In this last lecture we will explore basics of the sparsity theory in infinite graphs. We will also show how to obtain results for finite graphs from results for infinite graphs via a set-theoretic construction of *ultraproducts*, which enables easy translation between finite and infinite.

1 Ultrafilter and ultraproducts

Ultrafilters. Throghout this lecture we denote $\omega = \mathbb{N} = \{0, 1, 2, \ldots\}$.

Definition 1. An *ultrafilter* on ω is any family $\mathcal{U}$ of subsets of ω satisfying the following conditions:

(a) $\emptyset \notin \mathcal{U}$;

(b) if $X \subseteq Y \subseteq \omega$ and $X \in \mathcal{U}$, then also $Y \in \mathcal{U}$;

(c) if $X, Y \in \mathcal{U}$ then $X \cap Y \in \mathcal{U}$; and

(d) for each $X \subseteq \omega$, either X or $\omega - X$ belongs to $\mathcal{U}$.

One may view an ultrafilter $\mathcal{U}$ on ω also as a $\{0, 1\}$-measure on ω. Each subset of ω is considered either *large* (belongs to $\mathcal{U}$) or *small* (does not belong to $\mathcal{U}$). Then the conditions are natural: empty set is small, a superset of a large set is also large, the intersection of two large sets is large as well, and for each set, either this set or its complement is large. Observe that converse conditions for small sets hold: a subset of a small set is small and the union of two small sets is small.

It is very easy to give a trivial example of an ultrafilter $\mathcal{U}$ on ω: take any number $i \in \omega$ and declare a set $X \subseteq \omega$ large if $i \in X$. Such ultrafilters are called *principal* and will not be interesting for us. Conversely, an ultrafilter is *non-principal* if every singleton set does not belong to it (is small). Note that this entails that every finite set is small.

Lemma 1. *There exists a non-principal ultrafilter on ω.*

Proof. Consider the set of all *non-principal filters* on ω, that is, families of subsets of ω which satisfy conditions (a)–(c) and contain all co-finite sets, i.e., sets that lack only finitely many elements from ω. Note that this set is nonempty, because just the family of co-finite sets is a filter. Also, observe that by conditions (a) and (c), a non-principal filter does not contain any finite set.

Order non-principal filters by inclusion and observe that for any inclusion chain of non-principal filters, their union is also a non-principal filter. By Zorn's lemma, among non-principal filters there exists an inclusion-wise maximal one, say $\mathcal{U}$.

We now prove that $\mathcal{U}$ is an ultrafilter. Take any $X \subseteq \omega$ and for the sake of contradiction suppose neither X nor $\omega - X$ belongs to $\mathcal{U}$. At most one of X and $\omega - X$ is finite; suppose w.l.o.g. that X is infinite. Consider a family $\mathcal{W}$ of subsets of ω obtained by taking all sets of the form $X \cap Y$ for $Y \in \mathcal{U}$ and their supersets. Clearly $\mathcal{U} \subseteq \mathcal{W}$ and $X \in \mathcal{W}$, so in fact $\mathcal{U} \subsetneq \mathcal{W}$. We check that $\mathcal{W}$ is a non-principal filter, which contradicts the maximality of $\mathcal{U}$. Closure under taking supersets and under taking intersections (conditions (b) and (c)) is obvious. We are left with condition (a), so suppose $\mathcal{U}$ contains some set Y that is disjoint with X. Then $\omega - X$ is a superset of Y, so by (b) we have $\omega - X \in \mathcal{U}$, a contradiction. $\square$

From now on we fix some non-principal ultrafilter $\mathcal{U}$ on ω and whenever we say *large* or *small*, we measure it with respect to $\mathcal{U}$.

