Graph Theory

Part Two

Recap from Last Time

A ***graph*** is a mathematical structure for representing relationships.

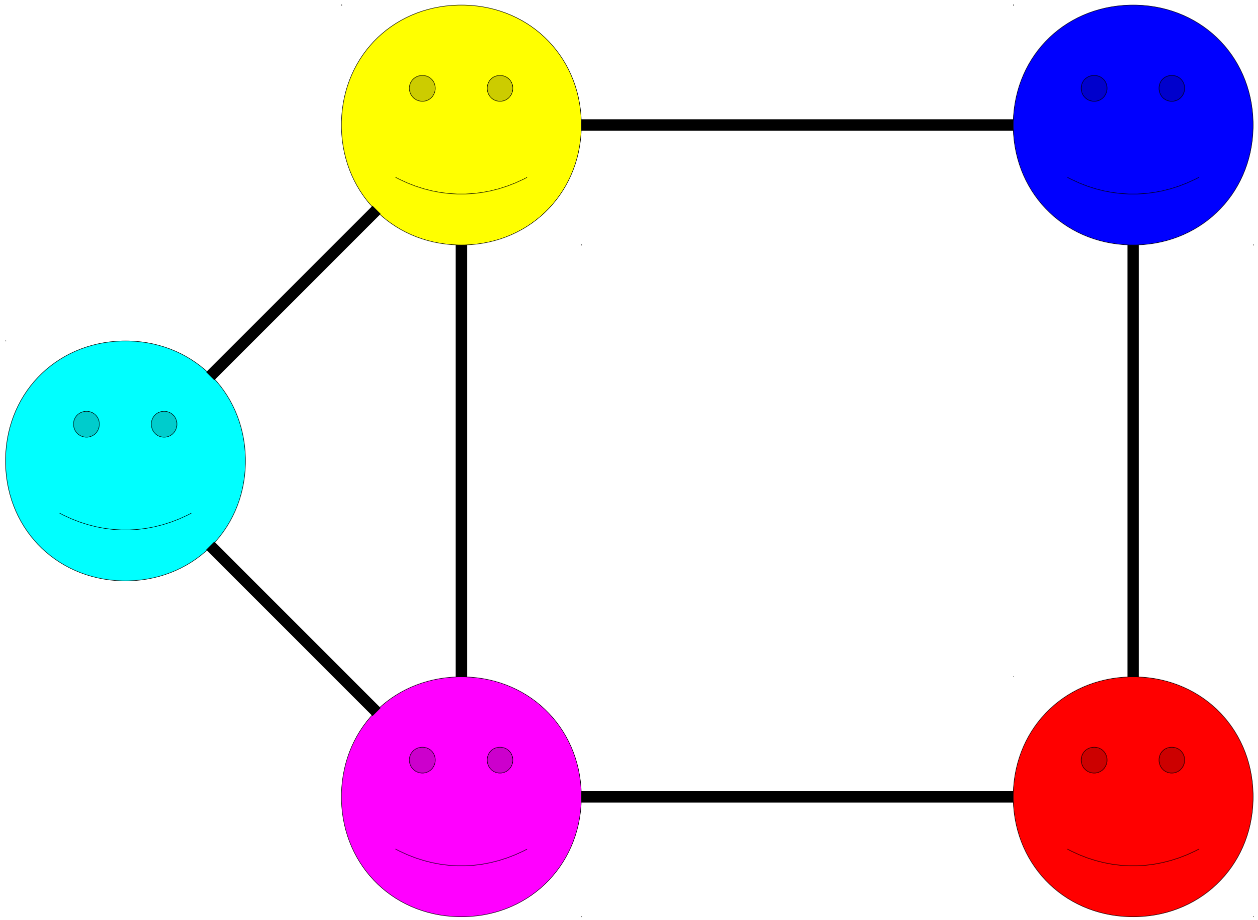

A graph consists of a set of ***nodes*** (or ***vertices***) connected by ***edges*** (or ***arcs***)

Adjacency and Connectivity

- Two nodes in a graph are called ***adjacent*** if there's an edge between them.
- Two nodes in a graph are called ***connected*** if there's a path between them.
    - A path is a series of one or more nodes where consecutive nodes are adjacent.

k-Colorability

- If $G = (V, E)$ is a graph, a ***k-coloring*** of G is a way of assigning colors to the nodes of G, using at most k colors, so that no two nodes of the same color are adjacent.

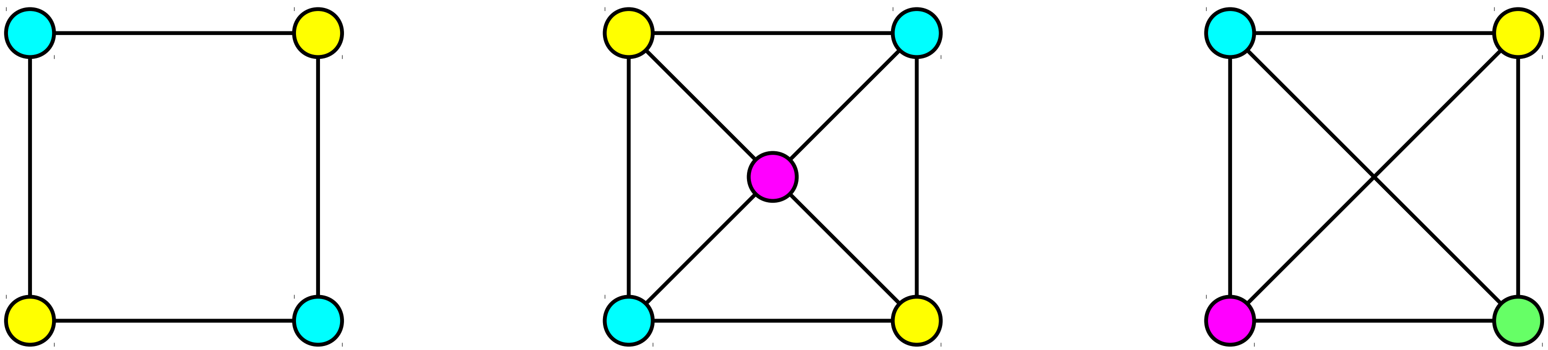

- The ***chromatic number*** of G, denoted $\chi(G)$, is the minimum number of colors needed in any k-coloring of G.
- Today, we're going to see several results involving coloring parts of graphs. They don't necessarily involve *k-colorings* of graphs, so feel free to ask for clarifications if you need them!

New Stuff!

The Pigeonhole Principle

The Pigeonhole Principle

- ***Theorem (The Pigeonhole Principle):*** If m objects are distributed into n bins and $m > n$, then at least one bin will contain at least two objects.

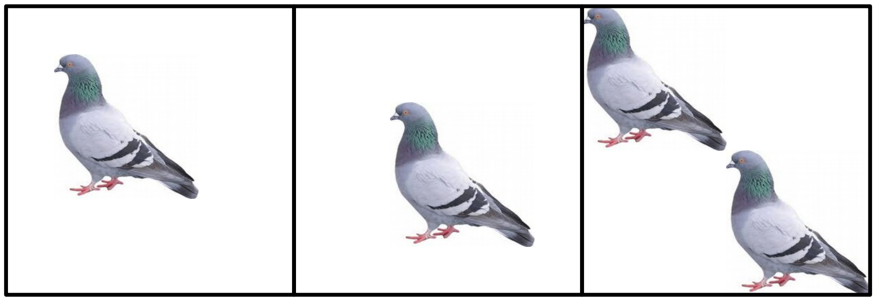

Some Simple Applications

- Any group of 367 people must have a pair of people that share a birthday.
  - 366 possible birthdays (pigeonholes)
  - 367 people (pigeons)
- Two people in San Francisco have the exact same number of hairs on their head.
  - Maximum number of hairs ever found on a human head is no greater than 500,000.
  - There are over 800,000 people in San Francisco.
- Each day, two people in New York City drink the same amount of water, to the thousandth of a fluid ounce.
  - No one can drink more than 50 gallons of water each day.
  - That's 6,400 fluid ounces. This gives 6,400,001 possible numbers of thousands of fluid ounces.
  - There are about 8,000,000 people in New York City proper.

Theorem: If m objects are distributed into n bins and $m > n$, then there must be some bin that contains at least two objects.

Proof: Suppose for the sake of contradiction that, for some m and n where $m > n$, there is a way to distribute m objects into n bins such that each bin contains at most one object.

Number the bins 1, 2, 3, …, n and let x_i denote the number of objects in bin i. There are m objects in total, so we know that

$$m = x_1 + x_2 + \ldots + x_n.$$

Since each bin has at most one object in it, we know $x_i \leq 1$ for each i. This means that

$$\begin{aligned} m &= x_1 + x_2 + \ldots + x_n \\ &\leq 1 + 1 + \ldots + 1 \quad \text{(}n\text{ times)} \\ &= n. \end{aligned}$$

This means that $m \leq n$, contradicting that $m > n$. We've reached a contradiction, so our assumption must have been wrong. Therefore, if m objects are distributed into n bins with $m > n$, some bin must contain at least two objects. ■

Pigeonhole Principle Party Tricks

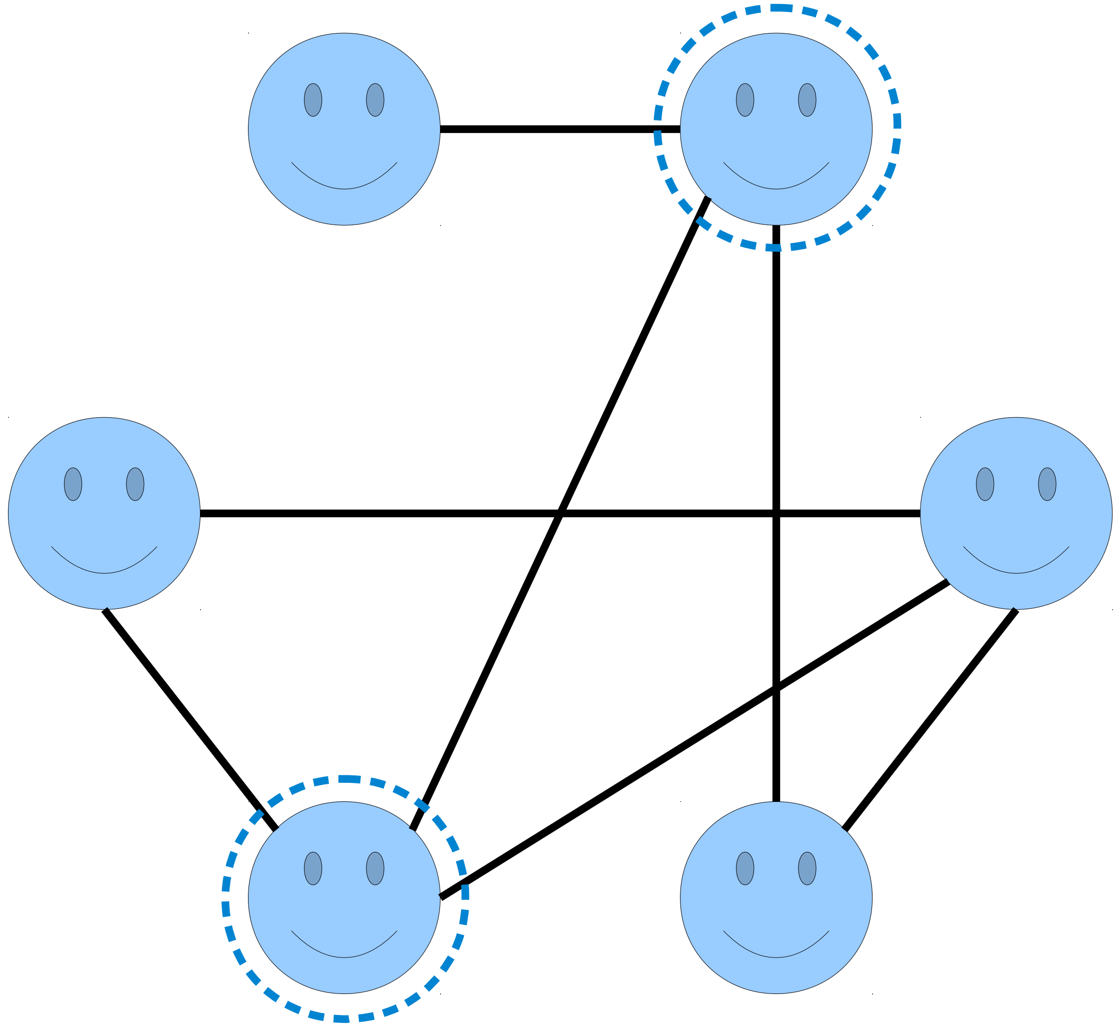

Degrees

- The ***degree*** of a node *v* in a graph is the number of nodes that *v* is adjacent to.

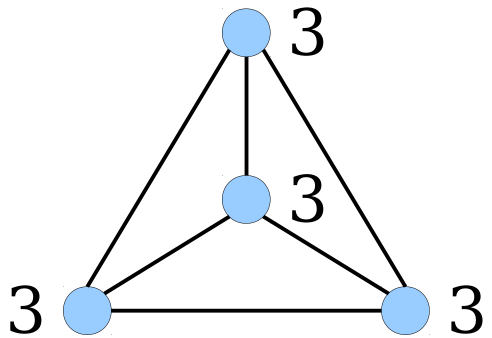

- ***Theorem:*** Every graph with at least two nodes has at least two nodes with the same degree.
  - Equivalently: at any party with at least two people, there are at least two people with the same number of Facebook friends at the party.

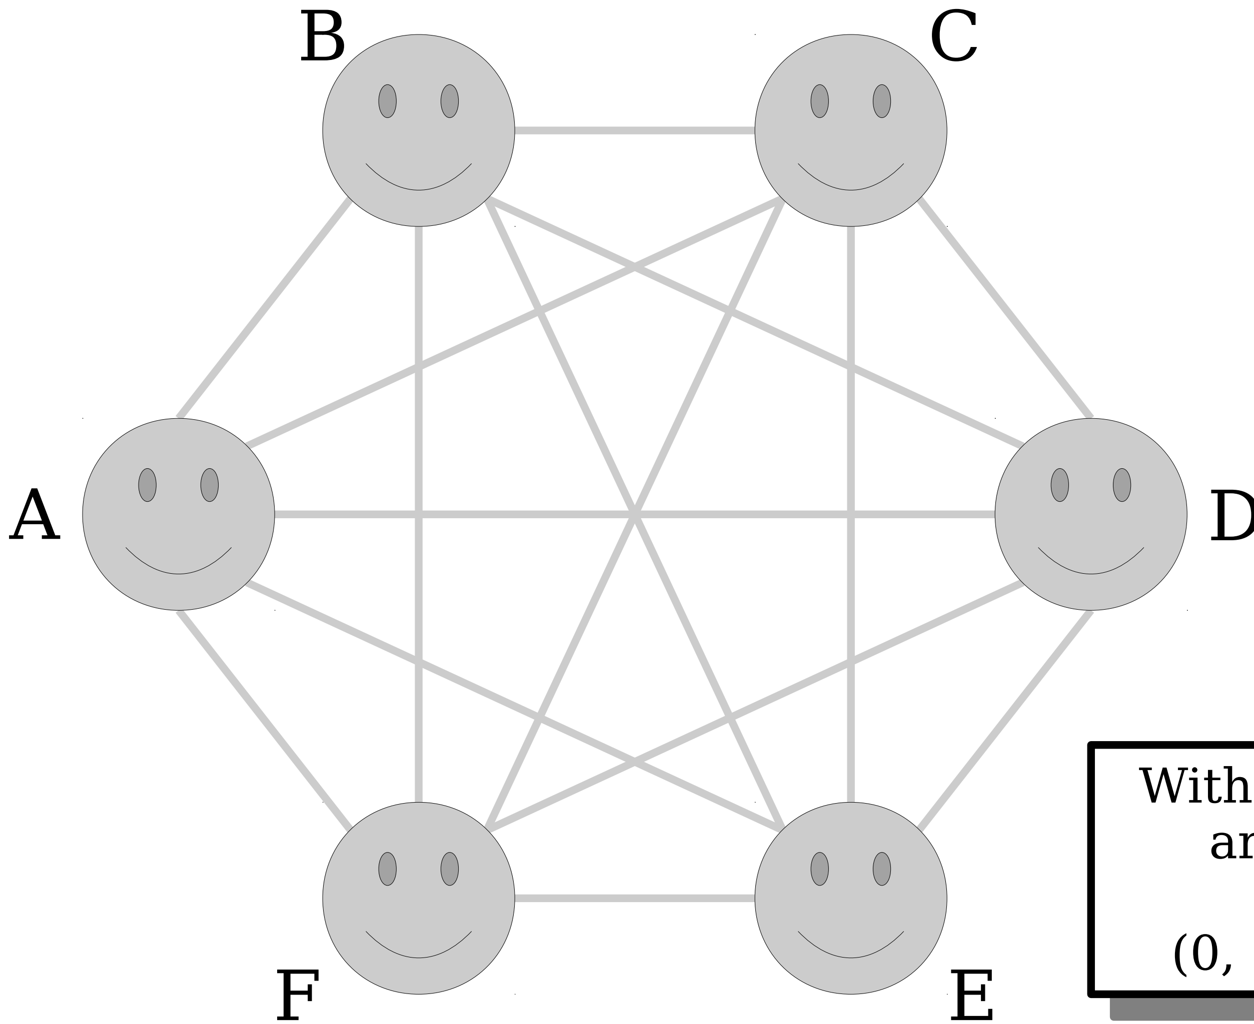

With n nodes, there are n possible degrees ($0, 1, 2, \ldots, n - 1$)

