REU 2007 · Transfinite Combinatorics · Lecture 3

Instructor: László Babai
Scribe: Damir Dzhafarov

July 23, 2007. Revised by instructor.
Last updated July 29, 10:40 PM

3.1 Directed Graphs

In a previous lecture, we assigned the non-trivial following exercise:

Exercise 3.1.1 (Erdős).** Prove that for all $g, k \in \mathbb{N}$, there exists a finite graph of girth $\geq g$ and chromatic number $\geq k$.

Here, we recall that the *girth* of a graph is the length of its shortest cycle, while the *chromatic number* refers to the least number of colors sufficient to color the vertices so that adjacent vertices never receive the same color. In what follows we present a solution to the problem for $g = 4$.

Definition 3.1.2. A *directed graph* G is a pair (V, E) where V is a set, called the set of *vertices*, and E, the set of *edges* or *arrows*, is an arbitrary subset of $V \times V \setminus \{(v, v) : v \in V\})$, i. e., an antireflexive binary relation on V.

Definition 3.1.3. Given a directed graph $G = (V, E)$, the *directed line graph of* G, written $\vec{L}(G)$, is the directed graph

$$(E, \{(e, f) \in E \times E : (\exists x, y, z \in V)[e = (x, y) \wedge f = (y, z)]\}).$$

In words, $\vec{L}(G)$ has as vertices the edges of G, and has a directed edge from one edge of G to another just in case the terminus in G of the former is the initial point of the latter.

Example 3.1.4. Let G be a directed graph and suppose vertex v has indegree k and outdegree ℓ, i. e., k edges end at v and ℓ edges start at v. Each of these $k + \ell$ edges will represent a vertex in the directed line graph $\vec{L}(G)$, and since the terminus of each of the former k edges is the initial point of each of the latter ℓ, in the line graph each of the former will have an edge directing to each of the latter, so we shall have a (directed) complete bipartite graph $\vec{K}_{k,\ell}$ in $\vec{L}(G)$.

Definition 3.1.5. Let $G = (V, E)$ be a directed graph, and let e_1, e_2, e_3 be distinct elements of E. We call the triple $\{e_1, e_2, e_3\}$ a *directed* 3*-cycle* if for some $v_1, v_2, v_3 \in V$, $e_1 = (v_1, v_2)$, $e_2 = (v_2, v_3)$ and $e_3 = (v_3, v_1)$. We call $\{e_1, e_2, e_3\}$ a *transitive triple* if for some $v_1, v_2, v_3 \in V$, $e_1 = (v_1, v_2)$, $e_2 = (v_2, v_3)$, and $e_3 = (v_1, v_3)$. A *3-cycle* is a sequence of edges which form an undirected 3-cycle if we ignore orientation.

Exercise 3.1.6. Show that in a directed graph, any 3-cycle is either a directed 3-cycle, or else a transitive triple.

Suppose $\{e_1, e_2, e_3\}$ is a directed 3-cycle in the directed graph G. Then the terminal point of e_1 is the initial point of e_2, so (e_1, e_2) is an edge in $\vec{L}(G)$. It follows similarly that (e_2, e_3) and (e_3, e_1) are edges in $\vec{L}(G)$, and so $\{(e_1, e_2), (e_2, e_3), (e_3, e_1)\}$ is a directed 3-cycle in $\vec{L}(G)$. Conversely, any directed 3-cycle in $\vec{L}(G)$ must be of the form $\{(e_1, e_2), (e_2, e_3), (e_3, e_1)\}$, where the e_i are edges in G, say with initial point a_i and terminal point b_i. Since (e_1, e_2) is an edge in $\vec{L}(G)$, it follows that $b_1 = a_2$, and similarly it follows that $b_2 = a_3$ and $b_3 = a_1$. Thus, $\{e_1, e_2, e_3\}$ is a directed 3-cycle in G, so we have a one-to-one correspondence between the directed 3-cycles in G and those in $\vec{L}(G)$.

Exercise 3.1.7. Show that, in fact, *all* the 3-cycles in $\vec{L}(G)$ are directed (i.e., that in $\vec{L}(G)$ there are no transitive triples). Thus, show that there is a one-to-one correspondence between the 3-cycles in $\vec{L}(G)$ and the directed 3-cycles in G.

Corollary 3.1.8. *If G has no directed* 3*-cycles then $\vec{L}(G)$ has no* 3*-cycles.*

Thus, to solve our problem, we can try to construct a directed graph G with no directed 3-cycles such that $\vec{L}(G)$ has large chromatic number. The key ingredient to doing this will be the following lemma, which we leave as an exercise:

Exercise 3.1.9. Let G be a directed graph. Show that $\chi(G) \leq 2^{\chi(\vec{L}(G))}$, where χ denotes the chromatic number.