Ultraproducts. Now, using the fixed ultrafilter $\mathcal{U}$ we define limits of sequences of graphs. Suppose $G_0, G_1, G_2, \ldots$ is a sequence of graphs, and for convenience denote $V_n = V(G_n)$. Consider the set $\prod_{n\in\omega} V_n$, that is, the set of all sequences of vertices where the nth vertex is taken from the nth graph. Consider the following relation $\sim$ on $\prod_{n\in\omega} V_n$:

$$(u_0, u_1, u_2, \ldots) \sim (v_0, v_1, v_2, \ldots) \qquad \text{if and only if} \qquad \{n\colon u_n = v_n\} \in \mathcal{U}.$$

In other words, two sequences of vertices are $\sim$-equivalent if and only if the set of indices on which they differ is small with respect to $\mathcal{U}$. Note that $\sim$ is an equivalence relation: if $(u_0, u_1, u_2, \ldots)$ differs from $(v_0, v_1, v_2, \ldots)$ on a small set of indices, and $(v_0, v_1, v_2, \ldots)$ differs from $(w_0, w_1, w_2, \ldots)$ on a small set of indices, then $(u_0, u_1, u_2, \ldots)$ differs from $(w_0, w_1, w_2, \ldots)$ on a small set of indices, because the union of two small sets is small. Define

$$\prod_{\mathcal{U}} V_n \coloneqq \prod_{n\in\omega} V_n \Big/ \sim .$$

In other words, $\prod_{\mathcal{U}} V_n$ comprises equivalence classes of $\sim$. We think of elements of $\prod_{\mathcal{U}} V_n$ as of sequences of vertices from consecutive graphs G_n, defined up to perturbation on a small set. For $\mathbf{u} \in \prod_{n\in\omega} V_n$, let $[\mathbf{u}] \in \prod_{\mathcal{U}} V_n$ be its equivalence class in $\sim$.

The set $\prod_{\mathcal{U}} V_n$ will be the vertex set of the *ultraproduct* $\prod_{\mathcal{U}} G_n$; it remains to define the adjacency relation. Take any two sequences $\mathbf{u} = (u_0, u_1, u_2, \ldots)$ and $\mathbf{v} = (v_0, v_1, v_2, \ldots)$, representing two vertices $[\mathbf{u}], [\mathbf{v}] \in \prod_{\mathcal{U}} V_n$. We make vertices $[\mathbf{u}]$ and $[\mathbf{v}]$ adjacent in $\prod_{\mathcal{U}} G_n$ if and only if the following condition holds:

$$\{n\colon u_n v_n \in E(G_n)\} \in \mathcal{U}.$$

In other words, the edge between u_n and v_n should be present in G_n for a large set of indices n. Note that this definition does not depend on the choice of representatives $\mathbf{u}$ and $\mathbf{v}$, as representatives differ from each other only on small sets.

Łoś' theorem. We now give a basic theorem for translation of properties from a sequence of graphs to their ultraproduct.

Theorem 2 (Łoś' Theorem)**.** *Let $\varphi(x^1, \ldots, x^k)$ be a first-order formula over graphs and $G_0, G_1, G_2, \ldots$ be graphs. Suppose $\mathbf{u}^1, \ldots, \mathbf{u}^k \in \prod_{n\in\omega} V_n$. Then the following conditions are equivalent:*

- *$\varphi([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ holds in $\prod_{\mathcal{U}} G_n$; and*
- *the set of indices n for which $\varphi(u^1_n, \ldots, u^k_n)$ holds in G_n is large with respect to $\mathcal{U}$.*

Proof. We proceed by induction on the structure of the formula. When φ is an atom $E(x, y)$, where E is the binary predicate denoting adjacency, then the claim follows immediately from the definition of the ultraproduct.

Suppose $\varphi(x^1, \ldots, x^k) = \psi_1(x^1, \ldots, x^k) \wedge \psi_2(x^1, \ldots, x^k)$. For $t = 1, 2$, let I_t be the set of those indices n for which $\psi_t(u^1_n, \ldots, u^k_n)$ holds in G_n. By induction hypothesis, $\psi_t([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ holds in $\prod_{\mathcal{U}} G_n$ if and only if I_t is large. Observe that the set of indices i for which $\varphi(u^1_n, \ldots, u^k_n)$ holds in G_n is exactly $I_1 \cap I_2$. Now if $\varphi([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ holds, then I_1, I_2 are both large and consequently $I_1 \cap I_2$ is large as well. However, if $\varphi([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ does not hold, then either I_1 or I_2 is small, and $I_1 \cap I_2$ is small as well. The case when $\varphi(x^1, \ldots, x^k) = \neg\psi(x^1, \ldots, x^k)$ is analogous.

