GIRTH AND CHROMATIC NUMBER OF GRAPHS

AARON HALPER

Abstract. This paper will look at the relationship between high girth and high chromatic number in both its finite and transfinite incarnations. On the one hand, we will demonstrate that it is possible to construct graphs with high oddgirth and high chromatic number in all cases. We will then look at a theorem which tells us why, at least in the transfinite case, it is impossible to generalize this to include even cycles. Finally, we will use the probabilistic method to show why it is possible to construct graphs of any given finite girth and finite chromatic number.

Contents

1. Introduction

We begin with a few basic definitions:

Definition 1.1. A *graph* G is defined as a set of vertices $V = \{v_1, v_2, ..., v_n\}$ and edges E that connect the vertices to each other. We signify that two vertices v, w are connected by writing $\{v, w\}$.

Note that even though G is here defined as having only a countable number of vertices, we will later allow G to have an uncountable number of vertices.

Definition 1.2. On any graph G, we can impose a *coloring* wherein we color each of the vertices in such a way that if two vertices are connected to each other, then they cannot have the same color.

Definition 1.3. The *Chromatic Number* χ of a graph is the smallest number of colors needed to fully color the graph.

Definition 1.4. A *subgraph* G' of a graph G is a graph that can be obtained by deleting edges and/or vertices from G.

Date: DEADLINE AUGUST 22, 2008.

Definition 1.5. An *n-cycle* is a graph on n distinct vertices $v_1, v_2, ..., v_n$ with the edges $\{v_1, v_2\}, \{v_2, v_3\}, \{v_3, v_4\}, \dots, \{v_{n-1}, v_n\}, \{v_n, v_1\}$. We say that G *contains a cycle* whenever there is a subgraph of G that is a cycle. For the present purposes, we will restrict our attention to n-cycles with $n \geq 3$. We occasionally call cycles with three vertices *triangles* and cycles with five vertices *pentagons*, etc..

Definition 1.6. The *girth* of a graph G, denoted $\text{girth}(G)$ is the size of the smallest cycle in the graph. As we might expect, the *oddgirth* of a graph is the size of the smallest odd cycle in the graph.

One might think that if we do not have any "small" cycles in our graph, then we could get away with coloring our graph with just a few colors. After all, if the girth is high, then at any point the graph will locally look like a tree (i.e., a graph without any cycles at all, which would always have a chromatic number of two); however, this paper will show that it is not the case that the chromatic number can be determined by looking at what happens locally around any point. The chromatic number is actually a much more global property of the graph.

2. Oddgirth and Chromatic Number

In this section we will present a construction of a graph with high chromatic number and high oddgirth. We begin this section with yet another

Definition 2.1. A *directed graph* $\vec{G}$ is a graph, each of whose edges have associated to it a direction. The directions of the edges is entirely independent of the coloring of the graph, i.e. directed graphs may be colored in the same way that non-directed graphs are colored. Because of this, we may still talk about the *chromatic number of a directed graph* $\vec{G}$ as being the size of the minimal coloring of $\vec{G}$.

Definition 2.2. A *directed acyclic graph* or *DAG* is a graph without any directed cycles, so beginning at a point v there is no directed path that leads back to v. Note that there may be cycles in the undirected sense, however.

We will look at the following

Construction 2.3. The *Line Graph* $\overrightarrow{L(G)}$ of a directed graph $\vec{G}$ is created by replacing each edge in $\vec{G}$ with a vertex in $\overrightarrow{L(G)}$. Two vertices representing directed edges from u to v and from w to x in $\vec{G}$ are connected by an edge from uv to wx in the line graph when $v = w$. That is, each edge in $\overrightarrow{L(G)}$ represents a length-two directed path in $\vec{G}$.

The goal of this section will be to prove that this construction does indeed generate graphs with a high oddgirth and high chromatic number. In order to do this we will need these important facts:

1. $2^{\chi(\overrightarrow{L(G)})} = \chi(\vec{G})$
2. If $\vec{G}$ has no directed cycles then $\text{oddgirth}(\overrightarrow{L(G)}) > \text{oddgirth}(\vec{G})$
3. $\forall$ cardinals $\zeta, \forall n < \omega, \exists G$ such that $\chi(G) \geq \zeta$ and $\text{oddgirth}(G) \geq n$.