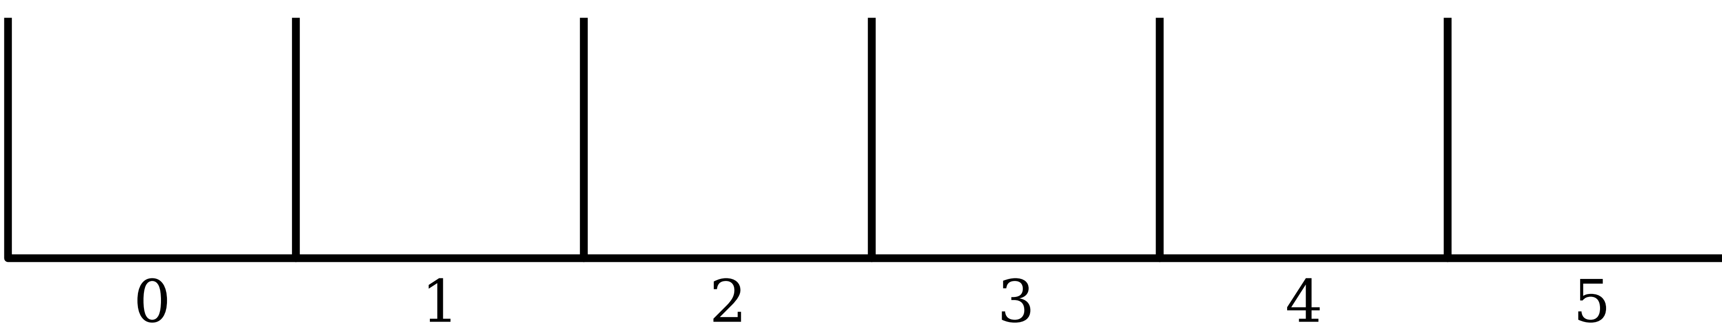

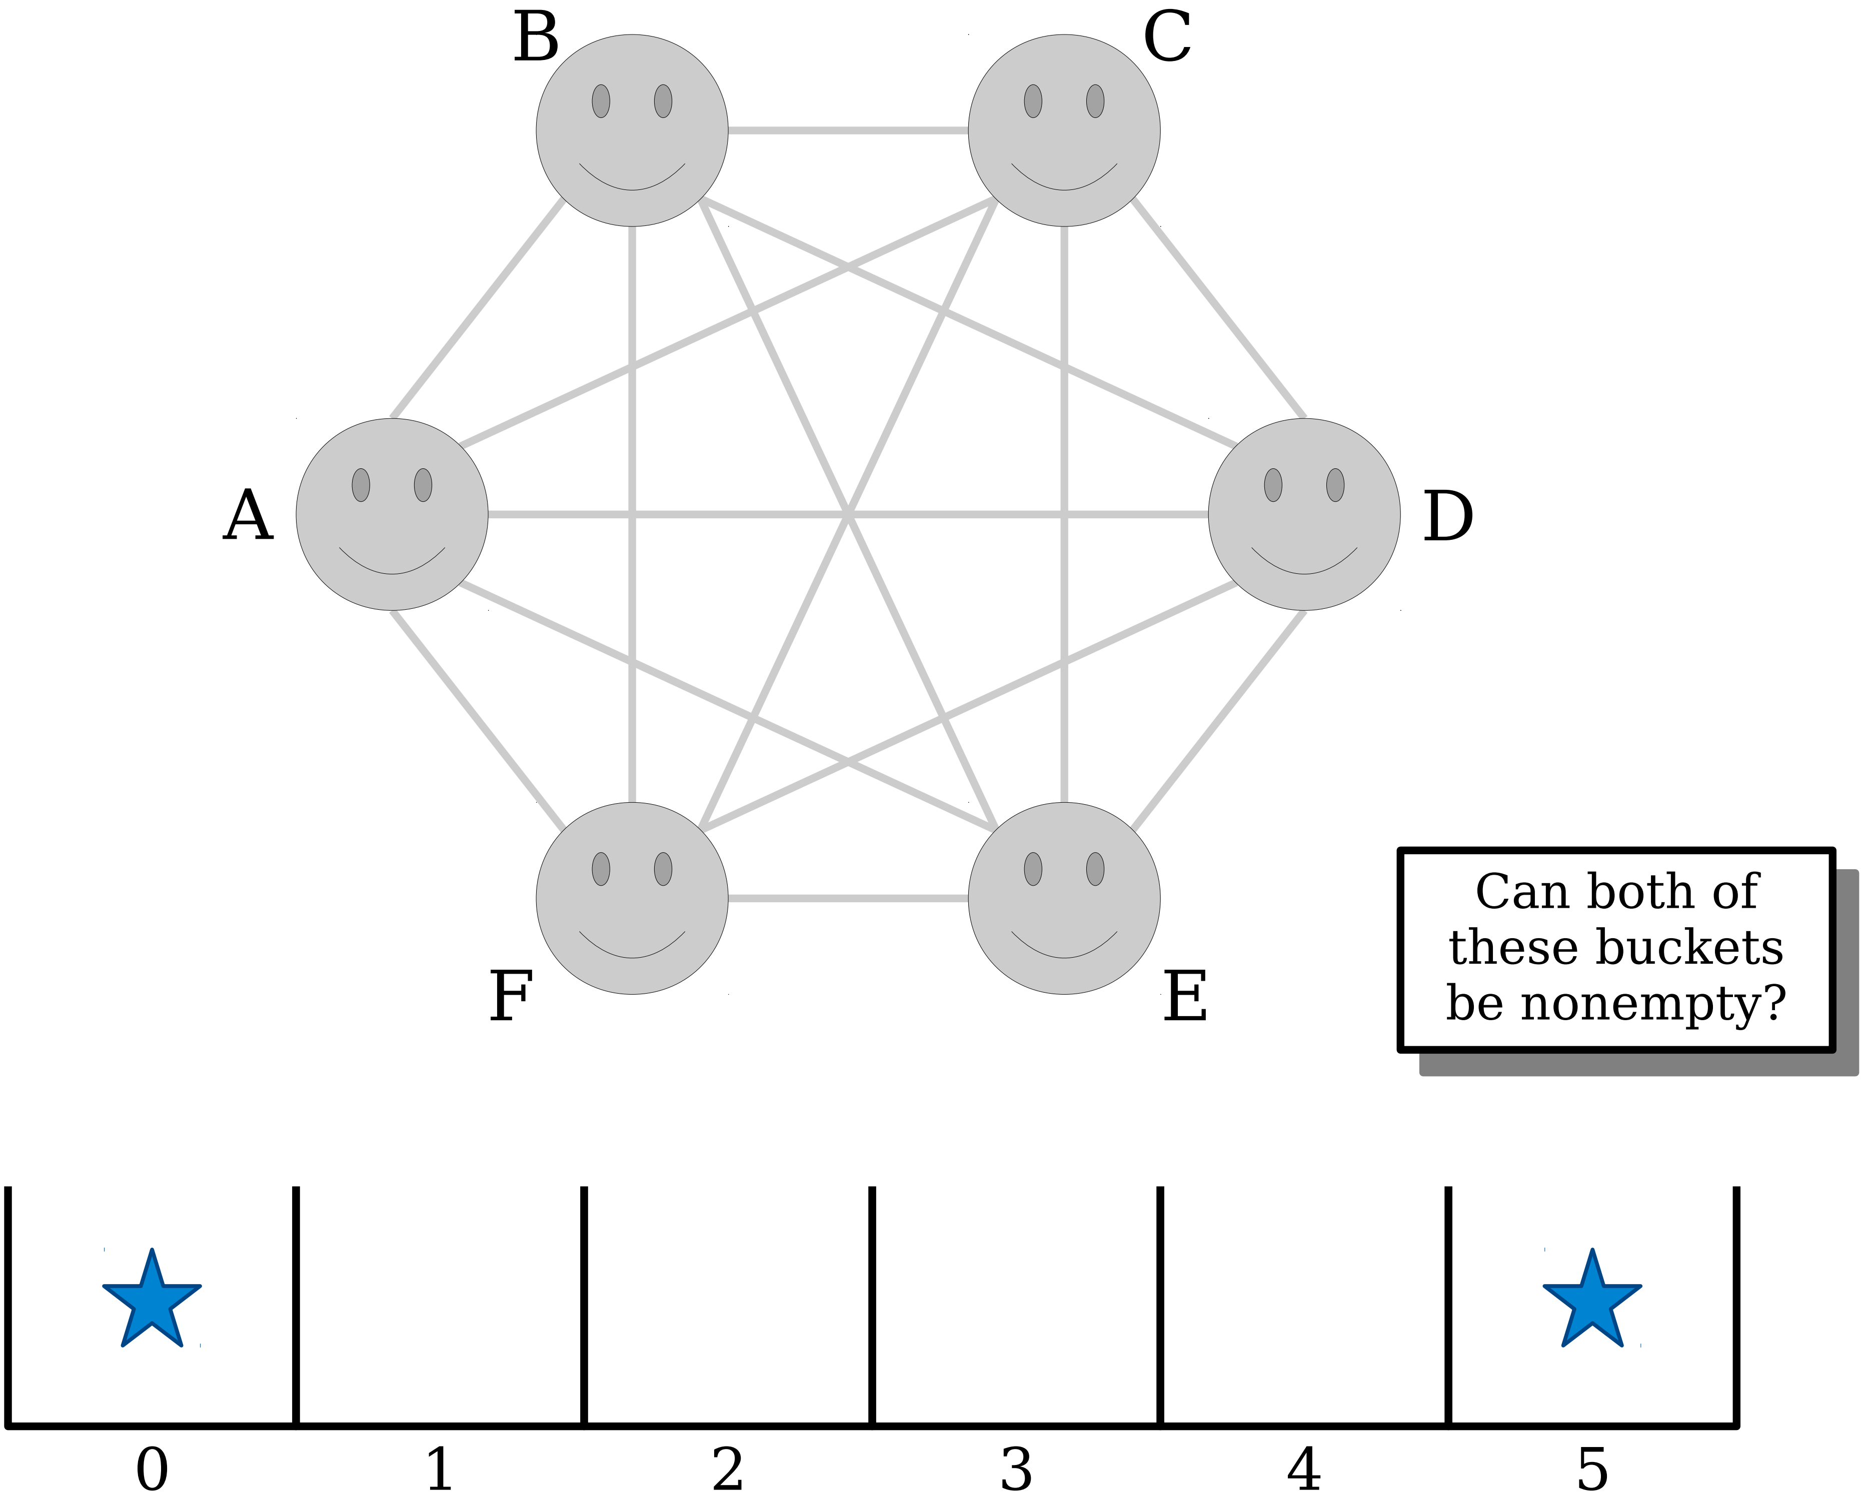
B
C
A
D
F
E
Can both of these buckets be nonempty?
0
1
2
3
4
5

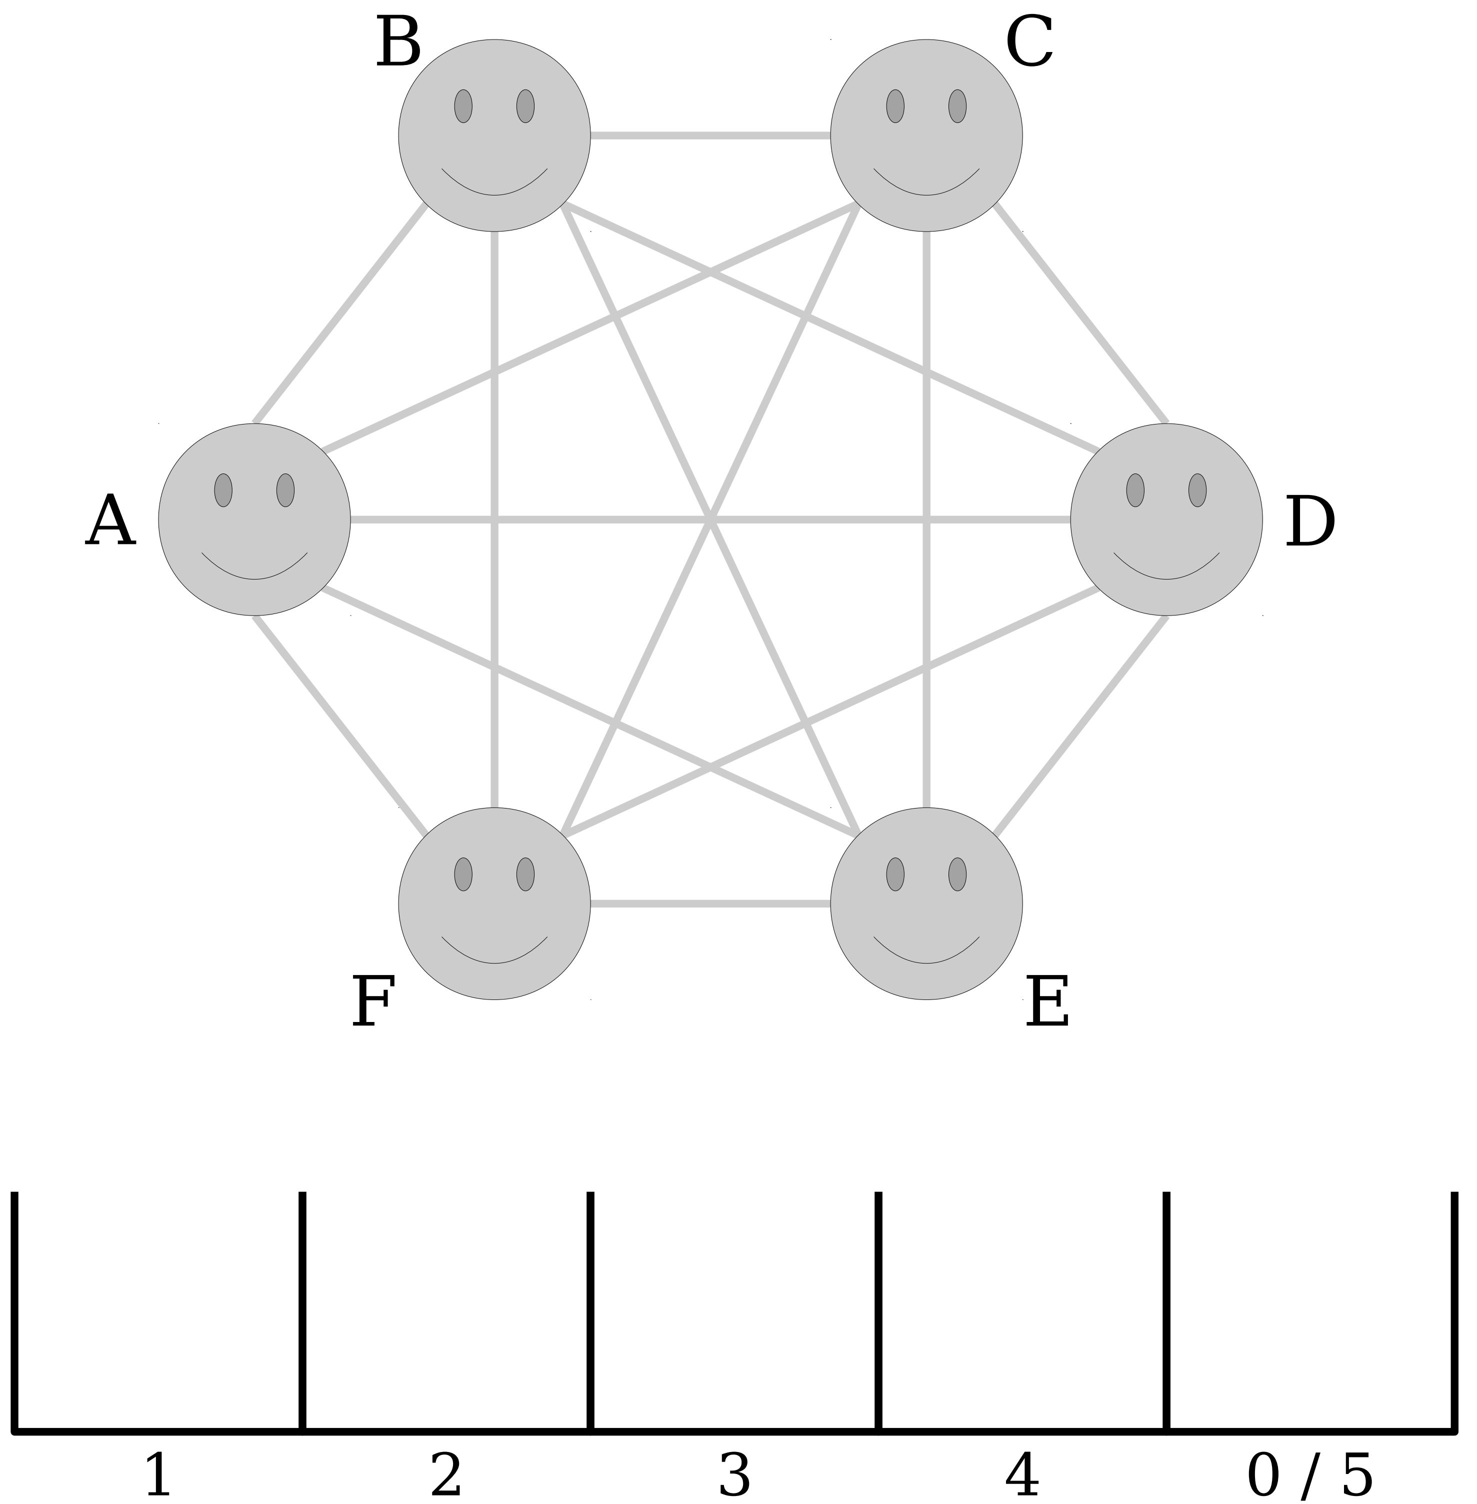
B
C
A
D
F
E
1
2
3
4
0 / 5

Theorem: In any graph with at least two nodes, there are at least two nodes of the same degree.

Proof 1: Let G be a graph with $n \geq 2$ nodes. There are n possible choices for the degrees of nodes in G, namely, 0, 1, 2, …, and $n - 1$.

We claim that G cannot simultaneously have a node u of degree 0 and a node v of degree $n - 1$: if there were such nodes, then node u would be adjacent to no other nodes and node v would be adjacent to all other nodes, including u. (Note that u and v must be different nodes, since v has degree at least 1 and u has degree 0.)

We therefore see that the possible options for degrees of nodes in G are either drawn from 0, 1, …, $n - 2$ or from 1, 2, …, $n - 1$. In either case, there are n nodes and $n - 1$ possible degrees, so by the pigeonhole principle two nodes in G must have the same degree. ■

Theorem: In any graph with at least two nodes, there are at least two nodes of the same degree.

Proof 2: Assume for the sake of contradiction that there is a graph G with $n \geq 2$ nodes where no two nodes have the same degree. There are n possible choices for the degrees of nodes in G, namely 0, 1, 2, …, $n - 1$, so this means that G must have exactly one node of each degree. However, this means that G has a node of degree 0 and a node of degree $n - 1$. (These can't be the same node, since $n \geq 2$.) This first node is adjacent to no other nodes, but this second node is adjacent to every other node, which is impossible.

We have reached a contradiction, so our assumption must have been wrong. Thus if G is a graph with at least two nodes, G must have at least two nodes of the same degree. ■

The Generalized Pigeonhole Principle

A More General Version

- The ***generalized pigeonhole principle*** says that if you distribute m objects into n bins, then
  - some bin will have at least $\lceil m/n \rceil$ objects in it, and
  - some bin will have at most $\lfloor m/n \rfloor$ objects in it.

$\lceil m/n \rceil$ means "m/n, rounded up."
$\lfloor m/n \rfloor$ means "m/n, rounded down."

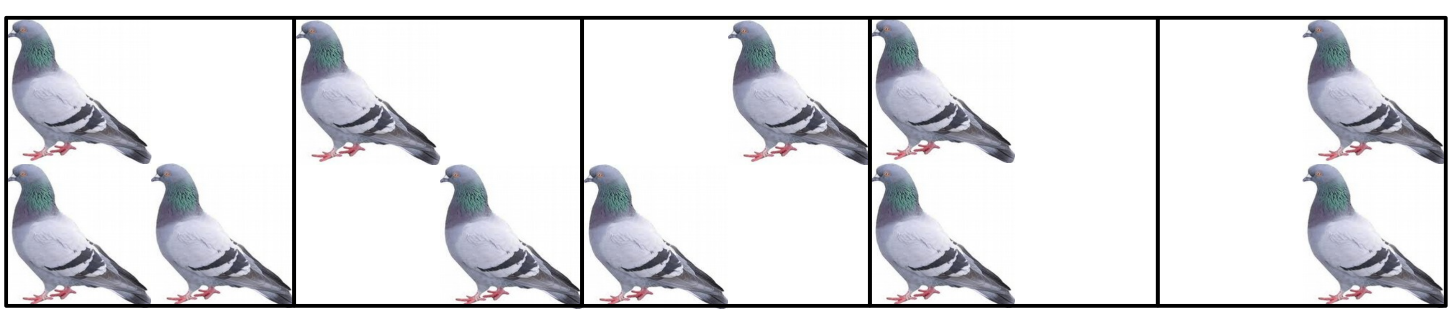

$m = 11$
$n = 5$

$\lceil m / n \rceil = 3$
$\lfloor m / n \rfloor = 2$

Theorem: If m objects are distributed into $n > 0$ bins, then some bin will contain at least $\lceil m/n \rceil$ objects.

Proof: We will prove that if m objects are distributed into n bins, then some bin contains at least m/n objects. Since the number of objects in each bin is an integer, this will prove that some bin must contain at least $\lceil m/n \rceil$ objects.