Finally, suppose $\varphi(x^1, \ldots, x^k) = \exists_y \, \psi(x^1, \ldots, x^k, y)$. If $\varphi([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ holds, then there is $[\mathbf{w}] \in \prod_{\mathcal{U}} V_n$ such that $\psi([\mathbf{u}^1], \ldots, [\mathbf{u}^k], [\mathbf{w}])$ holds. By induction hypothesis, $\psi(u^1_n, \ldots, u^k_n, w_n)$ holds for a large set of indices n, implying that so does $\varphi(u^1_n, \ldots, u^k_n)$ — vertices w_n serve as witnesses for existential quantification. Conversely, if $\varphi(u^1_n, \ldots, u^k_n)$ holds for a large set of indices I, then for indices $n \in I$ we may find $w_n \in V_n$ such that $\psi(u^1_n, \ldots, u^k_n, w_n)$. Setting $\mathbf{w} = (w_0, w_1, w_2, \ldots)$, where w_n is set in any way for $n \notin I$, by induction hypothesis we infer that $\psi([\mathbf{u}^1], \ldots, [\mathbf{u}^k], [\mathbf{w}])$ holds in $\prod_{\mathcal{U}} G_n$, and hence $\varphi([\mathbf{u}^1], \ldots, [\mathbf{u}^k])$ holds as well. □

2 Sparsity for infinite graphs

Observe that notions of depth-r minors and depth-r topological minors work equally well for (possibly) infinite graphs. Recall that for a class $\mathcal{C}$ and $r \in \omega$, by $\mathcal{C} \nabla r$ we denote the set of all depth-r minors of graphs from $\mathcal{C}$. Let us define

$$\mathcal{C} \nabla \omega \coloneqq \bigcup_{r \in \omega} \mathcal{C} \nabla r.$$

In other words, $\mathcal{C} \nabla \omega$ comprises graphs that can be observed as depth-r minors of graphs from $\mathcal{C}$ for some fixed finite depth r. Replacing depth-r minors with depth-r topological minors, we define $\mathcal{C} \nabla^{\mathrm{t}} r$ and $\mathcal{C} \nabla^{\mathrm{t}} \omega$ in the same way.

Note that if $\mathcal{C}$ consists only of finite graphs, then $\mathcal{C} \nabla \omega$ is simply the closure of $\mathcal{C}$ under taking minors. However, when $\mathcal{C}$ contains infinite graphs then this is not necessarily the case. This is because it may be that some infinite graph H is a minor of a graph $G \in \mathcal{C}$, but all minor models do not have uniformly bounded radius; in this case H is *not* included in $\mathcal{C} \nabla \omega$.

Let K_ω be the complete graph on ω vertices.

Definition 2. A class of (possibly infinite) graphs $\mathcal{C}$ is *limit nowhere dense* if $K_\omega \notin \mathcal{C} \nabla \omega$. Similarly, $\mathcal{C}$ is *limit topologically nowhere dense* if $K_\omega \notin \mathcal{C} \nabla^{\mathrm{t}} \omega$.

Note that in the above definition there are no annoying radii, parameters, or functions. This is the main advantage of going to the infinite setting.

We would now like to connect nowhere denseness for classes of finite graphs and limit nowhere denseness for infinite graphs. To this end we use ultraproducts.

Definition 3. Let $\mathcal{C}$ be a class of finite graphs. Then define the *limit class* $\mathcal{C}^\star$ as follows:

$$\mathcal{C}^\star \coloneqq \left\{ \prod_{\mathcal{U}} G_n \ : \ G_0, G_1, G_2, \ldots \in \mathcal{C} \right\}.$$

In other words, $\mathcal{C}^\star$ comprises all ultraproducts of sequences of graphs from $\mathcal{C}$.