(1) is the easiest to prove. Suppose that we are given a coloring of $\overrightarrow{L(G)}$ which colors each $s' \in \overrightarrow{L(G)}$ with a color $\text{col}'(s')$. Simply associate to each vertex v of $\vec{G}$ $\text{in}(v) = \{\text{edges } \{w, v\}, \text{ or edges going into } v\}$ and $\text{col}(v) = \{\text{col}'(e) \mid e \in \text{in}(v)\}$. Since edges in $\vec{G}$ are vertices in $\overrightarrow{L(G)}$, this associates to each vertex $v \in \vec{G}$ the set

of colors of all edges that end in v. This is an association between the colors of $\vec{G}$ and the power set of the colors of $\overrightarrow{L(G)}$. Since this latter has cardinality $2^{\chi(\overrightarrow{L(G)})}$, this will give us the inequality that we see in (1). We need only prove that this is well-defined, i.e. that this does indeed give us a legal coloring of $\vec{G}$.

Suppose that we are given two adjacent points $v, w \in \vec{G}$ with c connecting v to w.

$$\textcircled{v} \xrightarrow{\;c\;} \textcircled{w}$$

If we call c' the set of all vertices that go into v, then we wish to show that $\mathrm{col}(v) \neq \mathrm{col}(w)$. Suppose that this is not the case, that $\mathrm{col}(v) = \mathrm{col}(w)$. By the way we defined the coloring of $\vec{G}$, we know that $\mathrm{col}'(c) = \mathrm{col}(w)$, and that $\mathrm{col}'(c') = \mathrm{col}(v)$. Putting these together yields

$$\mathrm{col}(v) = \mathrm{col}(w) \Rightarrow \mathrm{col}'(c) = \mathrm{col}'(c').$$

This means that we have two points in $\overrightarrow{L(G)}$, namely c, c' that are connected but are the same color. This is a contradiction, proving that this coloring is indeed well-defined.

Proving (2) will be a bit more tricky. If we restrict our view to triangles, it is easy to see that all of the triangles in $\overrightarrow{L(G)}$ must have come from a directed triangle in $\vec{G}$. So by making $\vec{G}$ be acyclic, we can make $\overrightarrow{L(G)}$ have no triangles. Generalizing, suppose $\mathrm{oddgirth}(\overrightarrow{L(G)}) = 2k + 1$. We then need to show that the shortest odd cycle in $\vec{G}$ has length $\leq 2k-1$. Suppose that we are given the shortest odd cycle in $\overrightarrow{L(G)}$. Then we can pick the first point that has an edge from it, say v_1, and continue along that edge. Because we said that this was the smallest graph in $\overrightarrow{L(G)}$, we know that there can be nothing of the following form (We leave this fact as an easy exercise to the reader):

$$\cdots \xrightarrow{\;e_n\;} \textcircled{v_n} \xleftarrow[e_{n+1}]{} \textcircled{v_{n+1}} \xrightarrow{\;e_{n+2}\;} \cdots$$

Restricting our view of $\overrightarrow{L(G)}$ to this cycle, we can look at each of the points where the cycle "changes direction", a vertex in the cycle where there are two ingoing edges or two outgoing edges, vertices which we will call *turning points.* Because of the shape that this will impose on $\vec{G}$, the exercise amounts to showing that we will be able to construct a new, smaller cycle of $\overrightarrow{L(G)}$ if we have two turning points that are next to each other.

$\vec{G}$ being assumed to be without any directed cycles, we know that $\overrightarrow{L(G)}$ also has no directed cycles. Therefore, we know that our cycle must contain at least one turning point, else we would have a directed cycle in $\overrightarrow{L(G)}$. Not only that, but there must be an even number of turning points in the graph since every time that the cycle is turned around, that must be undone at some other turning point. We can select those points that are *not* turning points, and discover that the graph in $\vec{G}$ that corresponds to our cycle is made up of only these points (exercise). This gives us the shortest odd cycle in $\vec{G}$, a smaller cycle than the shortest odd cycle in $\overrightarrow{L(G)}$, proving (2).