To do this, we proceed by contradiction. Suppose that, for some m and n, there is a way to distribute m objects into n bins such that each bin contains fewer than m/n objects.

Number the bins 1, 2, 3, …, n and let x_i denote the number of objects in bin i. Since there are m objects in total, we know that

$$m = x_1 + x_2 + \dots + x_n.$$

Since each bin contains fewer than m/n objects, we see that $x_i < m/n$ for each i. Therefore, we have that

$$\begin{aligned} m &= x_1 + x_2 + \dots + x_n \\ &< m/n + m/n + \dots + m/n \quad (n \text{ times}) \\ &= m. \end{aligned}$$

But this means that $m < m$, which is impossible. We have reached a contradiction, so our initial assumption must have been wrong. Therefore, if m objects are distributed into n bins, some bin must contain at least $\lceil m/n \rceil$ objects. ■

An Application: Friends and Strangers

Friends and Strangers

- Suppose you have a party of six people. Each pair of people are either friends (they know each other) or strangers (they do not).
- ***Theorem:*** Any such party must have a group of three mutual friends (three people who all know one another) or three mutual strangers (three people where no one knows anyone else).

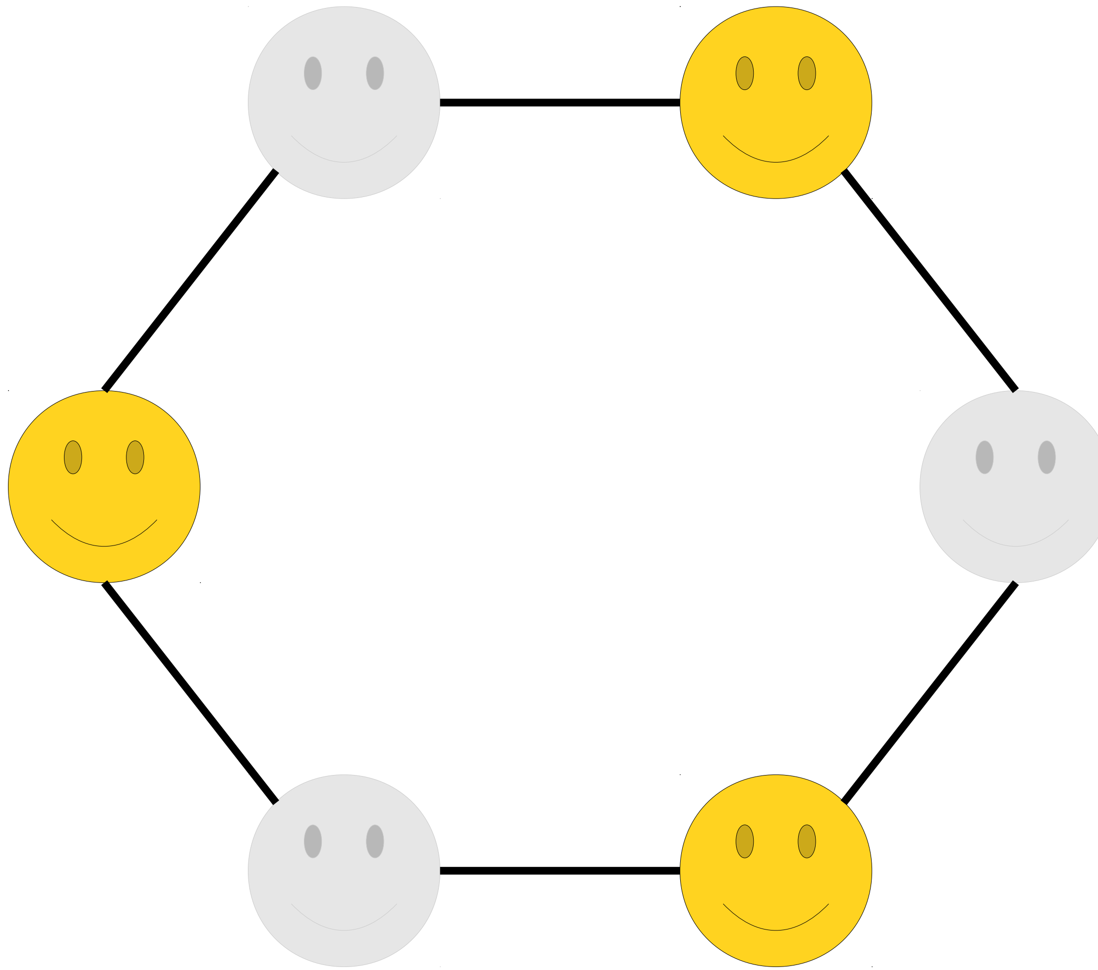

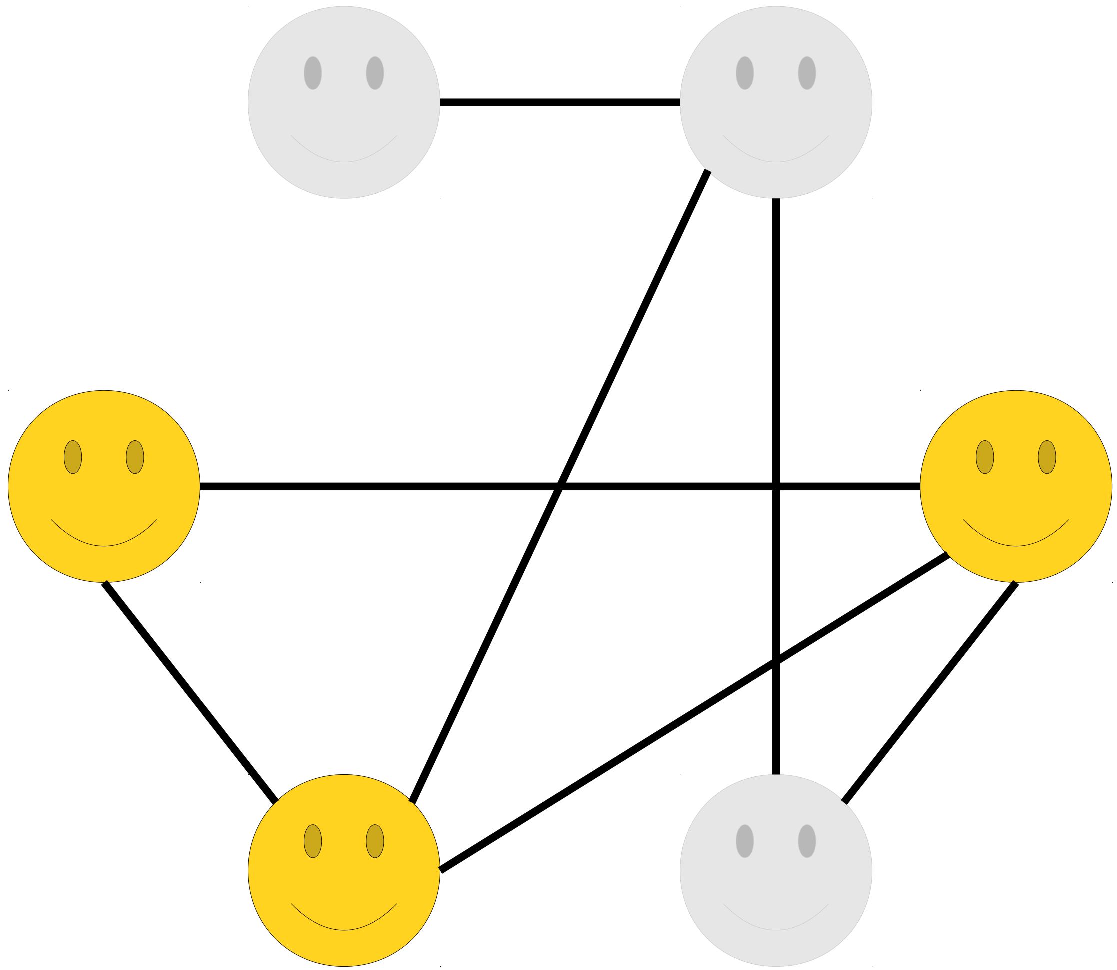

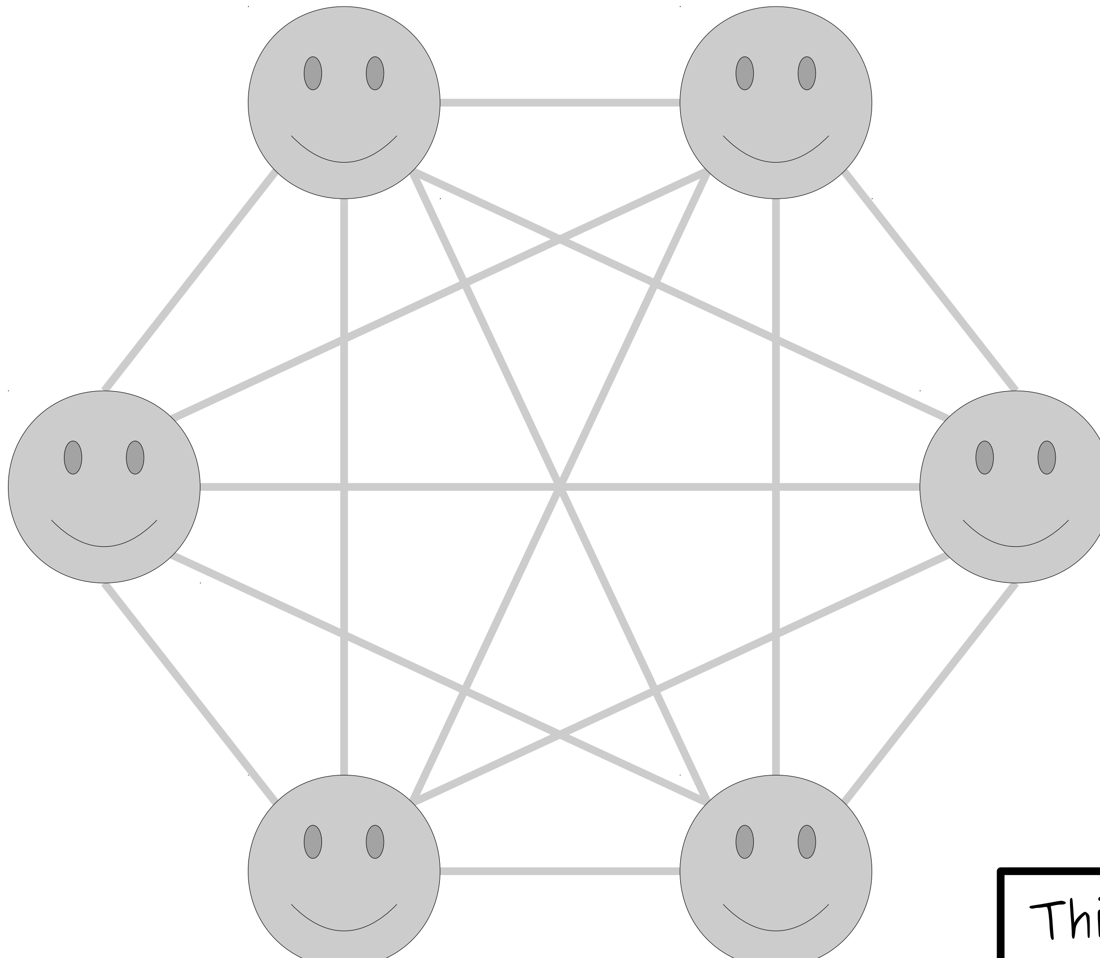

This graph is called a *6-clique*, by the way.

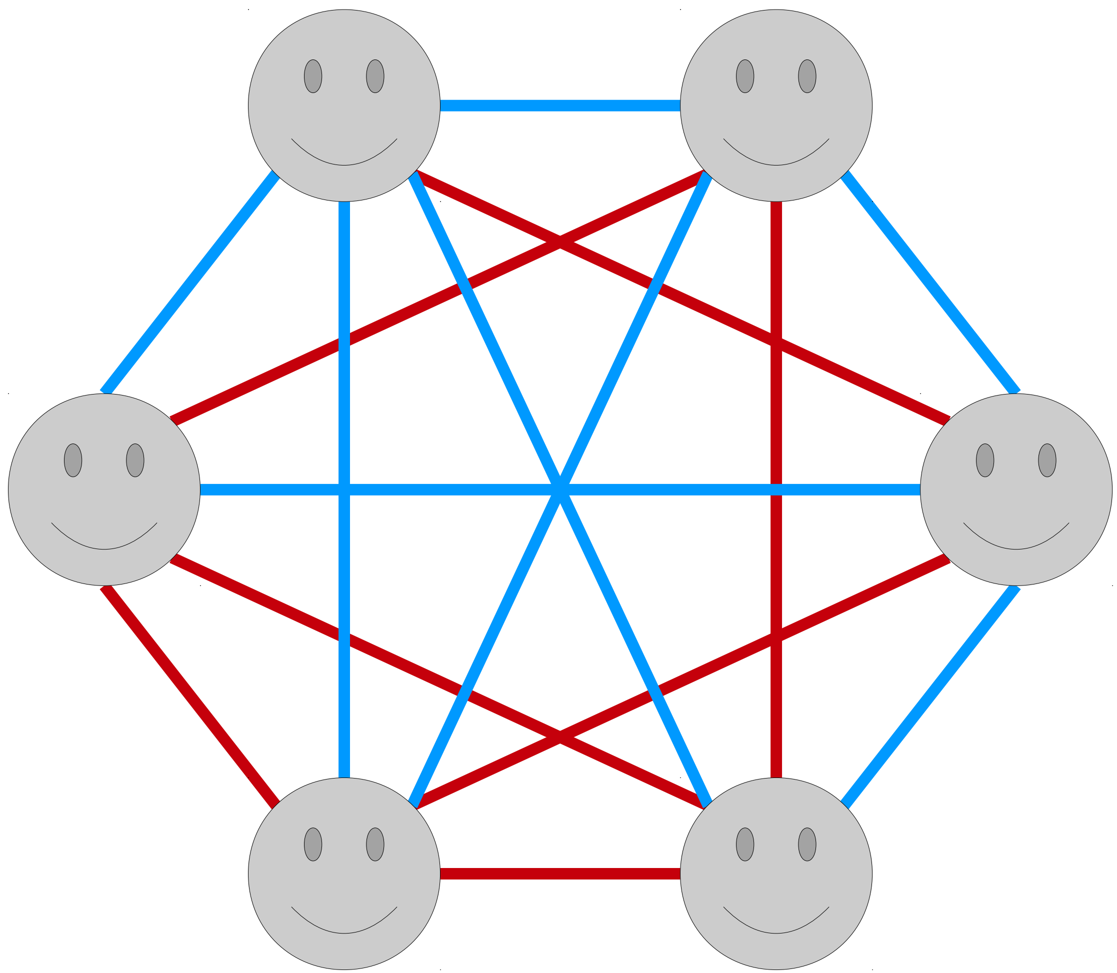

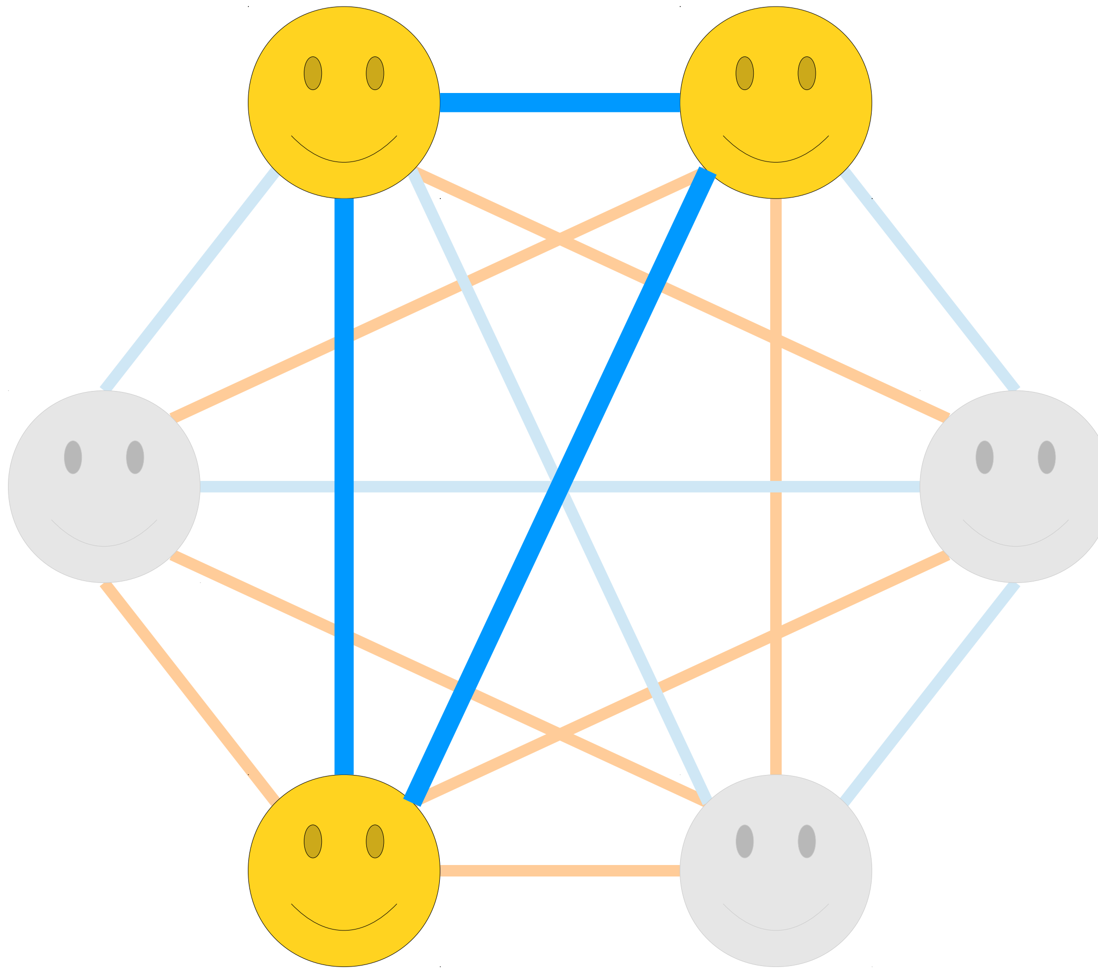

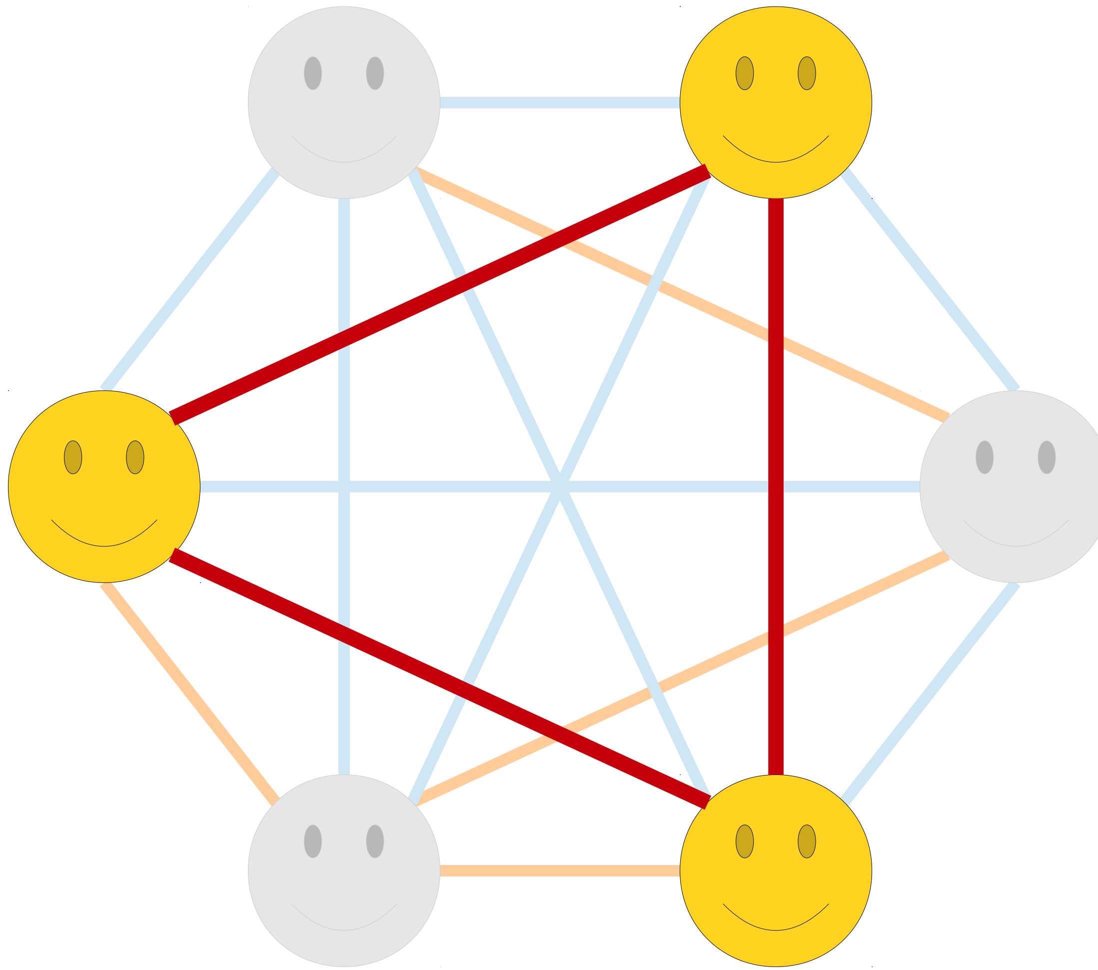

Friends and Strangers Restated

- From a graph-theoretic perspective, the Theorem on Friends and Strangers can be restated as follows:
- ***Theorem:*** Consider a 6-clique where every edge is colored red or blue. The the graph contains a red triangle, a blue triangle, or both.
- How can we prove this?