Lemma 3. *A class of finite graphs $\mathcal{C}$ is nowhere dense if and only if $\mathcal{C}^\star$ is limit nowhere dense.*

Proof. Suppose first $\mathcal{C}$ is nowhere dense. Observe that for all $t, r \in \omega$ there is a first-order sentence $\varphi_{t,r}$ that, when applied to a graph G, verifies whether K_t is a depth-r minor of G. Indeed, a depth-r

minor model of K_t can be assumed to use at most rt^2 vertices, hence the question whether K_t is a depth-r minor of G boils down to asking whether G admits one of a finite list of subgraphs on at most rt^2 vertices; then $\varphi_{t,r}$ makes a disjunction over these subgraphs. From Łoś' Theorem we infer that if $K_t \notin \mathcal{C}\nabla r$, then also $K_t \notin \mathcal{C}^\star\nabla r$. This implies that if $\mathcal{C}$ is nowhere dense, then for every $r \in \omega$ there is a finite bound on the sizes of cliques that can be found as depth-r minors of graphs from $\mathcal{C}^\star$. In particular, $K_\omega \notin \mathcal{C}^\star\nabla\omega$.

Suppose now $\mathcal{C}$ is somewhere dense. This means that there is $r \in \omega$ and a sequence of graphs $G_0, G_1, G_2, \ldots$ from $\mathcal{C}$ such that K_n is a depth-r minor of G_n, for all $n \in \omega$. For each $n \in \omega$, fix a depth-r minor model ϕ_n of K_n in G_n. We may assume that each branch set $\phi_n(i)$ is a tree of depth at most r rooted at a vertex c_n^i, formed by the union of $n-1$ paths connecting c_n^i with connection points to other branch sets. Let

$$c_n^i = u_n^{i,j,0}, u_n^{i,j,1}, \ldots, u_n^{i,j,r}$$

be the vertices on the path within $\phi_n(i)$ connecting c_n^i with the connection point to $\phi_n(j)$, where we put $u_n^{i,j,t} = u_n^{i,j,t+1}$ for some t in case this path is shorter than r.

Now, for each $i, j \in \omega$ with $i \neq j$ and $t \in \{0, 1, \ldots, r\}$ consider the sequence

$$\mathbf{u}^{i,j,t} := (u_n^{i,j,t})_{n\in\omega}.$$

where we put arbitrary vertices as $u_n^{i,j,t}$ when $n < i$ or $n < j$. Similarly define the sequence

$$\mathbf{c}^i := (c_n^i)_{n\in\omega}.$$

Observe that the following assertions hold whenever $i, j, i', j' \leqslant n$:

- $u_n^{i,j,0} = c_n^i$;
- $u_n^{i,j,t}$ is either equal or adjacent to $u_n^{i,j,t+1}$, for all $t \in \{0, 1 \ldots, r-1\}$;
- $u_n^{i,j,r}$ is adjacent to $u_n^{j,i,r}$; and
- $u_n^{i,j,t}$ is different from $u_n^{i',j',t'}$ whenever $i \neq i'$.

All the above properties are expressible in first-order logic, and for fixed i, j, i', j' they do not hold only for a finite number of indices n — which is a small set w.r.t. $\mathcal{U}$. Hence, by Łoś' Theorem we infer that they also hold for vertices $\mathbf{u}^{i,j,t}$ and $\mathbf{c}^i$ in $\prod_\mathcal{U} G_n$. It now immediately follows that vertices $\mathbf{u}^{i,j,t}$ and $\mathbf{c}^i$ for $i, j \in \omega$, $i \neq j$, form a depth-r minor model of K_ω in $\prod_\mathcal{U} G_n$. □

Recall that a class $\mathcal{C}$ of finite graphs is *topologically nowhere dense* if for every $r \in \omega$ there exists a number $t(r)$ such that no graph from $\mathcal{C}$ admits $K_{t(r)}$ as a depth-r topological minor. A very similar reasoning yields the following.

Lemma 4. *A class of finite graphs $\mathcal{C}$ is topologically nowhere dense if and only if $\mathcal{C}^\star$ is limit topologically nowhere dense.*

In the first lectures we have proved that nowhere denseness is equivalent to topological nowhere denseness. We will now reprove this fact by going to the infinite setting By Lemmas 3 and 4, it suffices to prove the following statement; note that it does not involve any ultraproducts, it is a statement purely about classes of infinite graphs.

Lemma 5. *A graph class $\mathcal{C}$ is limit nowhere dense if and only if it is limit topologically nowhere dense.*

Proof. The left-to-right implication is obvious: a depth-r topological minor is also a depth-r minor.