After all of that, in order to prove (3), we need only put (1) and (2) together. Let's do it first for the case where $n = 5$. Consider $\vec{H} = K_{(2^\zeta)^+}$, the complete graph on $(2^\zeta)^+$ many vertices, where $(2^\zeta)^+$ denotes the successor cardinal of 2^ζ. Then we can make this graph a DAG by orienting all of the edges in one direction, say left. In particular, $\vec{H}$ has no directed triangles, so we can be certain that $\overrightarrow{L(H)}$ does not have any triangles at all. Because $\vec{H}$ is a complete graph, $\chi(\vec{H}) = (2^\zeta)^+$. Then by (1), $\chi(\overrightarrow{L(H)}) > \zeta$, and by (2), $\overrightarrow{L(H)} > 3$, exactly what we wanted.

We can do this for $n > 5$ as well by taking repeated iterations of the line graph function. That is, given n we can take

$$\vec{H} = K_{(\underbrace{2^{2^{\cdot^{\cdot^{\cdot^{2^\zeta}}}}}}_{\frac{n-1}{2}\text{ times}})^+}$$

and consider

$$\vec{H}' = \underbrace{L(L(\cdots L}_{\frac{n-1}{2}\text{times}}(H))).$$

Then by repeated iterations of (1), $\chi(\vec{H}') = \zeta$ and by repeated iterations of (2), $\vec{H}'$ has no cycles of size n or smaller, exactly what we are looking for. □

3. Graphs with Uncountable Chromatic Number

In the last section we constructed graphs with a large *odd*girth and with high, even infinitely high, chromatic number. It is not immediately obvious though why it is that we cannot generalize this to include the even cycles as well, to girth from oddgirth.

The goal of this section will be to prove a theorem stating that any graph with an uncountable chromatic number, $\chi(G) \geq \aleph_1$ must also contain a four-cycle, $K_{2,2}$.

We first require the following lemma:

Lemma 3.1. *If $G = H_1 \cup H_2$ then $\chi(G) \leq \chi(H_1)\chi(H_2)$*

Proof. Let's take $\lambda_1 = \chi(H_1)$ and $\lambda_2 = \chi(H_2)$. Now, H_1 already has a coloring c_1 on it which takes the vertices and maps them to the set of colors, that is $c_1 : v_1 \to \lambda_1$, and simlarly for H_2 we have that $c_2 : v \to \lambda_2$. In order to combine the coloring of these two graphs we can take their euclidean product $|\lambda_1 \times \lambda_2|$ and take as our new coloring $c(v) = (c_1(v), c_2(v))$. Then this coloring is a legal coloring, an easy fact the proof of which we leave to the reader. □

The proof of this section's theorem will hinge on the concept of closure of an operation. Suppose that we are given an operation

$$\begin{aligned} \circ : A \times A &\to A \\ (a, b) &\mapsto a \circ b \end{aligned}$$

Even without imposing any further conditions such as associativity or commutativity on this operation, given $B \subset A, |B| = \lambda$, we can still look at the substructure generated by B, and ask about the size of this new substructure. This quantity, $|\langle B\rangle|$ will be found by looking at all possible words composed of the $a_i \in B$ for a given length k. Since we did not assume any associativity we must also look at the number of ways of inserting parentheses into the word. Summing over all the lengths k will tell us the final size. This method gives us that

$$|\langle B\rangle| = \sum_{k<\omega} \lambda = \lambda * \aleph_0 = \lambda$$

For λ any infinite ordinal. We are now ready to begin the proof of the theorem that is the star of this section.

Theorem 3.2. *(Erdős and Hajnal)*
If $\chi(G) \geq \aleph_1$, then $G \supseteq K_{2,2}$.

Proof. We will prove this thoerem in the equivalent contrapositive form: If G does not contain a four-cycle, i.e. $G \supseteq K_{2,2}$ then the chromatic number of G is at most countable, which we denote by $\chi(G) \leq \aleph_0$. In addition we will at first assume that $|V| = \aleph_1$.