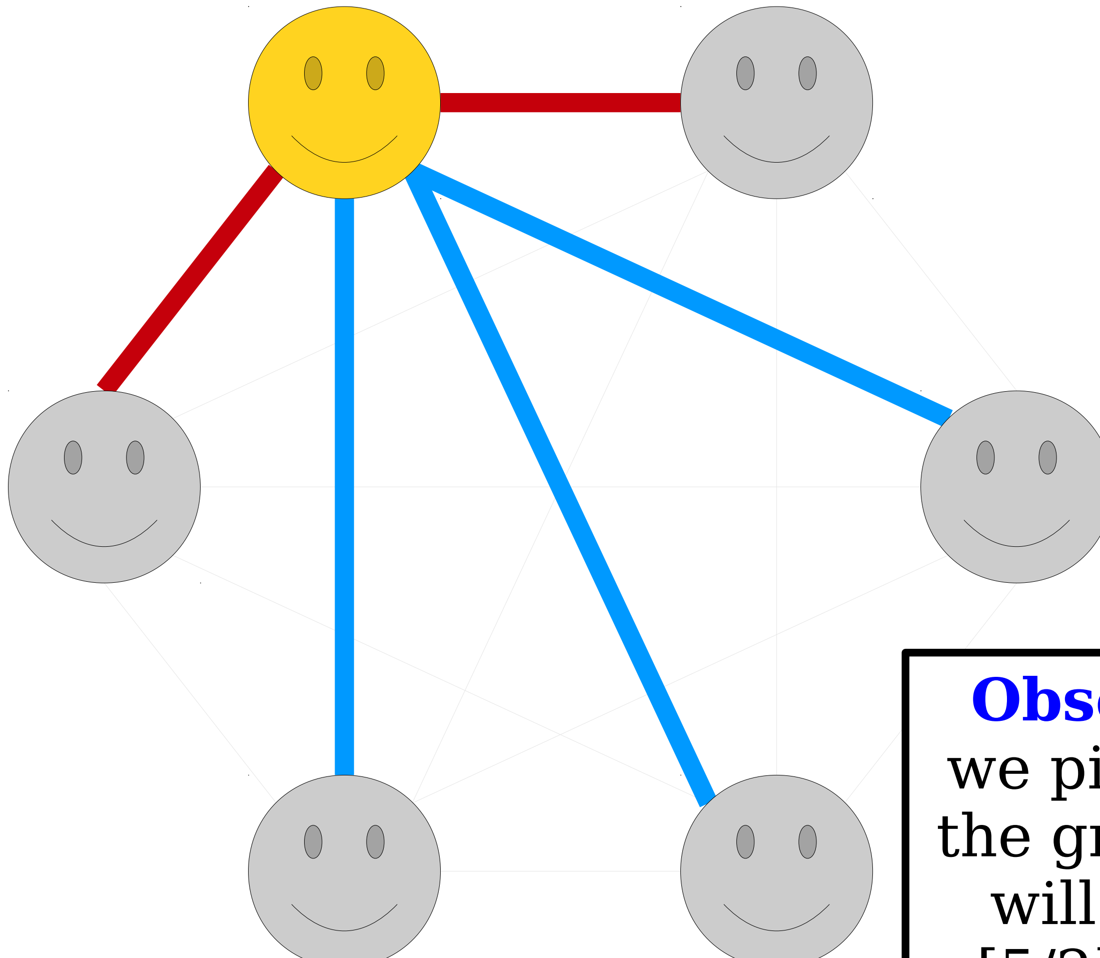

Observation 1: If we pick any node in the graph, that node will have at least $\lceil 5/2 \rceil = 3$ edges of the same color incident to it.

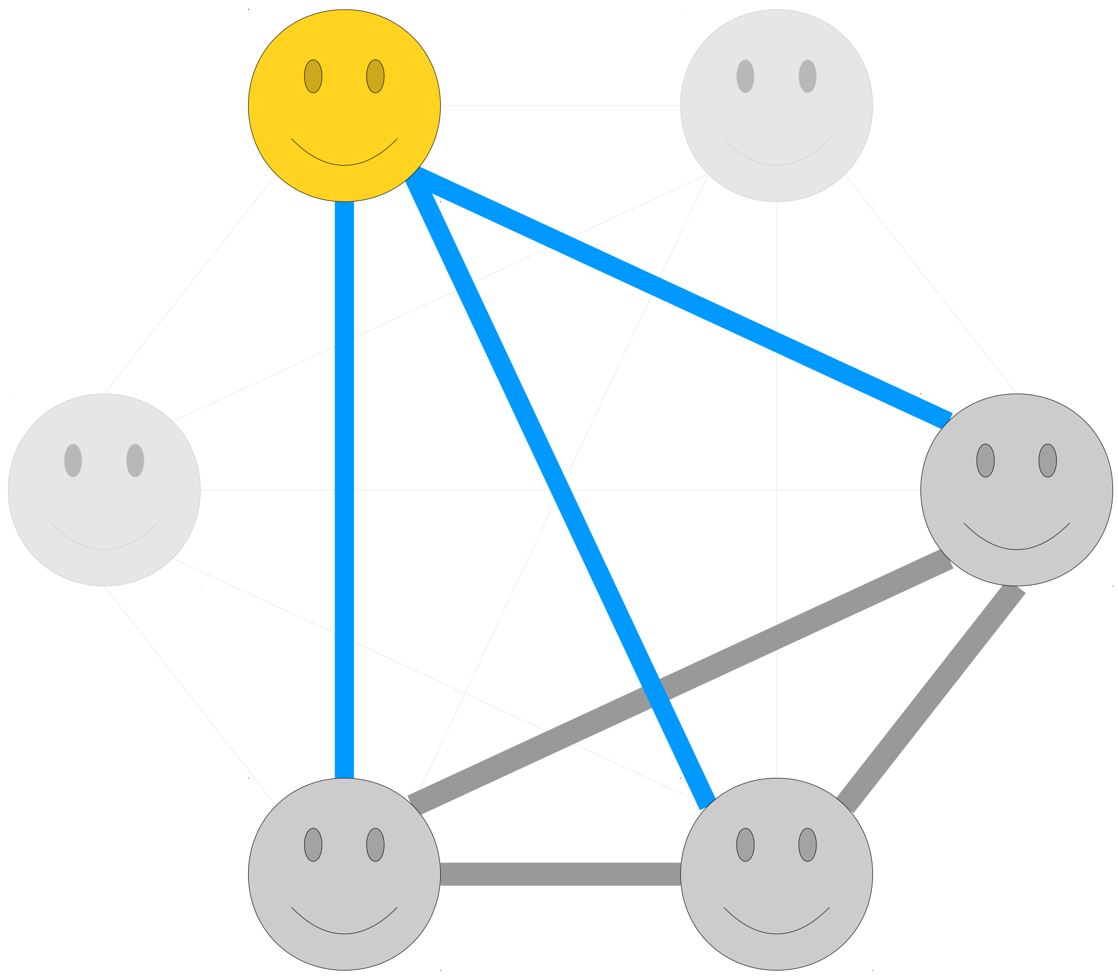

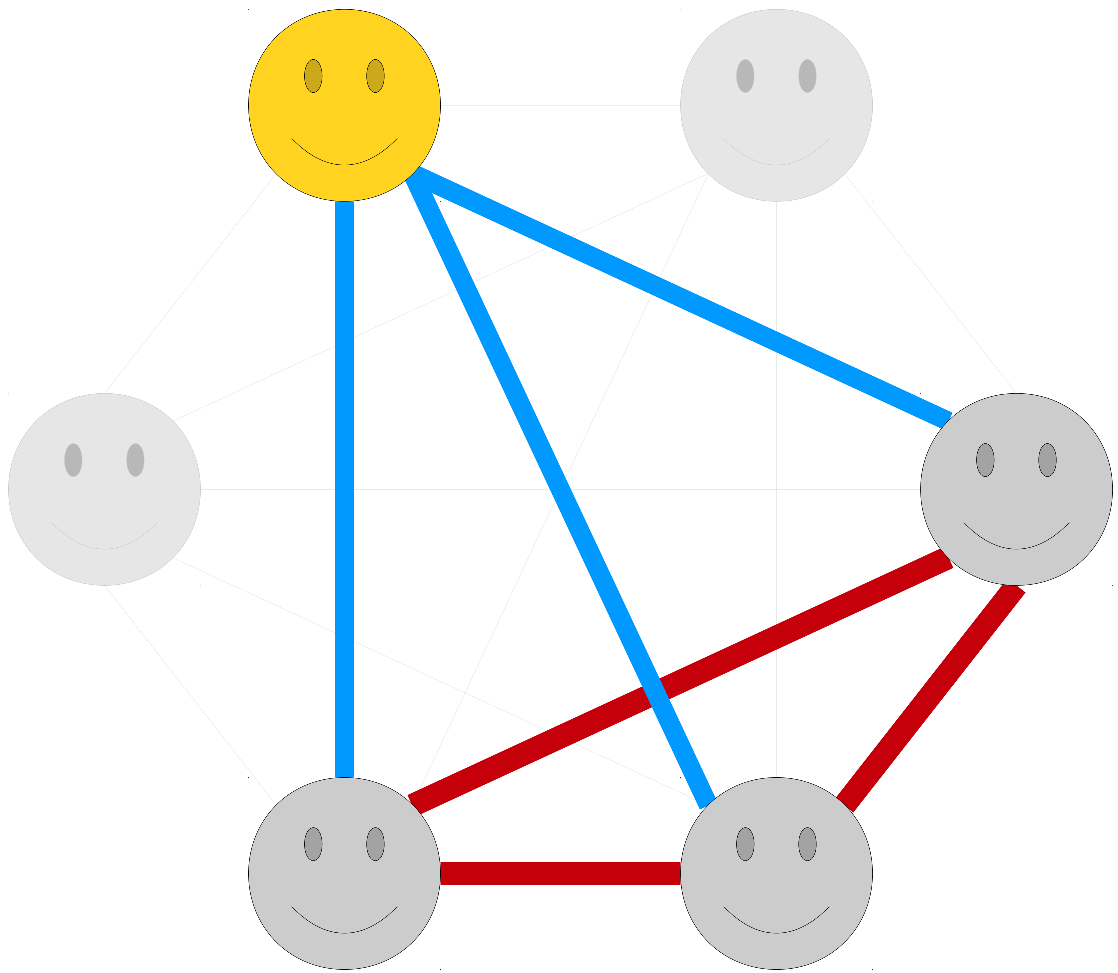

Theorem: Consider a 6-clique in which every edge is colored either red or blue. Then there must be a triangle of red edges, a triangle of blue edges, or both.

Proof: Color the edges of the 6-clique either red or blue arbitrarily. Let x be any node in the 6-clique. It is incident to five edges and there are two possible colors for those edges. Therefore, by the generalized pigeonhole principle, at least $\lceil 5/2 \rceil = 3$ of those edges must be the same color. Call that color c_1 and let the other color be c_2.

Let r, s, and t be three of the nodes connected to node x by an edge of color c_1. If any of the edges $\{r, s\}$, $\{r, t\}$, or $\{s, t\}$ are of color c_1, then one of those edges plus the two edges connecting back to node x form a triangle of color c_1. Otherwise, all three of those edges are of color c_2, and they form a triangle of color c_2. Overall, this gives a red triangle or a blue triangle, as required. ■

What This Means

- The proof we just did was along the following lines:

  If you choose a sufficiently large object, you are guaranteed to find a large subobject of type A or a large subobject of type B.

- Intuitively, it's not possible to find gigantic objects that have absolutely no patterns or structure in them – there is no way to avoid having some interesting structure.
- There are numerous theorems of this sort. The mathematical field of ***Ramsey theory*** explicitly studies problems of this type.

Time-Out for Announcements!

APPLY TO JOIN
BLACK IN CS CORE!

Applications are live

DEADLINE: Wednesday, May 3rd at 11:59 PM

As we prepare for our second year as a student group, you can be a crucial part of forming the vision and impact of Black in CS. As a core member, part of your duty is to continue to build the foundations for the student group and help cultivate a strong Black in CS community. Feel free to self-nominate or nominate others for positions in Black in CS. Candidates for Co-Presidents and Financial Officer positions can expect elections and for General Core Member positions, an application process.

Please contact a co-president, either Lindsey Redd (**l1redd@stanford.edu**) or Ekua Awotwi (**eawotwi@stanford.edu**), if you have any questions!

The Grace Hopper Celebration is the world's largest gathering of women in computing. It is designed to bring the interests of women in computing to the forefront.

Do you want to attend? The CS department can send up to 30 CS women to Grace Hopper this year. We'll pay for registration, travel, and hotel costs (within reason).

The Grace Hopper Celebration of Women in Computing
Dates: October 4-6, 2017
Location: Orlando, Florida
http://ghc.anitaborg.org

If you're interested, please fill out ***this link*** by ***Sunday, April 30***.

Midterm Exam

- The first midterm exam is next ***Tuesday, May 2nd***, from ***7:00PM – 10:00PM***. Locations are divvied up by last (family) name:
  - `Abb – Niu`: Go to Hewlett 200.
  - `Nor – Vas`: Go to Hewlett 201.
  - `Vil – Yim`: Go to Hewlett 102.
  - `You – Zuc`: Go to Hewlett 103.
- You’re responsible for Lectures 00 – 05 and topics covered in PS1 – PS2. Later lectures and problem sets won’t be tested.
- The exam is closed-book, closed-computer, and limited-note. You can bring a double-sided, 8.5” × 11” sheet of notes with you to the exam, decorated however you’d like.
- Students with OAE accommodations: you should have heard back from us with alternate exam logistics earlier today. Let us know if that’s not the case!

Midterm Practice

- To help you prep for the midterm, we've posted

  two practice midterm exams and

  three sets of extra practice problems,

  with solutions, up on the course website.
- Need more practice? Let us know what we can do to help out!
- Please feel free to ask questions on Piazza over the weekend. We're happy to help!

Problem Sets

- Problem Set Three was due at the start of class today.
- Problem Set Four goes out today.
  - Checkpoint problems are due on Monday right before class.
  - Remaining problems due on Friday right before class.
- Play around with infinite cardinalities, the limits of set theory, nifty properties of graphs, and some cool applications of the material!

Your Questions

“What gets you up in the morning?”
“Who inspires you?”

You do! I never cease to be amazed by what students here can do and what you'll accomplish. Like, really.

"Do you think political diversity is important? Do you think Stanford has enough political diversity?"

It's absolutely important. Lack of political diversity takes ideas and solution routes off the table.

Take a look at Salt Lake City's program for helping the homeless by giving them homes (which is now a national model) or San Antonio's approach to <u>mental health services</u> (again, a national model). These approaches combine ideas across the political spectrum and show a lot of promise.

Back to CS103!

Another View of Pigeonholing

- The pigeonhole principle is a result that, broadly speaking, follows this template:

  ***m* objects cannot be distributed into *n* bins without property *X* being true.**

- What other sorts of properties can we say about how objects get distributed?

Observation: The number of boxes containing an odd number of pigeons seems to always be even!

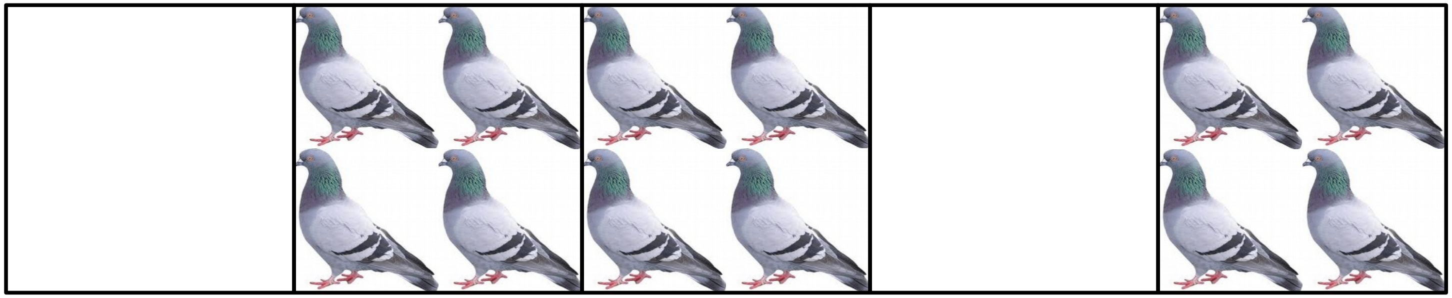

m = 12 pigeons
n = 5 boxes

Observation: Now the number of boxes containing an odd number of pigeons seems to always be odd!

