Problem Set 4

This fourth problem set explores strict orders, graph theory, and the pigeonhole principle. Over the course of working through these problems, you will get a better intuitive feel for these concepts, will see their applications in the real world, and will solidify your proofwriting techniques. We hope you have as much fun working through these problems as we did researching and developing them.

As always, please feel free to drop by office hours, email us, or ask on Piazza if you have any questions. We'd be happy to help out.

Good luck, and have fun!

Checkpoint due Monday, February 1 at the start of lecture.

Remaining problems due Friday, February 5 at the start of lecture.

Write your solutions to the following problem and submit them online by Monday, February 1st at the start of class. This problem will be graded on a 0/1/2 scale based on whether you have attempted to solve the problem rather than on correctness. We will try to get these problems returned to you with feedback on your proof style this Wednesday, February 3rd.

Checkpoint Problem: Strict Orders (2 Points)

We defined strict orders as binary relations that are irreflexive, asymmetric, and transitive. We proved in lecture that any relation that is asymmetric and transitive is also a strict order, meaning that we could have potentially left irreflexivity out of our definition of strict orders.

Interestingly, it turns out that we could have also left asymmetry out of our definition and just gone with irreflexivity and transitivity.

Prove that a binary relation R over a set A is a strict order if and only if R is irreflexive and transitive.

The remaining problems on this problem set are due on Friday, February 5th at the start of lecture.

Problem One: Covering Relations (3 Points)

Let $<_A$ be a strict order relation over a set A. We can define a new binary relation $\lessdot_A$, called the ***covering relation for*** $<_A$, as follows:

$$x \lessdot_A y \quad \text{if} \quad x <_A y \text{ and there is no } z \in A \text{ where } x <_A z \text{ and } z <_A y$$

This question explores properties of covering relations.

i. Consider the $<$ relation over the set $\mathbb{N}$. What is its covering relation? To provide your answer, fill in the blank below, then briefly justify your answer:

$$x \lessdot y \quad \text{if} \quad \underline{\hspace{6cm}}$$

ii. Prove that the relation $\lessdot$ you found in part (i) is ***not*** a strict order.

iii. Let $<_A$ be a strict order over a set A. There is a close connection between the Hasse diagram of $<_A$ and its covering relation $\lessdot_A$. What is it? No formal proof is required.

Problem Two: Graph Coloring (6 Points)

The ***degree*** of a node v in a graph G is the number of edges that have v as an endpoint. In other words, it's the number of edges v directly touches. Interestingly, there usually isn't much of a connection between the degree of the nodes in a graph and the number of colors necessary to color that graph.

i. Give an example of a 2-colorable graph where some node has degree seven. Briefly justify why your graph meets these criteria; no proof is necessary.

ii. Generalize your answer from part (i) by describing how, for any $n \geq 0$, you can build a 2-colorable graph where some node has degree at least n. This shows that there is no direct connection between the *maximum* degree of a node in a graph and the chromatic number of that graph.

iii. Give an example of a 2-colorable graph where every node has degree three. Briefly justify why your graph meets these criteria; no proof is necessary.

iv. Generalize your answer from part (iii) by describing how, for any $n \geq 0$, you can build a 2-colorable graph where every node has degree at least n. This shows that there is no direct connection between the *minimum* degree of a node in the graph and the chromatic number of that graph.

v. Give an example of a graph where every node has degree two but which is *not* 2-colorable. Briefly justify why your graph meets these criteria; no proof is necessary.

vi. Generalize your answer from part (v) by describing how, for any $n \geq 0$, you can build a connected graph with at least n nodes where every node has degree two but which is not 2-colorable. This shows that each part of a graph can look 2-colorable even though the graph as a whole is not.

Problem Three: Complements and Connectivity (4 Points)

If $G = (V, E)$ is an undirected graph, the ***complement of G***, denoted G^c, is a graph related to the original graph G. Intuitively, G^c has the same nodes as G, and its edges consist of all the edges missing from graph G. Formally speaking, G^c is the graph with the same nodes as G and with edges determined as follows: the edge $\{u, v\}$ is present in G^c if and only if $u \neq v$ and the edge $\{u, v\}$ is not present in G. As an example, here's a graph G and its complement graph G^c:

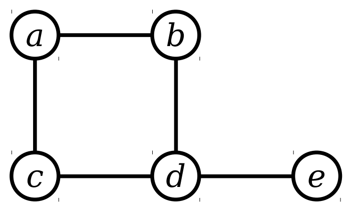

Graph G

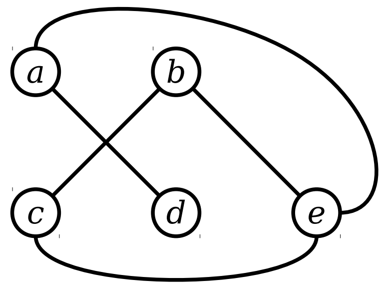

Graph G^c

Recall that a graph G is called ***connected*** if there is a path between any two nodes in G.