The proof proceeds by taking the so-called closure operation, which takes any two points and gives any points that are connected to both of them. Now, if two points were to share at least two common neighbors then these four points would form a four-cycle, which we have assumed do not appear in our graph. Thus, any two points may share at most one common neighbor, and hence the closure operation will give us at most one new point per pair of points. That is, given our set B we are taking any two points a and b and getting at most one new point $a \circ b$. If $|B| = \lambda$ and λ is an infinite cardinal, then there are λ^2 ways of doing this, so we form a new set B' which has at most $\lambda^2 = \lambda$ cardinality. We can iterate this to get B'' which will have $(\lambda^2)^2 = \lambda^4 = \lambda$ cardinality. After repeated iterations we get that the total cardinality is

$$\begin{aligned} \lambda + \lambda^2 + \lambda^4 + ... &= \\ \lambda + \lambda + \lambda + ... &= \\ \lambda * \aleph_0 &= \lambda \end{aligned}$$

This method of iterating the closure operation to ω we will call the closure process. The proof proceeds by creating C_α such that $C_1 \supset \omega$ and C_1 is closed, where we say that a set T is closed if $\forall x, y \in T, z\ x, y \Rightarrow z \in T$. To make C_2, we take $v_1 = \min\{\omega \setminus C_1\}$ and apply the closure process to both C_1 and v_1 to form C_2. Along these lines, we define

$$C_\alpha = \begin{cases} \text{closure}(C_\beta \cup v_\beta) & \text{for } \alpha = \beta + 1 \\ \bigcup_{\beta<\alpha} C_\beta & \text{for } \alpha \text{ a limit ordinal} \end{cases}$$

This defines C_α for all ordinals $\alpha < \omega_1$. From this definition we can see that $C_1 \subset C_2 \subset C_3 \subset ...$ and we can look at this transfinitely. From this we make the following claim:

$$|C_\alpha| = \aleph_0,\ \forall\ \alpha.$$

We will prove this claim by making use of transfinite induction. Because we are using induction, we can assume that the claim is true for all ordinals less than α. Case 1: If α is a succesor ordinal, then $\alpha = \beta + 1$. Then $C_\alpha = \text{closure}(C_\beta \cup v_\beta)$ And by the induction hypothesis, C_β is countable, hence C_α is the closure of two countable sets, hence is itself countable, as we have seen. Case 2: If α is a limit ordinal, then C_α is defined as the union of countably many sets, hence is again itself countable as we have seen.

The next fact about these C_α's that we will prove is that they "take up" the entire set ω_1. We will show, in other words, that $\bigcup_{\alpha<\omega_1} C_\alpha = \omega_1$. To see this, we first need to understand why it is that, for all ordinals γ, $v_\gamma \geq \gamma$. We formulate this claim as a lemma.

Lemma 3.3. *For any ordinal* $\gamma, v_\gamma \geq \gamma$.

Proof. To prove this, we need first show that for any given ordinals γ, δ; $\gamma \leq \delta \Rightarrow v_\gamma \leq v_\delta$. This is accomplished by recalling that, $\forall \alpha, C_1 \subset C_2 \subset ... \subset C_\alpha \subset ...$ and that the v_β's are chosen by picking the point that is just *after* the set C_β. So whenever $\gamma \leq \delta$, so too $C_\gamma \leq C_\delta$, hence we will choose v_γ before we choose v_δ, so we see that $v_\gamma \leq v_\delta$.

We are now ready to prove the lemma, by transfinite induction. The base case here is easy, for $1 \in \omega \subset C_1$, and v_1 was chosen as the element just after C_1, so v_1 must be greater than C_1, in particular greater than 1. Now suppose that $\forall \rho < \gamma, v_\rho \geq \rho$, the induction hypothesis. Then, by what we just proved, $v_\gamma > v_\rho \geq \rho$. Since this is true for all $\rho < \gamma$, we have proved that $v_\gamma \geq \gamma$. □

Essentially what this means is that our method of creating C_α's does not "back-track" on itself, but always picks up more points as it continues, telling us that we will have exhausted $\aleph_1$ after ω_1 many steps (and not sooner since the cofinality of ω_1 is ω_1).