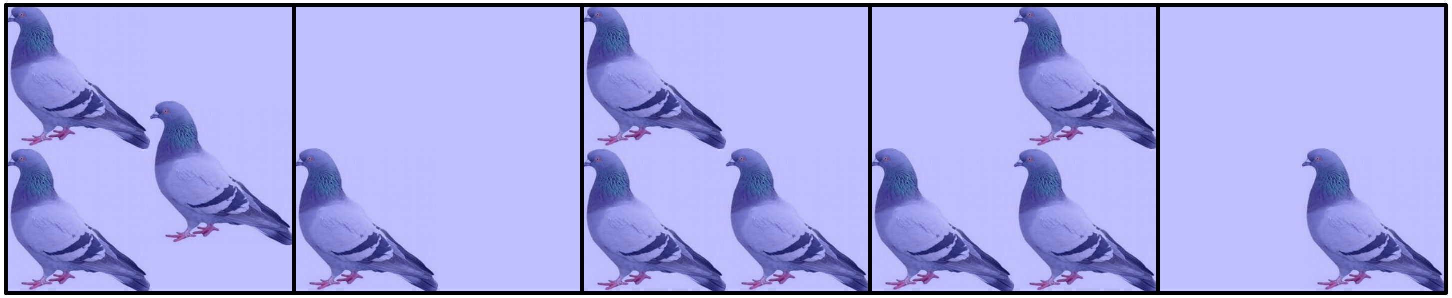

m = 11 pigeons
n = 5 boxes

Theorem: Suppose m objects are distributed into some number of bins. Let k be the number of bins containing an odd number of objects. Then m and k have the same parity.

Proof: Let m be an arbitrary natural number and suppose that m objects are distributed across some number of bins. Let k be the number of bins with an odd number of objects. We will prove that k has the same parity as m.

Denote the numbers of objects in each of the even-size bins as $2r_1$, $2r_2$, …, and $2r_h$ and the numbers of objects in the odd-size bins as $2s_1+1$, $2s_2+1$, …, and $2s_k+1$. Then, since each object is placed into some bin, we have that

$$m = (2r_1 + 2r_2 + \ldots + 2r_h) + ((2s_1 + 1) + (2s_2 + 1) + \ldots + (2s_k + 1)).$$

There are k copies of the +1 term in the second group, so we see

$$m = (2r_1 + 2r_2 + \ldots + 2r_h) + (2s_1 + 2s_2 + \ldots + 2s_k) + k.$$

Regrouping the terms to isolate k yields

$$m - 2(r_1 + r_2 + \ldots + r_h + s_1 + s_2 + \ldots + s_k) = k.$$

If m is even, then k is the difference of two even numbers, so k is even. Otherwise, m is odd. Then k is the difference of an odd number and an even number, so k is odd as well. In both cases, we see that k has the same parity as m, as required. ■

A Pretty Nifty Theorem:
Sperner's Lemma

- Begin with a triangle.

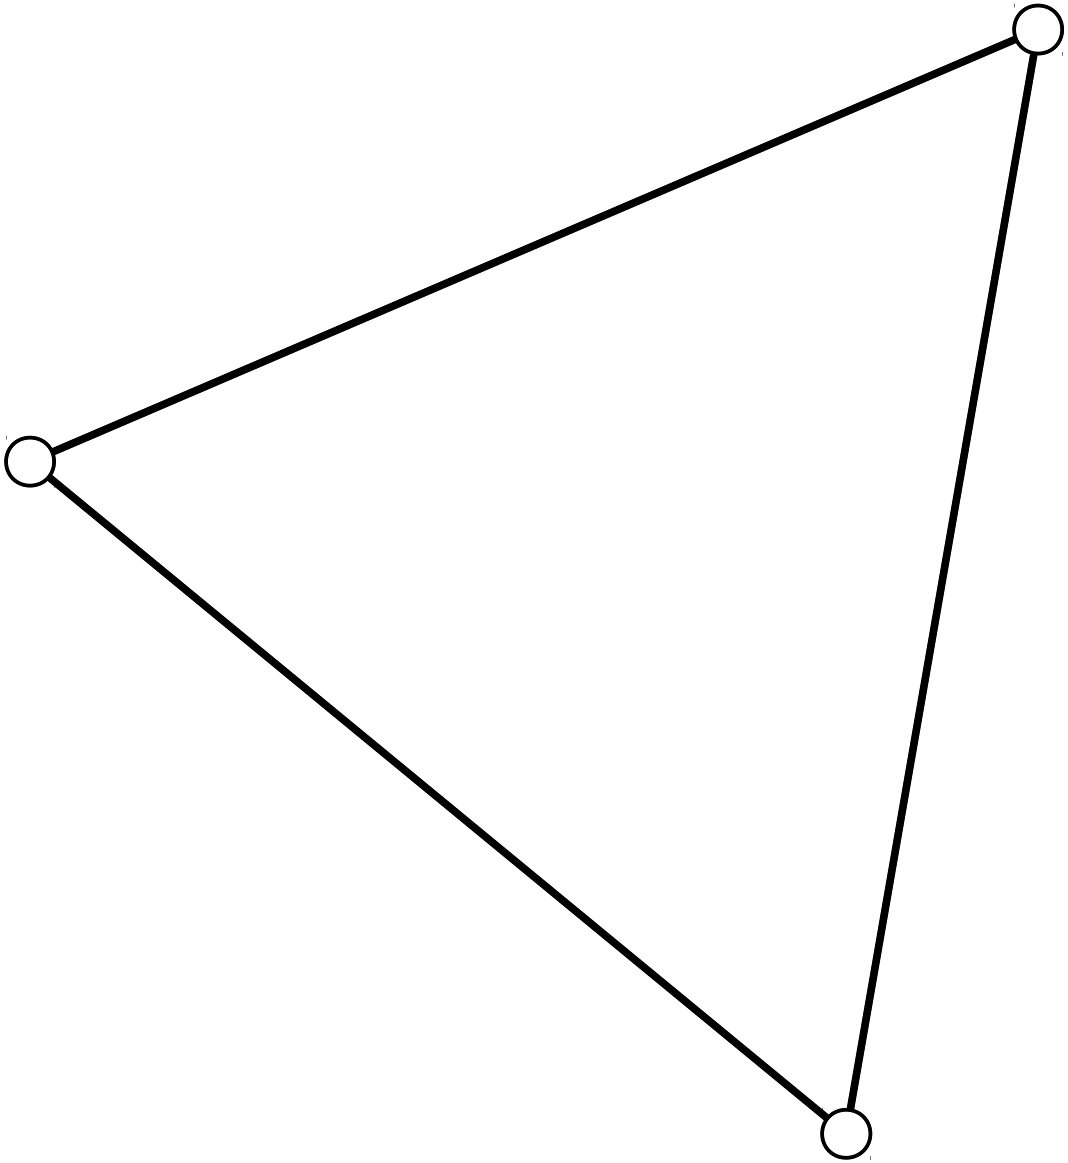

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)

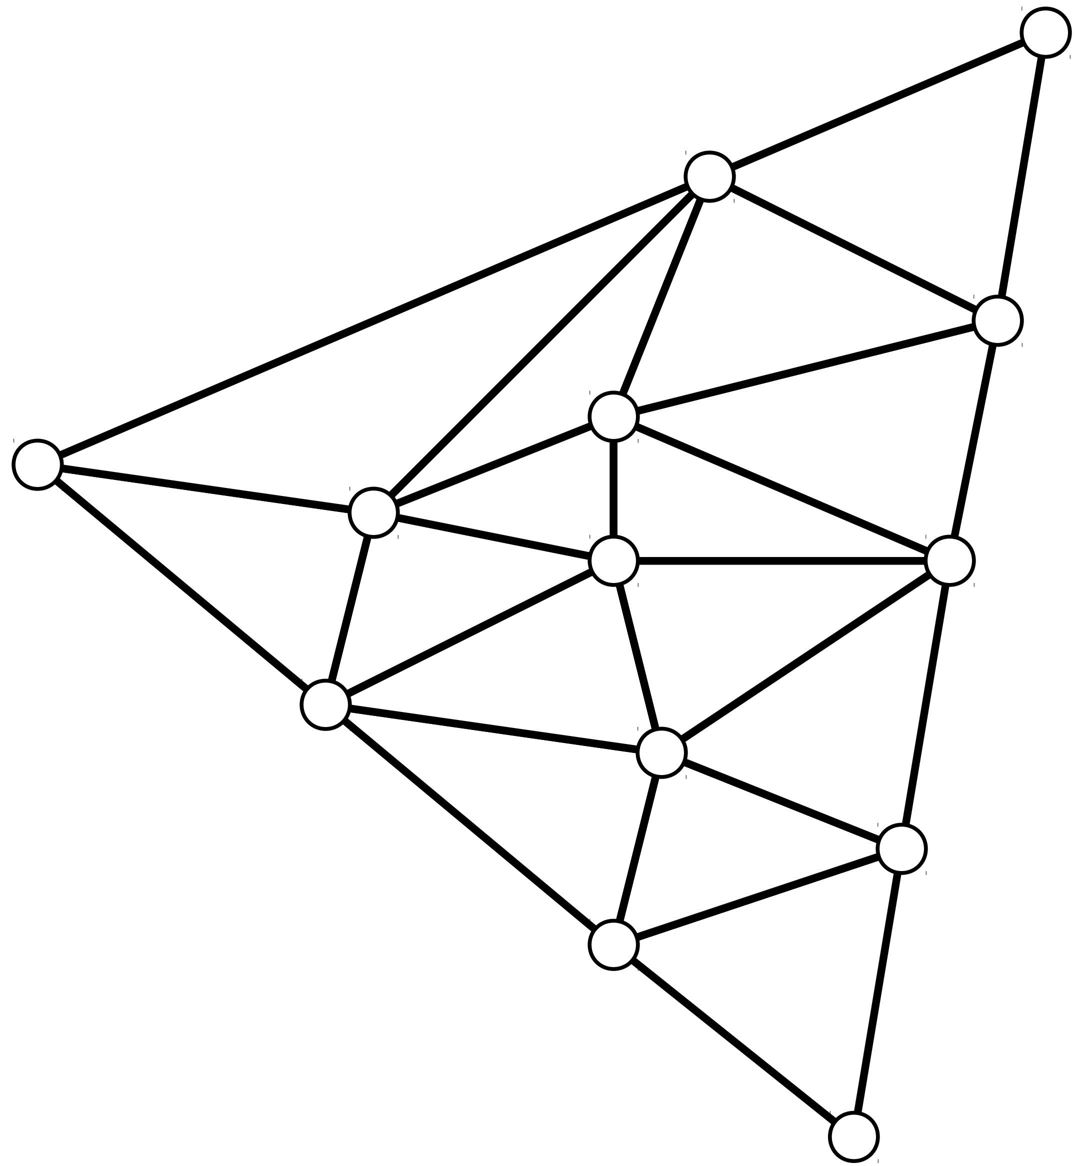

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.

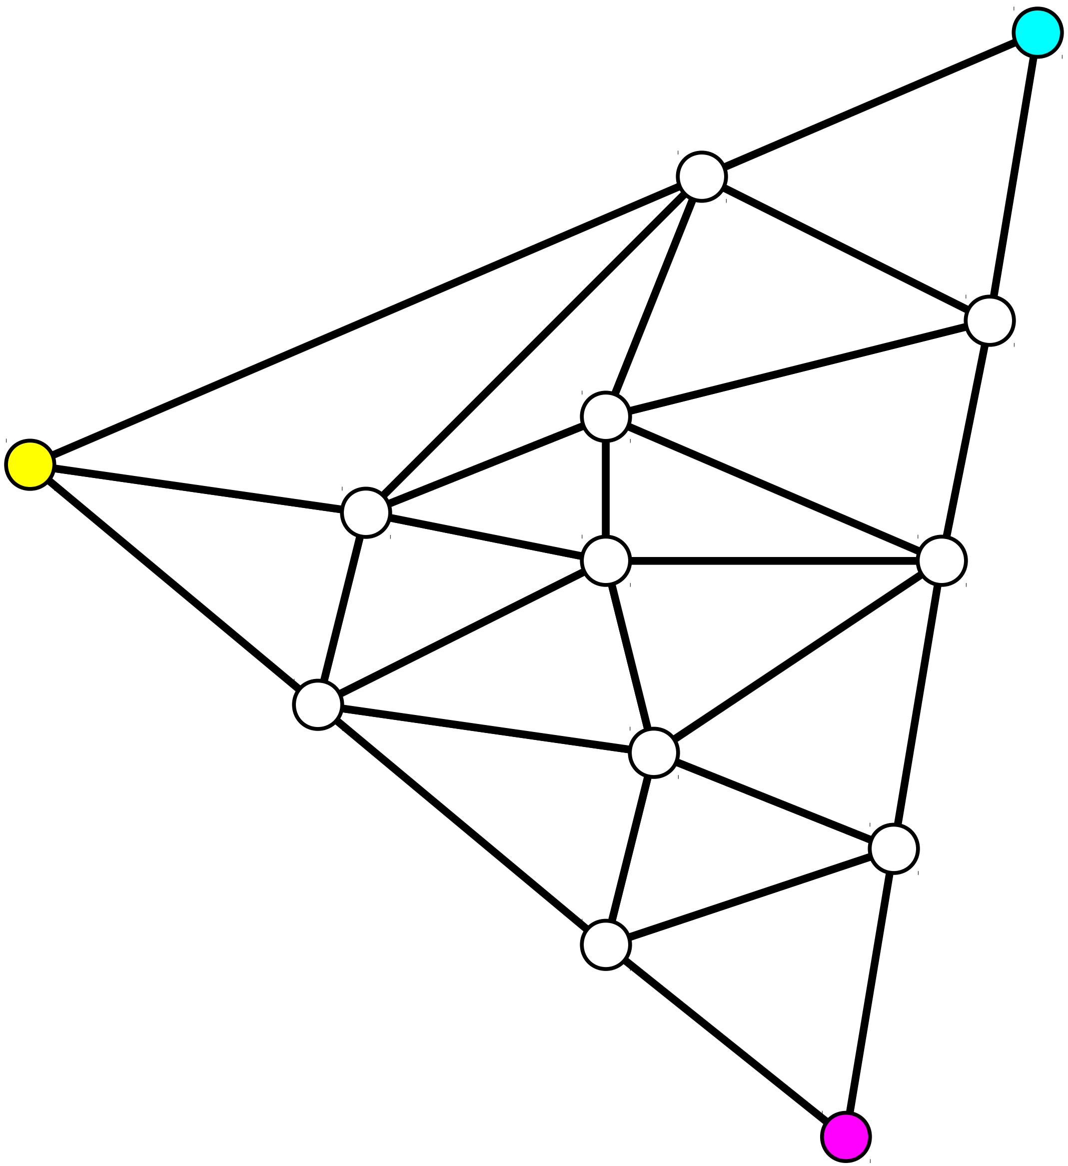

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.

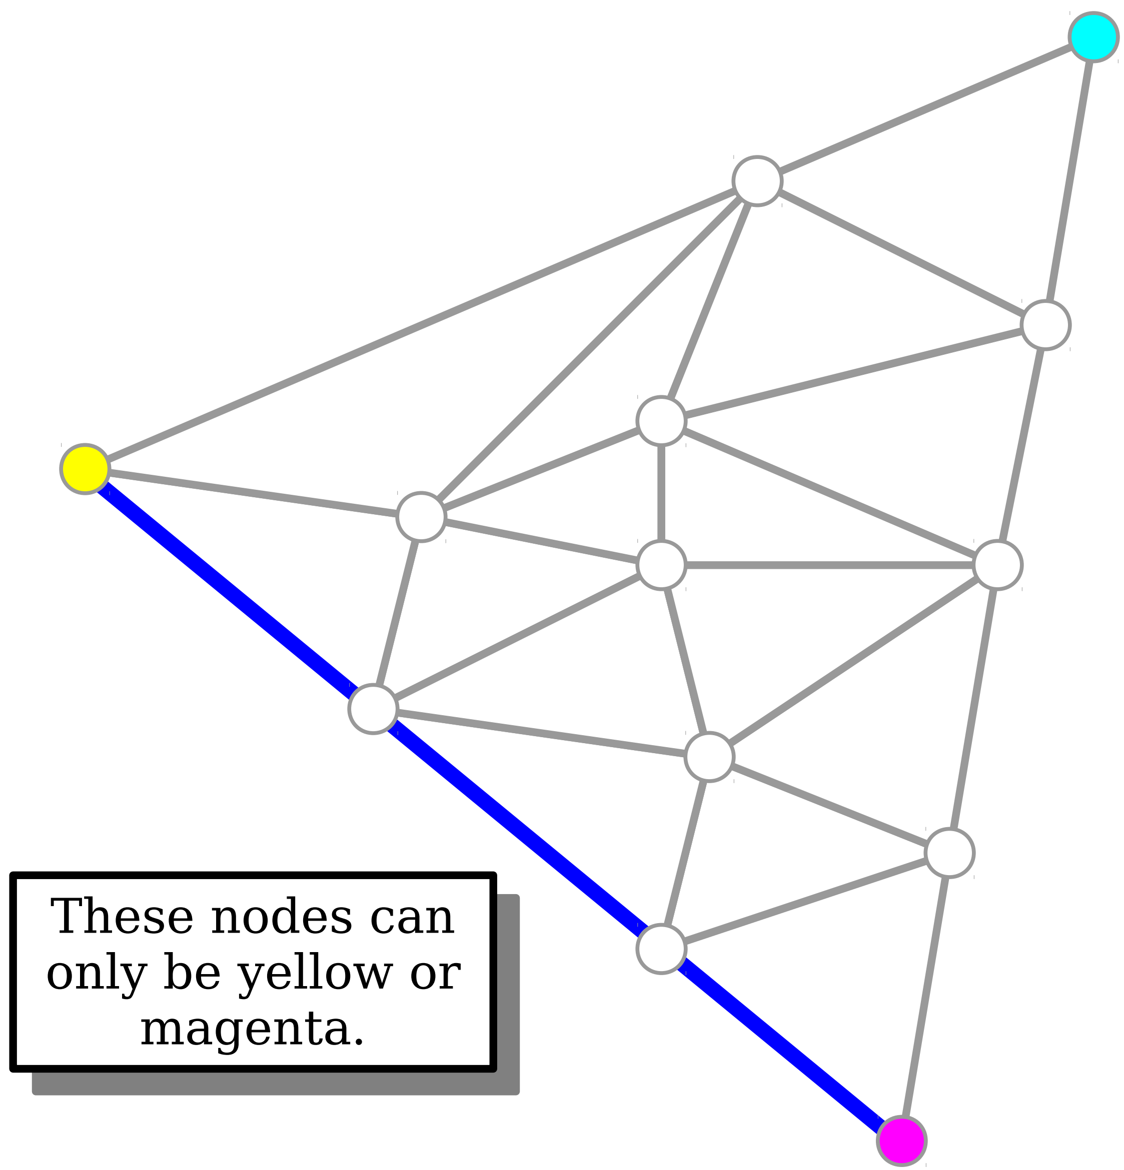

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.