The above lemma gives an upper bound on the chromatic number of an arbitrary directed graph. However, we shall prefer to think of it as a lower bound on the chromatic number of a directed line graph. Suppose, for example, that we wish to construct a directed graph G such that $\chi(\vec{L}(G)) \geq 100$. Recalling that $\chi(K_n) = n$ for every n, we can simply take any orientation of $G = K_{2^{100}}$, obtaining that

$$2^{100} = \chi(G) \leq 2^{\chi(\vec{L}(G))},$$

and hence that $\chi(\vec{L}(G)) \geq 100$. But we are still left with the problem that the particular orientation we chose of $K_{2^{100}}$ may have lots of directed 3-cycles. The following theorem resolves this problem.

Observation 3.1.10. Any graph $G = (V, E)$ admits an orientiation which contains no directed 3-cycles.

Proof. Put a linear order on the set of vertices and orient every edge from smaller to bigger.

For infinite sets, this proof requires the existence of a linear order. In fact, by the Well Ordering Theorem, there exists a well-ordering, which is more than what we need. $\square$

Definition 3.1.11. Let $\vec{L}^0(G) = G$, and for $k \in \mathbb{N}$, let $\vec{L}^{k+1}(G) = \vec{L}(\vec{L}^k(G))$.

Exercise 3.1.12. Show that if G has the orientation derived from a linear order (as in Observation 3.1.10) then $\vec{L}^k(G)$ has no odd cycles of length $\leq 2k+1$.

We thus proved a major result:

Theorem 3.1.13. *Given any integer g and any (finite or infinite) cardinal number κ, there exists a graph of chromatic number $\geq \kappa$ which has no* odd *cycles shorter than g.*

3.2 Borsuk's Theorem

In this section, we present yet another method by which to construct a graph with large chromatic number which avoids short odd cycles. Recall that the *n-sphere* is the set $S^n = \{\mathbf{x} \in \mathbb{R}^{n+1} : \|\mathbf{x}\| = 1\}$. Fix n and $\varepsilon > 0$, and define two points $\mathbf{x}$ and $\mathbf{y}$ in S^n to be *adjacent* if $\mathrm{dist}(\mathbf{x}, \mathbf{y}) \geq 2 - \varepsilon$ (i.e., $\mathbf{x}$ and $\mathbf{y}$ are *near antipodes*). With this definition, S^n can be regarded as a graph (of size $\mathfrak{c}$), and colorings of vertices may be defined as usual. Note that for small enough ϵ, the chromatic number of this graph is the minimum number of closed sets of diameter less than 2 that cover S^n.

Let us consider the case $n = 1$.

Exercise 3.2.1. Show that $\chi(S^1) > 2$.

In fact, $\chi(S^1) = 3$. Indeed, inscribe an equilateral triangle T in S^1, each side of which is labeled by a different color, and color the point $\mathbf{x} \in S^1$ by the color of the side of T intersected by the line segment from $\mathbf{x}$ to $\mathbf{0} \in \mathbb{R}^2$. For $n = 2$, we can argue similarly to show that $\chi(S^2) \leq 4$. Inscribe in S^2 a tetrahedron, with each side labeled by a color, and define a coloring of S^2 analogously to the way we did above. The general case is described by the following:

Theorem 3.2.2 (Borsuk). *For $n \geq 1$, $\chi(S^n) = n + 2$.*

We thus know of another class of graphs which easily give us large chromatic numbers. Together with the following exercise, we get another solution to our opening problem:

Exercise 3.2.3. Show that by taking ε sufficiently small, S^n will contain no short odd cycles.

Borsuk's Theorem is equivalent to the following theorem:

Theorem 3.2.4 (Borsuk-Ulam). *Let $f : S^n \to \mathbb{R}^n$ be continuous. Then there exists an $\mathbf{x} \in S^n$ such that $f(\mathbf{x}) = f(-\mathbf{x})$.*

Exercise 3.2.5. Establish the equivalence of Theorems 3.2.2 and 3.2.4.

It may be tempting to think that Exercise 3.1.1 can be generalized to the transfinite, perhaps something to the effect that for every integer g and every infinite cardinal κ, there exists a graph with chromatic number $\geq \kappa$ but no cycle of length $< g$. But in fact, we have already had an exercise which witnesses that this generalization *is false*, failing for $\kappa = \aleph_1$ and $g = 5$:

Exercise 3.2.6 (Erdős-Hajnal).** If a graph has uncountable chromatic number, then it contains a 4-cycle.

3.3 More on Ordinals and Cardinals

Recall that a *well order* is an ordered set $(A, \leq)$ with no infinite $\leq$-descending sequence (a previously-assigned exercise asked to show that this was equivalent to every non-empty subset of A having a $\leq$-minimum). We defined an *ordinal* as the order type of a well ordered set, and we remarked that:

Theorem 3.3.1. *Given any two ordinals, one is necessarily order-isomorphic to a prefix of the other.*

Theorem 3.3.2. *Any set of ordinals is well-ordered under the relation of being isomorphic to a prefix.*

Proof. The prefix relation is obviously reflexive and transitive, and in view of Theorem 3.3.1, it is total.