We are now ready to attempt to color in the graph. From the C_α's, we define $D_\alpha = C_\alpha \setminus C_{\alpha-1}$. These D_α's cover ω_1, but each D_α is countable. Let's color all of the *edges* that connect two points within a D_α red, and all of the *edges* connecting points that are in different D_α's blue.

The chromatic number of the red graph is not difficult to determine. Because each D_α contains only a countable number of points, we can color each D_α with at most countable colors. We'll color each D_α as a separate componenent, and since we can color all of the D_α's separately, we can thereby color their union with at most countable colors. We now have only the coloring of the blue graph to worry about.

We will prove that the blue graph is, in fact, two-colorable. We picked the blue graph as the connections between the D_α's, and how many of these can there be? Well, suppose that there is a vertex $w_\alpha \in D_\alpha$ connected to both $w_\beta \in D_\beta$ and $w_\gamma \in D_\gamma$ with $\beta, \gamma < \alpha$. This means that $\exists \delta$ such that $w_\beta, w_\gamma \in C_\delta$, but $w_\alpha \notin C_\delta$. This is a contradiction with the way that the C_α's were constructed. so from this argument we can see that any point w_α in the blue graph can only be connected to one point less than itself in the blue graph.

Suppose we had a cycle in the blue graph. Then every point in the cycle would have degree at least two. In particular, we can look at the maximum point in the cycle. This max point must be connected to two other points in the cycle, but because it is the max point, these other two points must be less than itself. Thus, there is a point connected to two points less than itself, contradicting the previous argument. Therefore, the blue graph can have no cycles in it, hence it forms a tree, and is therefore two-colorable.

We have thus far shown that the whole graph is the union of two graphs: the red graph and the blue graph. By Lemma 3.1,

$$\chi(G) \leq \chi(\text{Blue Graph}) * \chi(\text{Red Graph}) = 2 * \aleph_0 = \aleph_0,$$

thereby completing the proof in the case that $|G| = \aleph_1$.

To generalize this for $|G| = \aleph_2$ we need to repeat the same proof with an important modification. We choose our new $C_1 \supset \omega_1$ instead of ω_0. Then by the case that we have just proven C_1 is countably-colorable, and so we can cover all of $\aleph_2$ with these new C_α's, and repeat the rest of the proof. And in a similar fashion we can prove the theorem for all uncountable cardinalities. □

4. The Probabilistic Method

We still want to construct graphs with high girth and high chromatic number, even if that is impossible in the case of uncountable chromatic number. Is it possible to construct *finite* graphs with high girth and high chromatic number? By the previous theorem, any method that could generate such graphs could not be generalizable to the transfinite case. For this reason, a method such as the one that we attempted in section 2 will not work. What does work is method called the probabilistic method. The following theorem, due to Erdös, proves that there are graphs with as high a finite chromatic number as we want even if we demand that the graph not have any cycles less than a given size.

Theorem 4.1. *(Erdős)*
Given any $k, l \in \mathbb{N}$ there is a graph G such that $\chi(G) > k$ and girth(G) $> l$.

Proof. The proof makes use of a technique called the probabilistic method. This proceeds as follows: rather than actually constructing such a graph, the idea will be to look at a *random* graph and prove that the probability that this graph will satisfy the hypotheses is greater than zero. Even though we will not actually construct such a graph, since the graph exists with positive probability, then we know that it must exist.

To do this, we first pick a θ so that $0 < \theta < 1/l$. Then we take our graph G on n vertices, and we say that the probability of any two vertices being connected is $p = n^{\theta-1}$. This forms a random graph on n vertices and we now want to look at X, the number of cycles in the graph that are of length at most l. Now, since we are taking our graph on n vertices, the number of cycles of a given length i is certainly no more than n^i and the probability of each of them occuring is p^i, so

$$E[X] \leq \sum_{i=3}^{l} n^i p^i$$

where $E[X]$ is the expected value of X. This yields the following geometric series:

$$E[x] \leq \sum_{i=3}^{l} n^i n^{(\theta-1)*i} = \sum_{i=3}^{l} n^{i\theta} \leq \frac{n^{l\theta}}{1 - n^{-\theta}}$$