Prove that if G is an undirected graph, then G is connected or G^c is connected (or both). As a hint, look at Handout 14 and our advice about how to prove a statement of the form $P \vee Q$.

Problem Four: Bipartite Graphs (5 Points)

The ***bipartite graphs*** are a special class of graphs with applications throughout computer science. An undirected graph $G = (V, E)$ is called ***bipartite*** if there is a way to partition the nodes V into two sets V_1 and V_2 so that every edge in E has one endpoint in V_1 and the other in V_2.

To help you get a better intuition for bipartite graphs, let's consider an example. Suppose that you have a group of people and a list of restaurants. You can illustrate which people like which restaurants by constructing a bipartite graph where V_1 is the set of people, V_2 is the set of restaurants, and there's an edge from a person p to a restaurant r if person p likes restaurant r.

Bipartite graphs have many interesting properties. One of the most fundamental is this one:

An undirected graph is bipartite if and only if it contains no cycles of odd length.

Intuitively, a bipartite graph contains no odd-length cycles because cycles alternate between the two groups V_1 and V_2, so any cycle has to have even length.

The trickier step is proving that if G contains no cycles of odd length, then G has to be bipartite. For now, assume that G has just one connected component; if G has multiple connected components, we can treat each one as a separate graph for the purposes of determining whether G is bipartite. (You don't need to prove this, but I'd recommend taking a minute to check why this is the case.)

Suppose G is an undirected graph with no cycles of odd length. Choose any node $v \in V$. Let V_1 be the set of all nodes that are connected to v by a path of odd length and V_2 be the set of all nodes connected to v by a path of even length. (Note that these paths do not have to be simple paths).

i. Prove that V_1 and V_2 have no nodes in common.

ii. Using your result from part (i), prove that if G has no cycles of odd length, then G is bipartite.

Problem Five: Outerplanar Graphs (3 Points)

If G is a graph, the ***augmentation*** of G, denoted Aug(G), is formed by adding a new node $v_\star$ to G, then adding edges from $v_\star$ to each other node in G. For example, below is a graph G and its augmentation Aug(G). To make it easier to see the changes between G and Aug(G), we've drawn the edges added in Aug(G) using dashed lines:

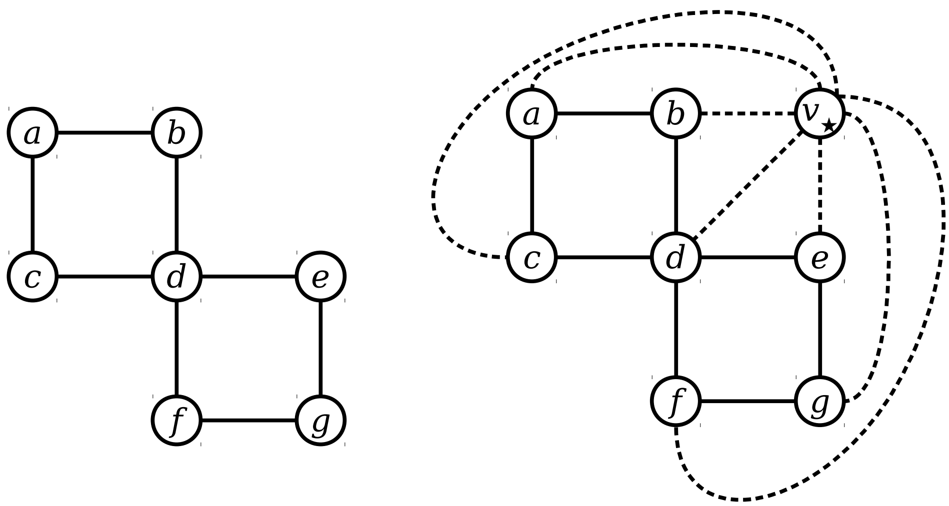

Graph G Aug(G)

Here's one additional definition: an undirected graph G is called an ***outerplanar graph*** if Aug(G) is a planar graph. In other words, if Aug(G) is a *planar* graph, then the original graph G is an *outerplanar* graph.