Observation 3.3.3. No ordinal is isomorphic to a proper prefix of itself. Indeed, if $f : \alpha \to \alpha$ were a non-surjective monotone injection, we could find some $x \in \alpha - f(\alpha)$, so that $f(x) < x$ by definition of prefix. But since f is monotone, this would mean that $\cdots < fff(x) < ff(x) < f(x) < x$ is an infinite descending sequence in α, which is a contradiction.

Antisymmetry now follows by the Observation and transitivity. Finally, suppose S is a non-empty set of ordinals, and suppose f is a monotone injection $\omega^* \to S$. Since $f(n+1) < f(n)$ for all n, it follows that there exists a non-bijective monotone injection $g_n : f(n+1) \to f(n)$. Let x_0 be the least element of $f(0) \setminus g_0(f(1))$, and for $n \in \mathbb{N}$, let x_{n+1} be the least element of $f(0) - g_0(g_1(\cdots g_{n+1}(f(n+1))\cdots))$. Then $\cdots < x_2 < x_1 < x_0$ is an infinite descending sequence in $f(0)$, which is a contradiction. $\square$

We have already seen that the reals are uncountable. By the Well Ordering Theorem, $\mathbb{R}$ can be well ordered, and hence there exists an uncountable ordinal. Therefore there exists a smallest among all uncountable ordinals; we denote this ordinal by ω_1 and denote its cardinality $\aleph_1$. More generally, we define:

Definition 3.3.4. (by transfinite recursion) Let $\omega_0 = \omega$. Let $\alpha > 0$ be an ordinal. Let ω_α be the least ordinal which is not of cardinality $|\omega_\beta|$ for any $\beta < \alpha$. Let $\aleph_\alpha = |\omega_\alpha|$.

Exercise 3.3.5 (Cantor). Every cardinal number is equal to $\aleph_\alpha$ for some α.

Recall that for $\alpha > 0$ and β ordinals, we defined α^β by transfinite recursion as 1 if $\beta = 0$; $\alpha^\gamma \alpha$ if $\beta = \gamma + 1$; and $\sup\{\alpha^\gamma : \gamma < \beta\}$ if β is a limit ordinal.

Warning 3.3.6. $|2^\omega| \neq 2^{\aleph_0}$!

Indeed, we already saw that $2^{\aleph_0} = |\mathbb{R}| = \mathfrak{c} > \aleph_0$. On the other hand, $2^\omega = \sup\{2^n : n < \omega\} = \omega$. In fact, it can be shown that ω^ω, ω^{ω^ω}, $\omega^{\omega^{\omega^\omega}}$, ... are all countable. What if the tower of ω's is taken to the limit?

Definition 3.3.7. Let $\zeta_0 = \omega$, and for $n \in \mathbb{N}$, let $\zeta_{n+1} = \omega^{\zeta_n}$. Define $\varepsilon_0 = \sup\{\zeta_n : n \in \mathbb{N}\}$. Alternatively, let ε_0 denote the least ordinal such that $\omega^{\varepsilon_0} = \varepsilon_0$.

Exercise 3.3.8. Prove that these two definitions of ε_0 are equivalent.

Definition 3.3.9. Let ε_1 denote the least ordinal $\neq \varepsilon_0$ such that $\omega^{\varepsilon_1} = \varepsilon_1$.

Exercise 3.3.10. Show that ε_0 and ε_1 are both countable.

Exercise 3.3.11. (a) For α an ordinal, define ε_α. (b) Show that if α is countable, so is ε_α. (c) Show that there exists a countable ordinal α such that $\varepsilon_\alpha = \alpha$. (d) Determine ε_{ω_1}.

Next we state an amusing result.

Theorem 3.3.12 (Base-ω Number System). *Let α be* any *ordinal. There exists a $k \in \mathbb{N}$, ordinals $\alpha_1 > \ldots > \alpha_k$, and natural numbers $n_1, \ldots, n_k$, such that $\alpha = \omega^{\alpha_1} n_1 + \cdots + \omega^{\alpha_k} n_k$. Moreover, this representation of α is unique.*

We conclude by a weak version of what, in set theory, is known as *Fodor's Lemma.*

Definition 3.3.13. Given an ordinal α, a function $f : \alpha \to \alpha$ is called *regressive* if $f(a) < a$ for all non-zero $a \in \alpha$.

Theorem 3.3.14 (Fodor). *Let $f : \omega_1 \to \omega_1$ be regressive. There exists a $y \in \omega_1$ such that $|f^{-1}(y)| = \aleph_1$.*

Exercise 3.3.15. Prove the above theorem, and relate it to an assigned puzzle problem.