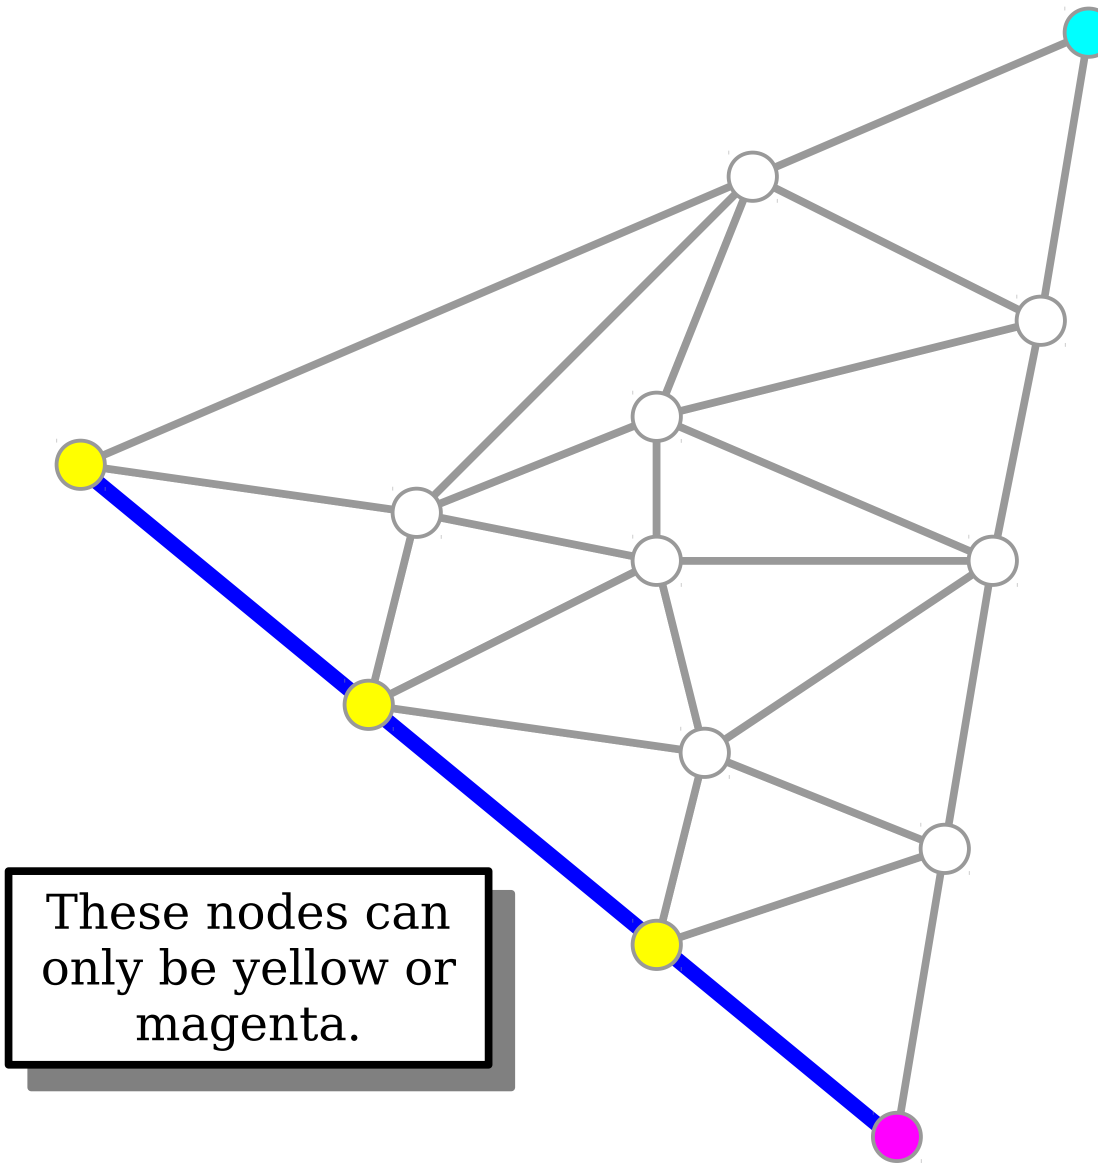

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.

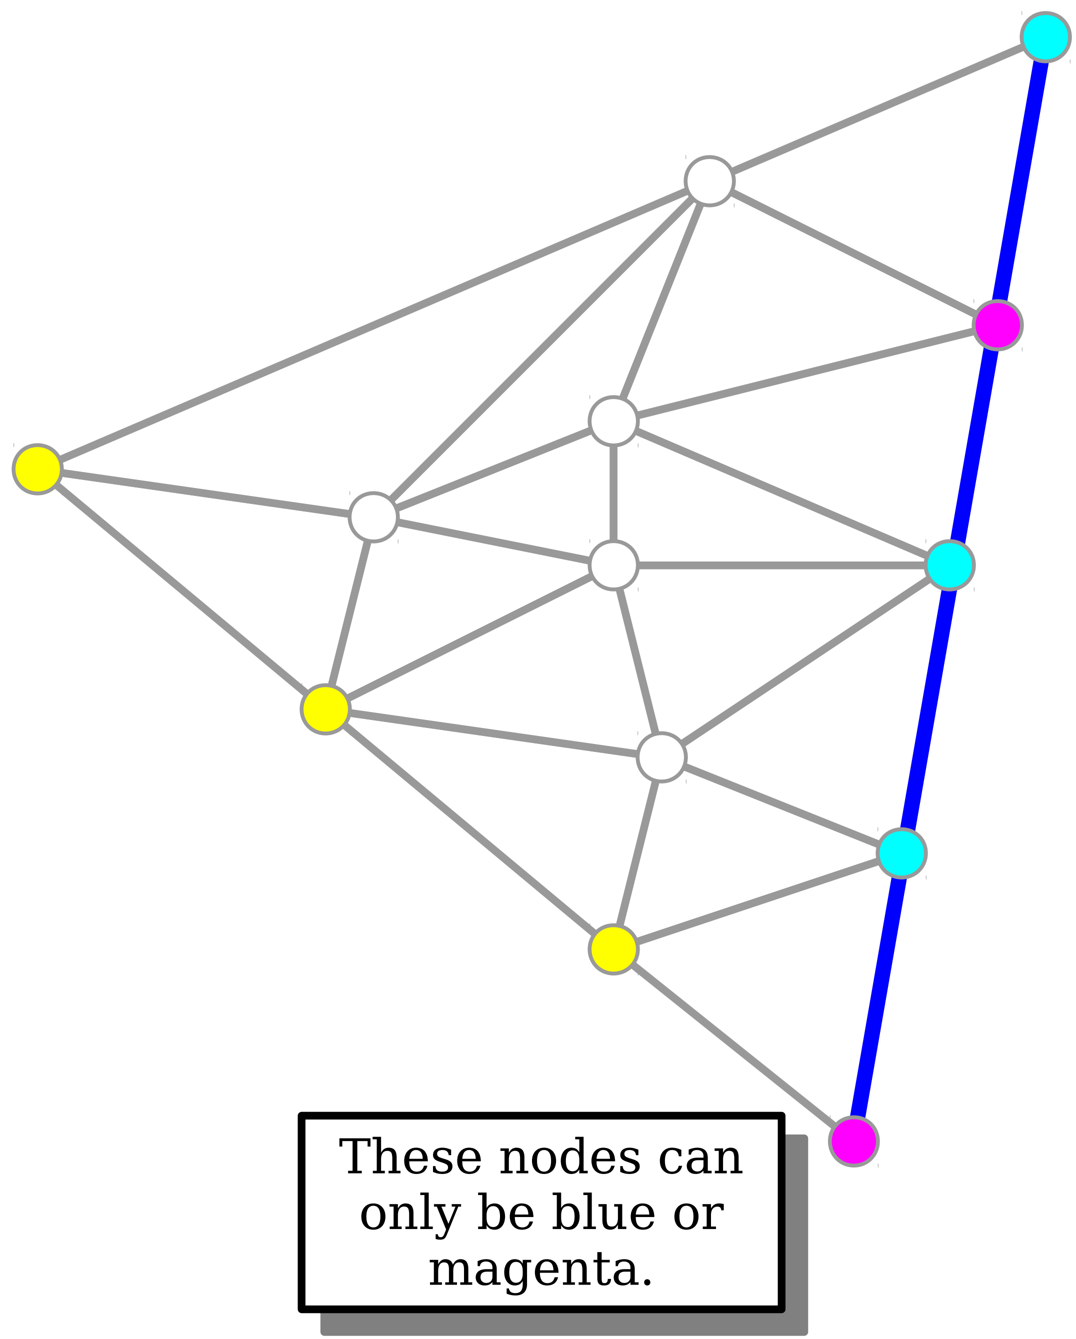

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.

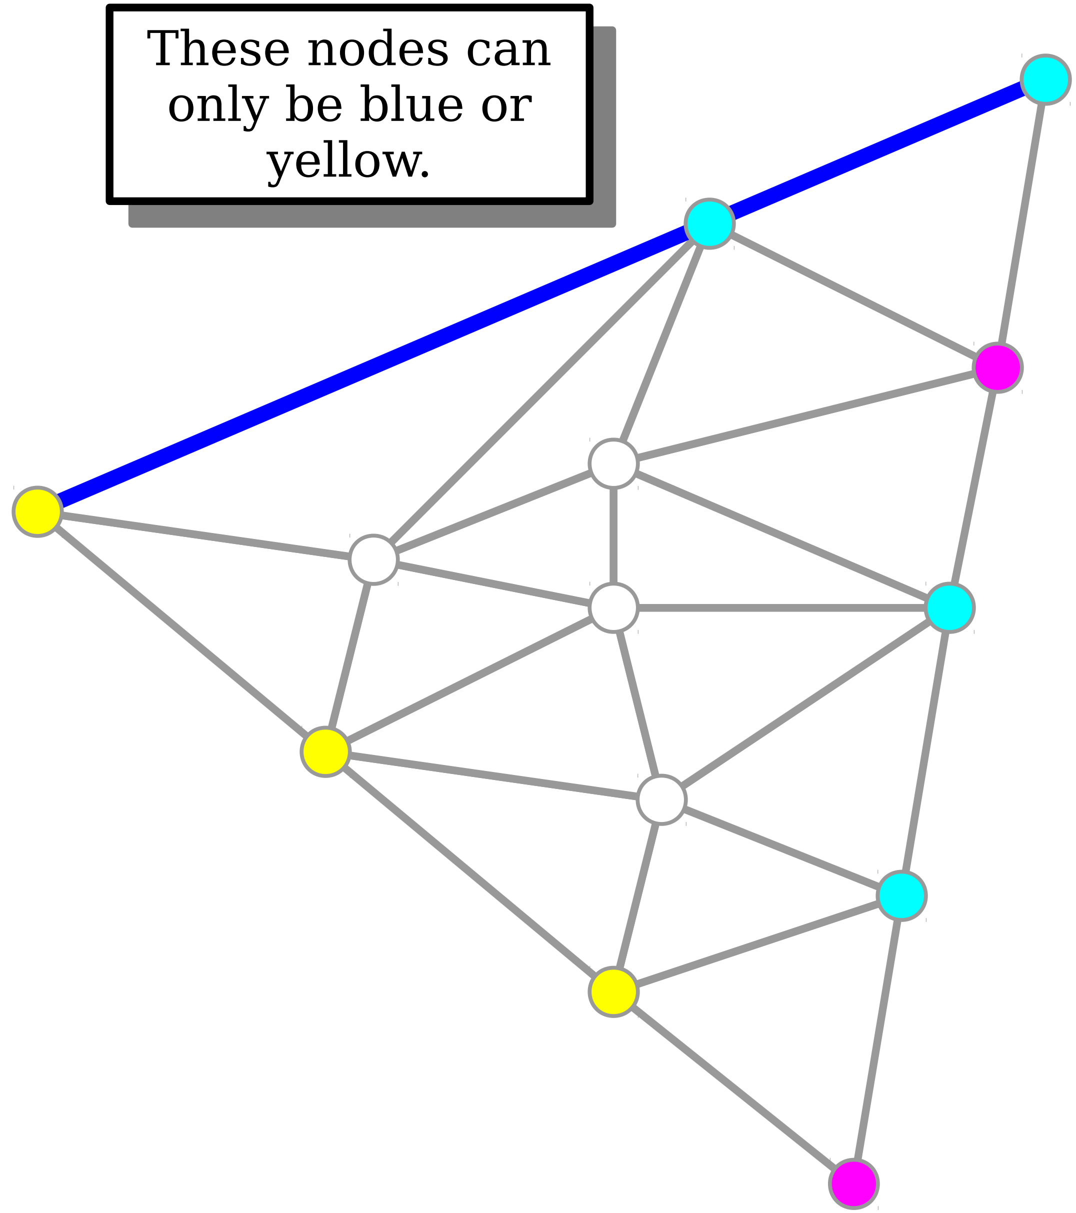

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.

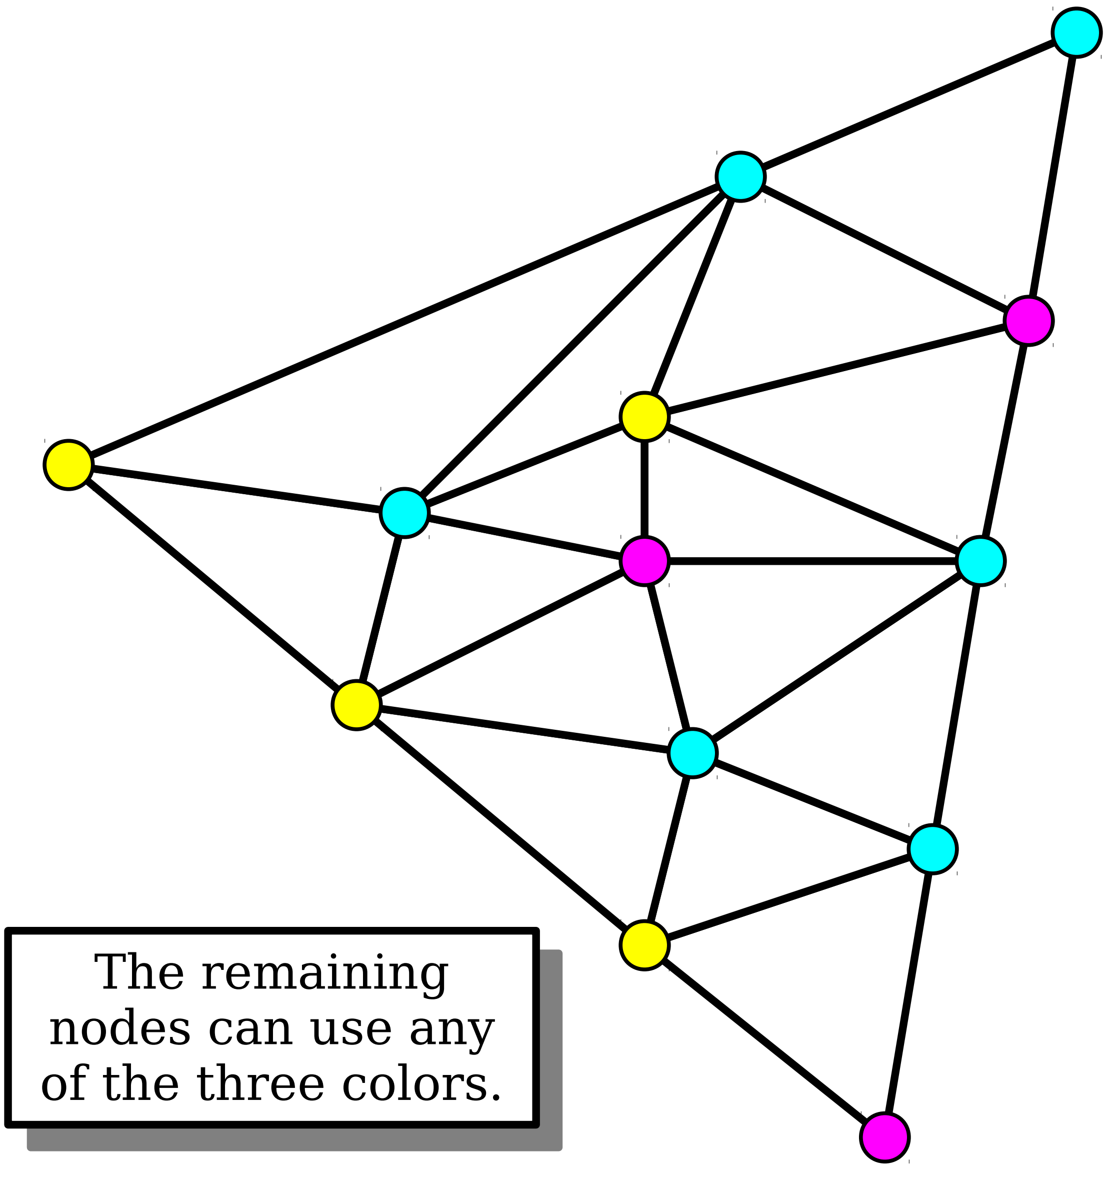

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.
- ***Sperner's Lemma:*** At least one of the smaller triangles will have corners of all three colors.

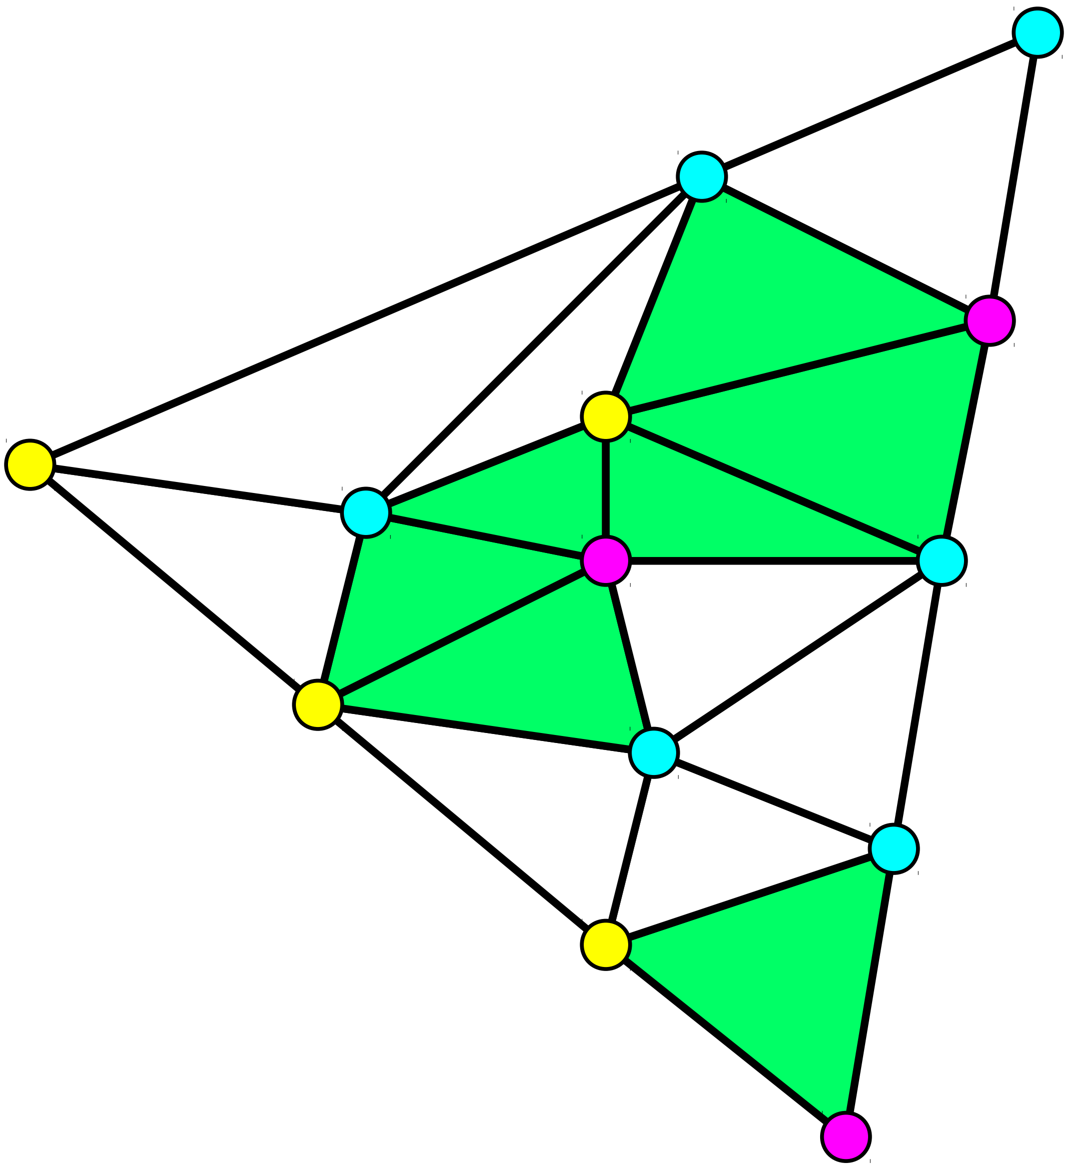

- Begin with a triangle.
- Subdivide the triangle into smaller triangles that meet corner-to-corner (that is, with no corner of one triangle intersecting the edge of another)
- Color the three corners of the big triangle three different colors.
- Color each other vertex using those three colors however you'd like, *except* that each vertex on an edge of the big triangle must not be the same color as the opposite corner of the big triangle.
- ***Sperner's Lemma:*** At least one of the smaller triangles will have corners of all three colors.