Since $l\theta < 1$, we see that this fraction approaches zero as $n \to \infty$. Therefore, we know that there is some n so large that $E[X] < n/4$. We then use Markov's Inequality to relate this expected value to the actual probability. Markov's Inequality states that for a random variable Y, and for $t > 0$,

$$Pr[Y \geq t] \leq \frac{E[Y]}{t}.$$

Applying this to the situation at hand, we see that

$$Pr[X \geq n/2] < \frac{n/4}{n/2} = 1/2.$$

What about chromatic number? Instead of looking at the chromatic number directly, we will only look at $\alpha(G)$, the size of the largest independent set in G. An independent set is a set of vertices in G that are not connected to each other. For example, if we colored a graph, then all of the red points in our graph would form an independent set. Of course, the other way around may not work: it may be that we can pick an independent set that cannot be colored by only one color. But there cannot be more points of any one color than there are points in the largest independent set. This leads us to the following formula:

$$\chi(G) \geq \frac{|V|}{\alpha(G)} \tag{4.2}$$

How do we find a bound on α? In any set of size a the probability that any two points will be unconnected is $1-p$, and there are $\binom{a}{2}$ ways of picking out two points within this set of a points. Since there are $\binom{n}{a}$ such sets of size a, we see that

$$Pr[\alpha(G) \geq a] \leq \binom{n}{a}(1-p)^{\binom{a}{2}}$$

We can get a bound on the right hand side of this by realizing that $\binom{n}{a} \leq n^a$ and that $(1-p)^{\binom{a}{2}} \leq e^{-pa(a-1)/2}$. Putting these together tells us that

$$Pr[\alpha(G) \geq a] \leq n^a e^{-\frac{pa(a-1)}{2}}.$$

To further refine this, we note that so far we have left a undefined. If we set $a = \lceil \frac{3}{p} \ln n \rceil$ then we get that

$$Pr[\alpha(G) \geq a] \leq n^a n^{-\frac{3(a-1)}{2}}$$

And we can see that as n gets very large, the right hand side of this inequality gets closer and closer to zero. Therefore, there is some n large enough so that $Pr[\alpha(G) \geq a] < 1/2$. We now have a bound for both the number of cycles that have length at most l, and a bound for $\alpha(G) \geq a$, telling us that we can take the union bound to get that

$$Pr[X \geq n/2 \text{ and } \alpha(G) \geq a] < \frac{1}{2} + \frac{1}{2} = 1$$

Because this probability is less than one, the probability that it won't take place will be greater than zero, telling us that there is a graph H with the number of short cycles is $X < n/2$ and the size of the largest independent set is $\alpha(G) < a$. From this graph we look at all those cycles, and randomly pick one of them out, which we then delete. The resulting graph H' will have no cycles of length less than l and since there are less than $n/2$ such cycles, H' will have at least $n/2$ vertices left in it. There is no way that deleting vertices in this way will ever increase the size of the large independent set, so we get that $\alpha(H') < a$. Using the formula that we saw before, 4.2, we get

$$\chi(H') \geq \frac{|V(H')|}{\alpha(H')} \geq \frac{n/2}{3n^{1-\theta}\ln n} = \frac{n^{\theta}}{6\ln n}.$$

By making n very large, we can increase the size of the right hand size of this equation, thereby increasing chromatic number of our graph. In particular, we can make so large that $\chi(H') > k$. This tells us that there is a graph H' without any cycles of length l or less and with a chromatic number greater than k, exactly what we wanted. □

Acknowledgements. It is a pleasure to thank my mentor, Nguyen T. T. Nguyen for all of the encouragement along the way. I would also to thank Professor László Babai for all of his help and guidance.

References

[1] Paul Erdős and András Hajnal. On Chromatic Number of Graphs and Set-Systems. Acta Mathematica Hungarica 17 (1966): 61-99.

[2] Noga Alon, Joel H. Spencer, and Paul Erdős. The Probabilistic Method. John Wiley & Sons, Inc. 1992.