One of the landmark results about planar graphs is the ***four-color theorem***: every planar graph is 4-colorable. Prove the ***three-color theorem***: every outerplanar graph is 3-colorable.

Problem Six: Chromatic and Independence Numbers (5 Points)

Let's introduce a new definition. An ***independent set*** in an undirected graph $G = (V, E)$ is a set $I \subseteq V$ such that if $x \in I$ and $y \in I$, then $\{x, y\} \notin E$. Intuitively, an independent set in G is a set of nodes where no two nodes in I are adjacent. The ***independence number*** of a graph G, denoted $\alpha(G)$, is the size of the largest independent set in G.

Consider the following graph G:

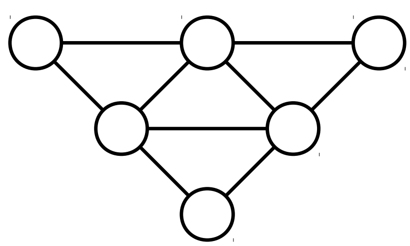

i. What is $\chi(G)$? What is $\alpha(G)$? No justification is required.

ii. Let r and s be arbitrary positive natural numbers. Prove that if G is an undirected graph with $rs+1$ nodes, then $\chi(G) \geq r+1$ or $\alpha(G) \geq s+1$ (or both). As a hint, we again recommend looking at Handout 14 and our advice about how to prove a statement of the form $P \vee Q$.

Problem Seven: Chains and Antichains (7 Points)

Let A be a set and $<_A$ be a strict order over A. A ***chain in*** $<_A$ is a series of elements $x_1, \ldots, x_k$ drawn from A such that

$$x_1 <_A x_2 <_A \ldots <_A x_k.$$

Intuitively, a chain is a series of values in ascending order according to the strict order $<_A$. The ***length*** of a chain is the number of elements in that chain.

i. Consider the $\subset$ relation over the set $\wp(\{a, b, c\})$. What is the length of the longest chain in this strict order? Give an example of a chain with that length. No justification is necessary. *(Hint: Draw the Hasse diagram and see if you can find a visual intuition for the definition of a chain.)*

Now, let's cover a new definition. An ***antichain*** is a set $X \subseteq A$ such that no two elements in X can be compared by the $<_A$ relation. In other words, a set $X \subseteq A$ is an antichain if

$$\forall a \in X. \forall b \in X. (a \not<_A b \quad \land \quad b \not<_A a)$$

The ***size*** of an antichain X is the number of elements in X.

ii. Consider the $\subset$ relation over the set $\wp(\{a, b, c\})$. What is the size of the largest antichain in this strict order? Give an example of an antichain with that size. No justification is necessary. *(Hint: Draw the Hasse diagram and see if you can find a visual intuition for the definition of an antichain.)*

Given an arbitrary strictly ordered set, you can't say anything a priori about the size of the largest chain or antichain in that strict order. However, you can say that at least one of them must be relatively large relative to the strictly ordered set.

Let r and s be natural numbers. We're going to ask you to prove the following result: if $|A| = rs+1$, then either A contains a chain of length $r+1$ or an antichain of size $s+1$.

iii. For each element $a \in A$, we'll say that the ***height*** of a is the length of the longest chain whose final element is a. Prove that if A does not contain a chain of length $r+1$ or greater, then there must be at least $s+1$ elements of A at the same height.

iv. Your result from part (iii) establishes that there must be a collection of $s+1$ elements of A at the same height as one another. Let X be any set of $s+1$ such elements. Prove that X must be an antichain.

Intuitively speaking, if $<_A$ is a strict order over A that represents some prerequisite structure on a group of tasks, a chain represents a series of tasks that have to be performed one after the other, and an antichain represents a group of tasks that can all be performed in parallel (do you see why?) In the context of parallel computing, the result you've proved states that if a group of tasks doesn't contain long dependency chains, that group must have a good degree of parallelism.

Extra Credit Problem: *k*-Regular Graphs (1 Point Extra Credit)

An undirected graph G is called ***k-regular*** if every node in G has degree exactly k. The ***girth*** of a graph is the length of the shortest simple cycle in G. If G has no cycles, its girth is infinite.

Prove that any k-regular graph with girth five has at least $k^2 + 1$ nodes.