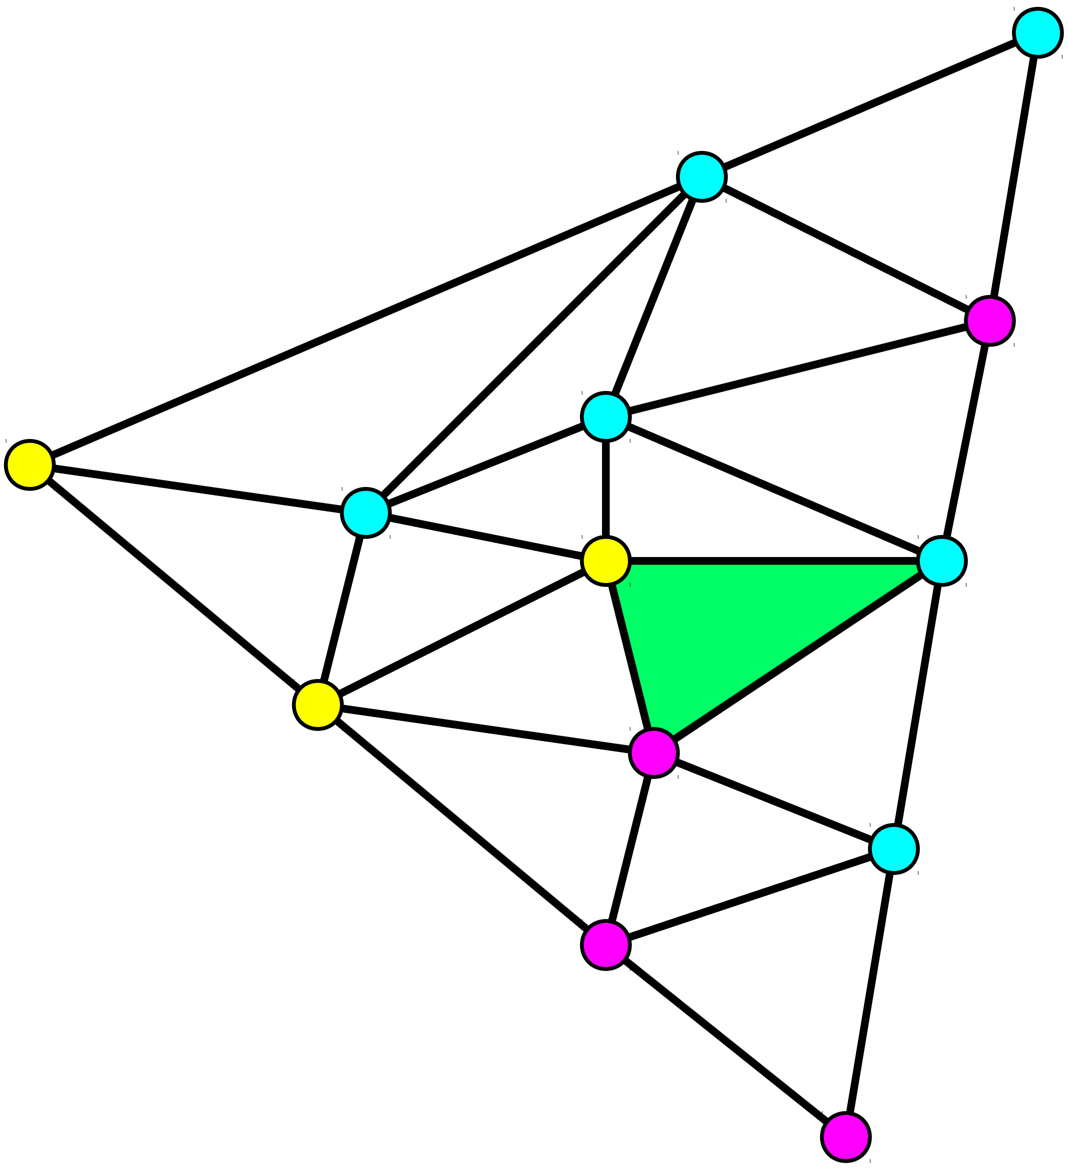

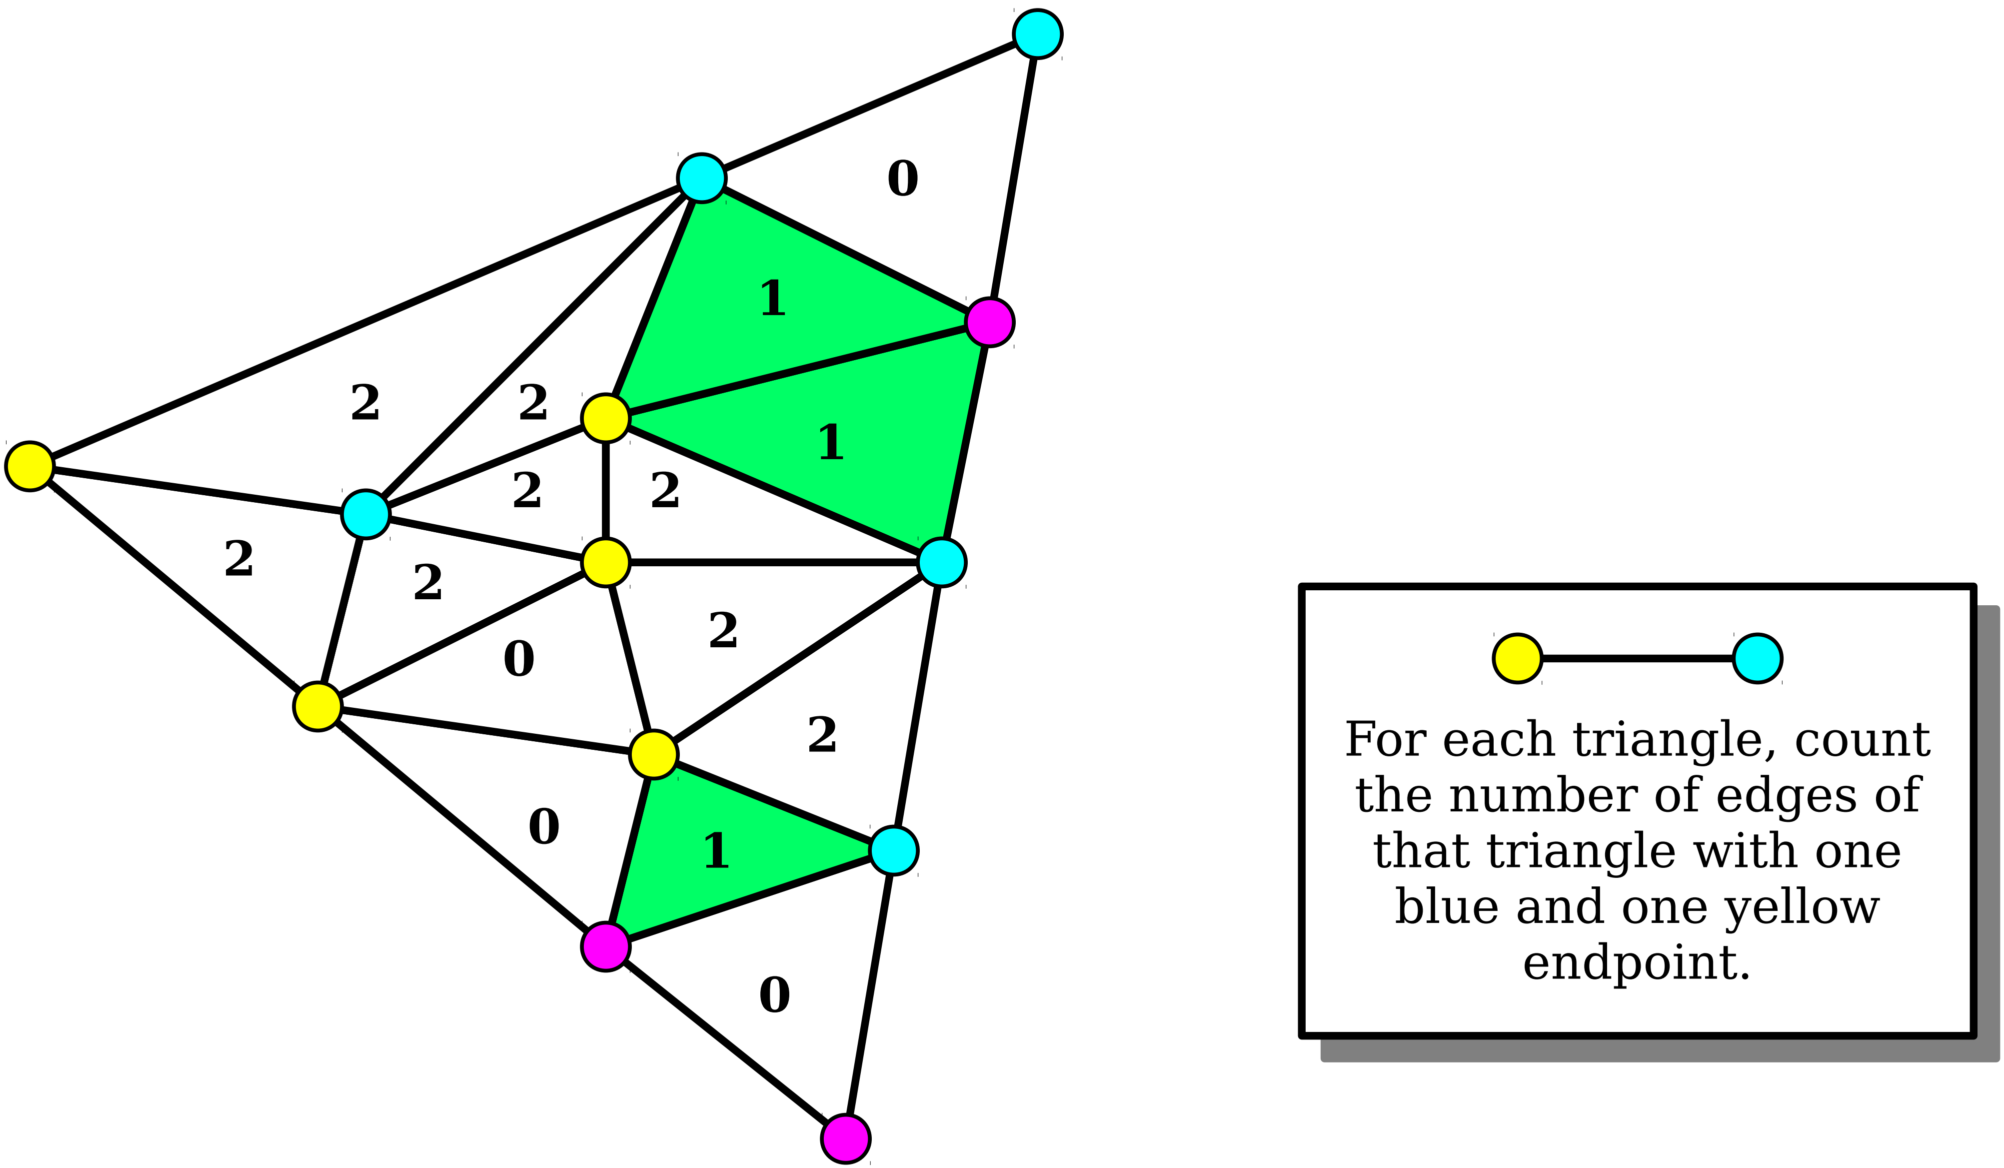
For each triangle, count
the number of edges of
that triangle with one
blue and one yellow
endpoint.

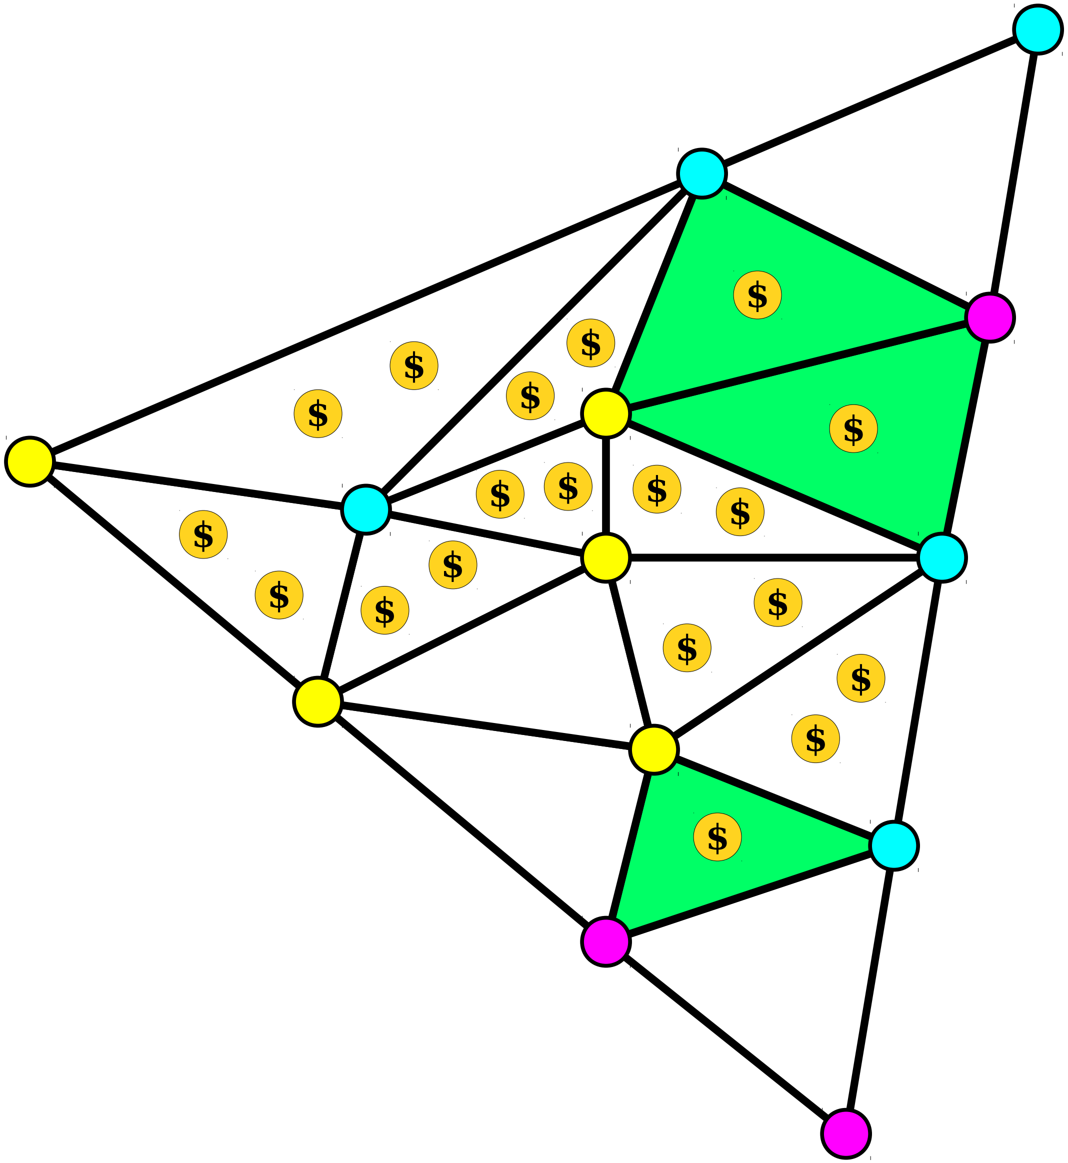

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

$

To make things a bit easier to see, imagine putting one coin into each triangle for each of its blue/yellow edges.

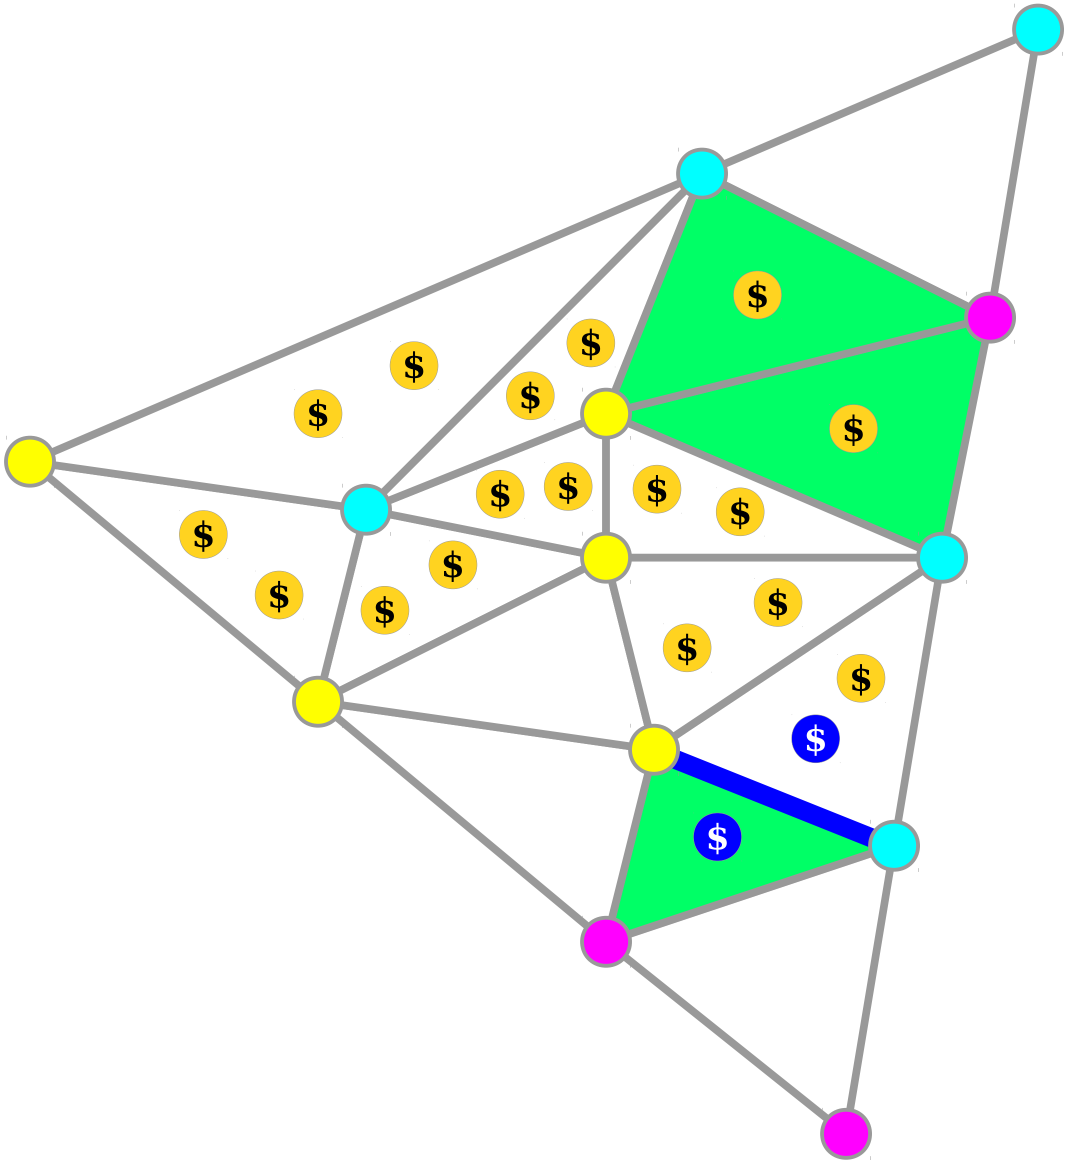

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

How many total coins are there across all of these triangles?