For the second implication, suppose $\mathcal{C}$ is not limit nowhere dense. This means that for some $r \in \omega$ there is a depth-r minor model ϕ of K_ω in some graph $G \in \mathcal{C}$. Thus, for each $i \in \omega$, the branch set $\phi(i)$ is a subgraph of G of radius at most r, and subgraphs $\phi(i)$ are pairwise disjoint and pairwise adjacent. For each $j \neq i$ there is some vertex $u_{j,i}$ in $\phi(j)$ that is adjacent to a vertex $u_{i,j}$ in $\phi(i)$. Let $\overline{\phi}(i)$ be $\phi(i)$ augmented by adding $u_{j,i}$ together with the edge $u_{j,i}u_{i,j}$, for all $j \neq i$. By removing unnecessary vertices and edges, we may assume that each $\overline{\phi}(i)$ is a tree of depth at most $r+1$ rooted at a vertex c_i, whose leaves are exactly vertices $u_{j,i}$ for $j \neq i$.

Since $\overline{\phi}(i)$ is a rooted tree of finite depth and infinitely many leaves, we may find a vertex ℓ_i in $\overline{\phi}(i)$ that has infinitely many children. For each child w of ℓ_i in $\overline{\phi}(i)$, arbitrarily select $j \in \omega$ such that the leaf $u_{j,i}$ is a descendant of w. Call all indices j defined in this way *buddies* of i. Thus, every index i has infinitely many buddies $j \neq i$.

We now define a depth-$(3r+1)$ topological minor of K_ω in G by induction. Start with the empty model, and during the induction maintain the following invariant: after the nth step we have defined a depth-$(3r+1)$ topological minor model ψ_n of K_n whose every principal vertex is equal to ℓ_i for some $i \in \omega$. In the inductive step we extend ψ_{n-1} to ψ_n as follows. Note that ψ_{n-1} involves only a finite number of vertices, so there exist some $m \in \omega$ such that $\phi(m)$ is disjoint with the model ψ_{n-1}. We start defining ψ_n from ψ_{n-1} by mapping n to ℓ_m.

Now, for consecutive $j = 1, \ldots, n$ we show how to find a suitable path for the image of the edge jn. Let k be such that $\psi(j) = \ell_k$. Observe that m has infinitely many buddies and j has infinitely many buddies, and only finitely many branch sets of ϕ intersect the model ψ_n so far. Hence we may find a buddy x of j and a buddy y of n such that $\phi(x)$ and $\phi(y)$ do not intersect the model ψ_n. Define path P_{jn} as follows: starting from ℓ_m go to $u_{m,x}$ within $\phi(m)$, then proceed to $u_{x,m}$, then travel within $\phi(x)$ to $u_{x,y}$, then proceed to $u_{y,x}$, then travel within $\phi(y)$ to $u_{y,k}$, then proceed to $u_{k,y}$, and finally go from $u_{k,y}$ to ℓ_k within $\phi(k)$. This path has length at most $6r+3$, which is allowed in a depth-$(3r+1)$ topological minor model. Since we picked buddies whose branch sets were disjoint with the current model, it is easy to see that the path P_{jn} is internally disjoint with all the other paths present in the current model. Hence we may set $\psi(jn)$ to be P_{jn} and proceed to the next index j.

Thus, we have inductively defined a depth-$(3r+1)$ topological minor model of K_ω in G, which means that $\mathcal{C}$ is not limit topologically nowhere dense. $\square$

As we argued, the combination of Lemmas 3, 4, and 5 gives an indirect way of obtaining equivalence of nowhere denseness and topological nowhere denseness for classes of finite graphs. Taking apart the quite straightforward translation using Łoś' theorem, the proof relies on essentially the same combinatorial idea but is conceptually somewhat easier. Namely, instead of using multiple numerical parameters that need to be weighted against each other in the proof, we replace every intuitive occurrence of "large" with formal "infinite" and every intuitive occurrence of "small" with "finite". The same methodology can be used to rework the proofs of other statements concerning nowhere dense classes, for example the equivalence of nowhere denseness and uniform quasi-wideness, or the equivalence of nowhere denseness and the existence of a winning strategy for Splitter in the Splitter game.