Each ***internal*** blue/yellow edge contributes two coins, one to the triangle on each side of the edge.

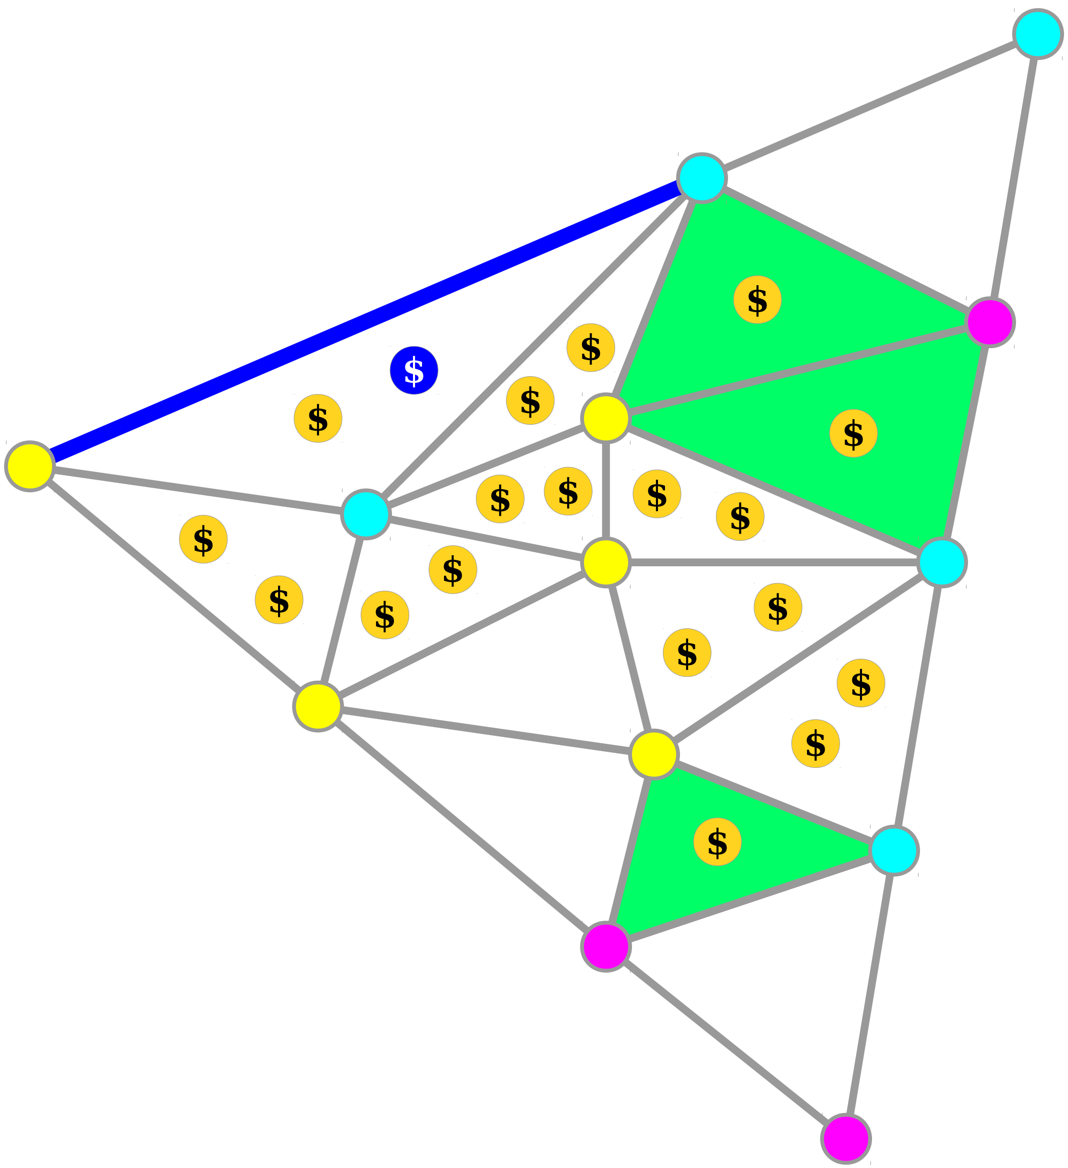

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

How many total coins are there across all of these triangles?

Each ***internal*** blue/yellow edge contributes two coins, one to the triangle on each side of the edge.

Each ***external*** blue/yellow edge contributes one coin to the triangle it's on.

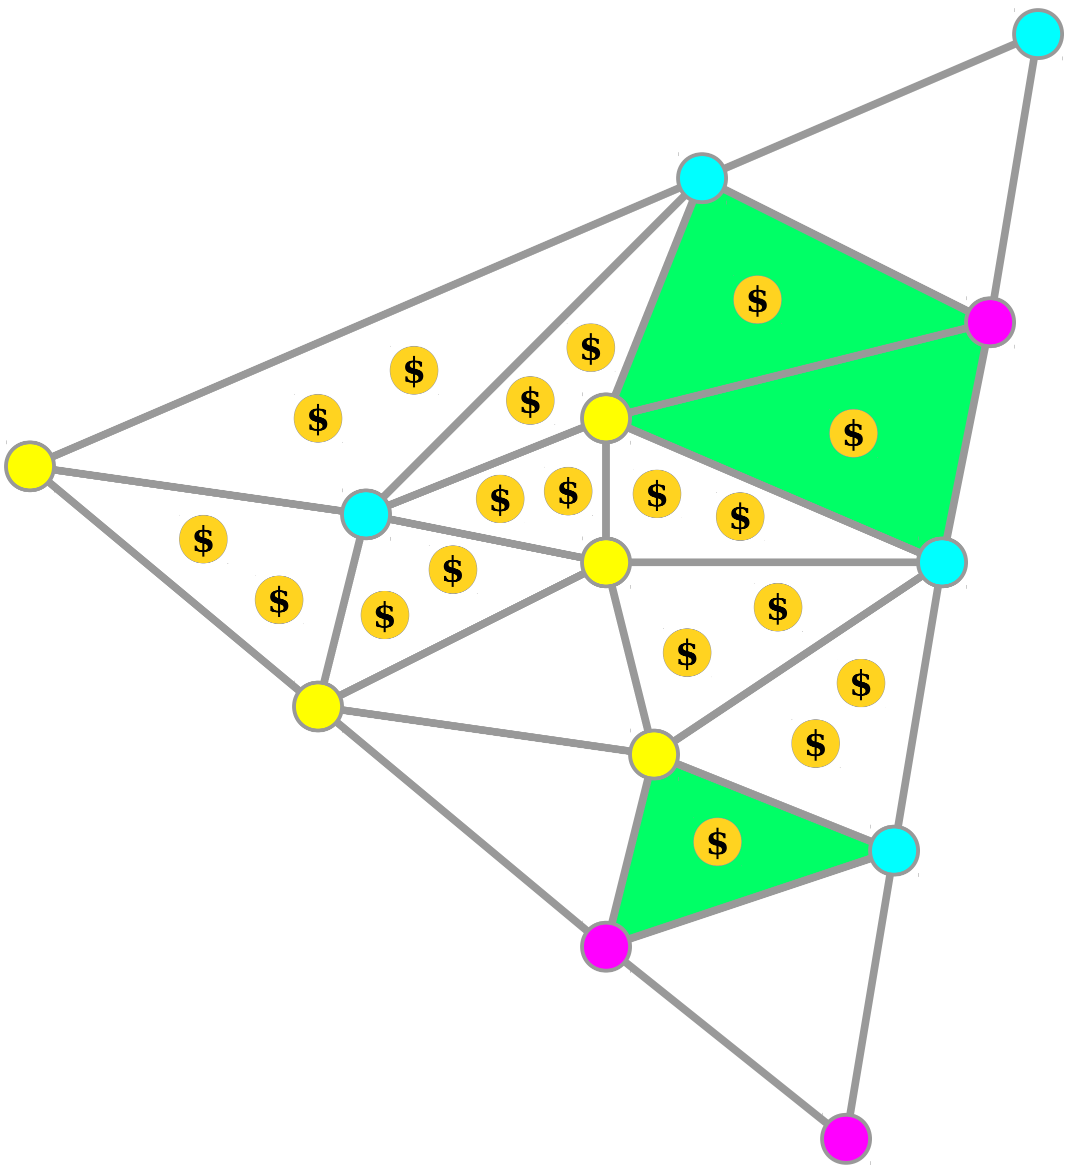

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

Lemma 3: The total number of coins is equal to $2I + E$, where I is the number of internal blue/yellow edges and E is the number of external blue/yellow edges.

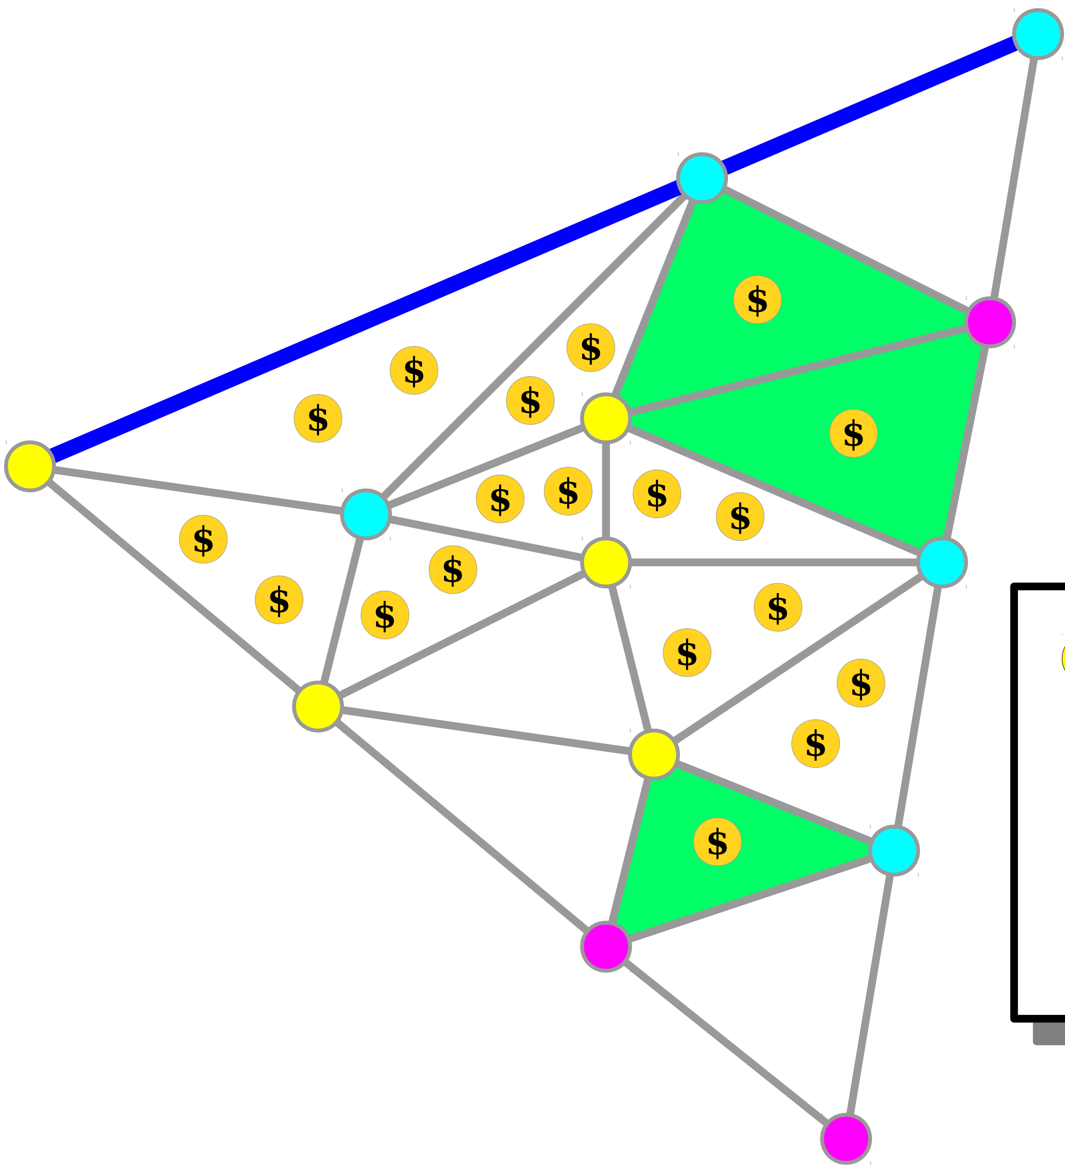

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

Lemma 3: The total number of coins is equal to $2I + E$, where I is the number of internal blue/yellow edges and E is the number of external blue/yellow edges.

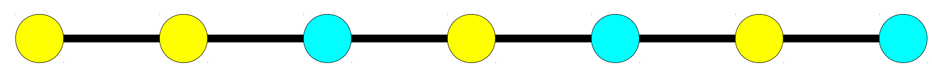

There have to be an odd number of blue/yellow edges on this side, since we start yellow, end blue, and each blue/yellow edge toggles the color.

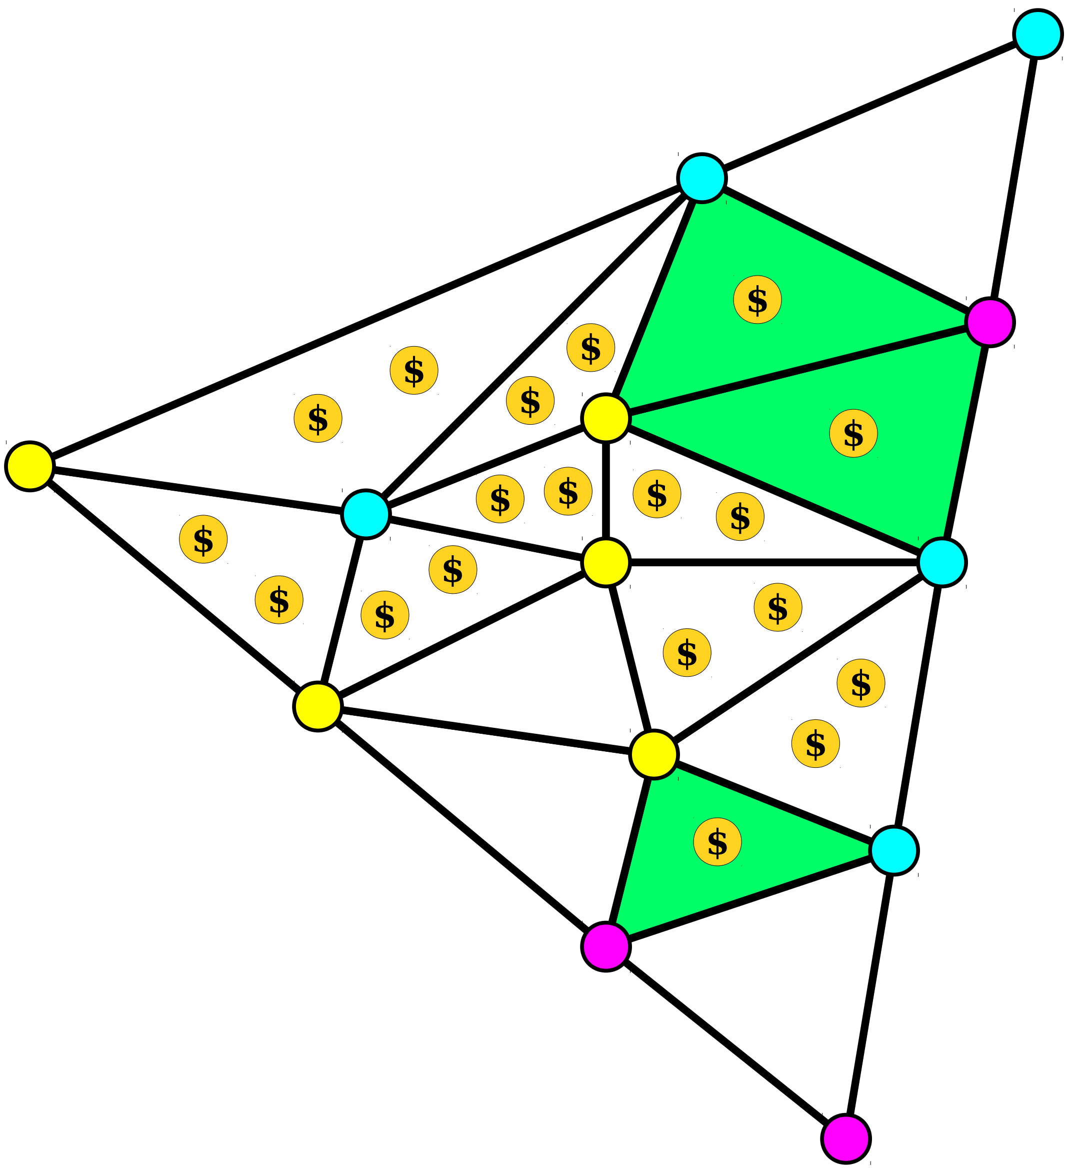

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

Lemma 3: The total number of coins is equal to $2I + E$, where I is the number of internal blue/yellow edges and E is the number of external blue/yellow edges.

Lemma 4: E is odd.

Lemma 5: The total number of coins distributed across all the triangles is odd.

Proof of Sperner's Lemma:
Consider any triangle subdivided and colored according to the rules. Place one coin into each triangle for each blue/yellow edge of that triangle.

By Lemma 5, there are an odd number of coins distributed across the triangles, so by our earlier theorem there must be an odd number of triangles containing an odd number of coins.

This means that there must be at least one triangle containing an odd number of coins, since otherwise there would be zero such triangles and zero is not odd.

Therefore, by Lemma 2, there must be at least one triangle with one corner of each color. ■

Lemma 1: A triangle has corners of three different colors iff it has exactly one coin in it.

Lemma 2: A triangle has corners of three different colors iff it has an odd number of coins in it.

Lemma 3: The total number of coins is equal to $2I + E$, where I is the number of internal blue/yellow edges and E is the number of external blue/yellow edges.

Lemma 4: E is odd.

Lemma 5: The total number of coins distributed across all the triangles is odd.

Some Nifty Applications

- Sperner's lemma is a key step in many ***fair division protocols***, ways of splitting up chores or rent in a way that everyone is happy.
  - See this ***New York Times*** article, for example.
- It's also a key step in ***Monsky's theorem***, which says that you can't split a square into an odd number of triangles of equal area.
  - This theorem has an ***unusual history*** and relates several seemingly unrelated branches of mathematics together.
- *And* it's used in an elegant proof of the ***Brouwer fixed-point theorem***, which says (among other things) that no matter how you stir your coffee, there must be some point in the coffee that ***stays in the same place that it started***.
- *And they're all based on the simple idea of looking at what happens if you divide an odd number of things into some buckets!*

Next Time

- ***Mathematical Induction***
  - Proofs on stepwise processes
- ***Applications of Induction***
  - … to numbers!
  - … to data compression!
  - … to puzzles!
  - … to algorithms!