Graph Theory

Part Two

Recap from Last Time

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

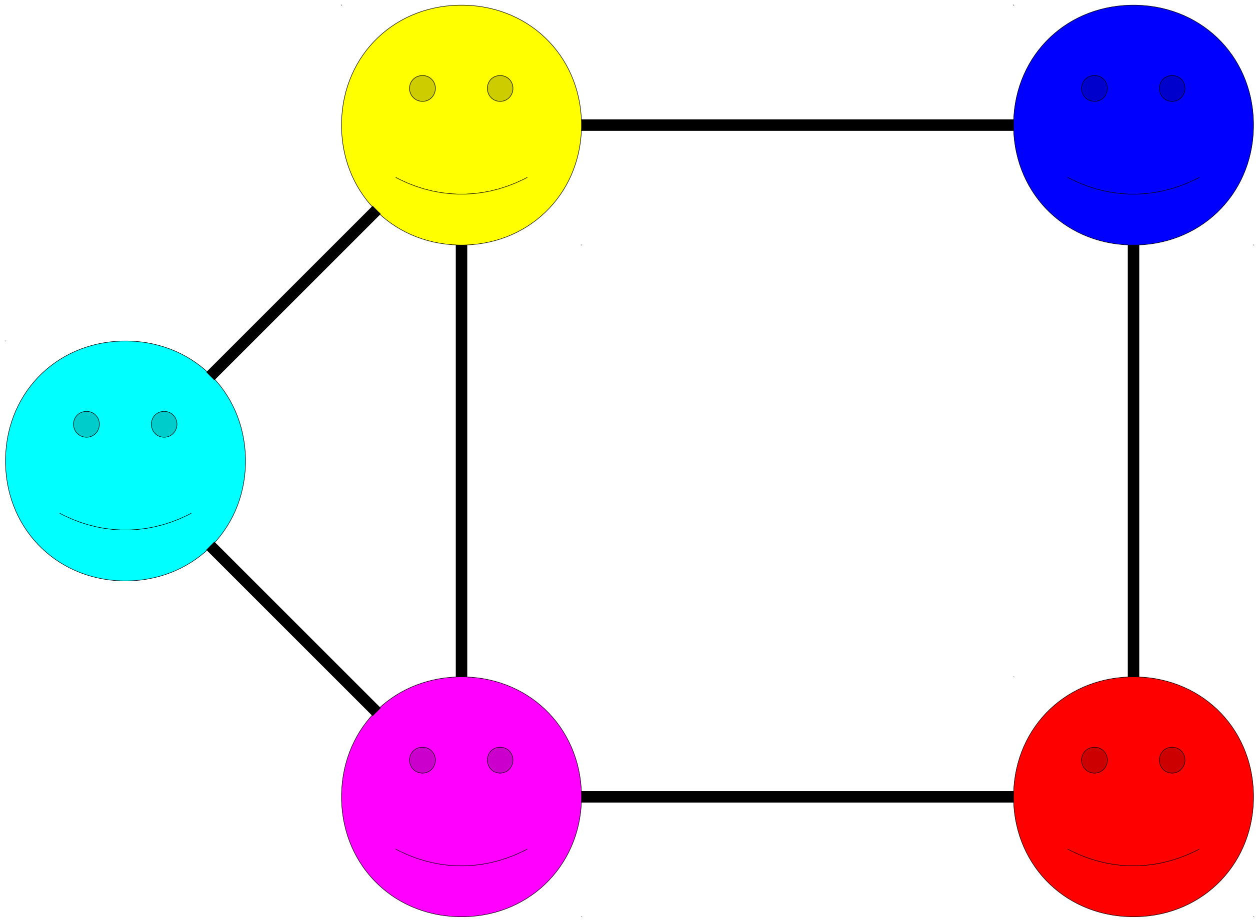

A graph consists of a set of ***nodes*** (or ***vertices***) connected by ***edges*** (or ***arcs***)

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

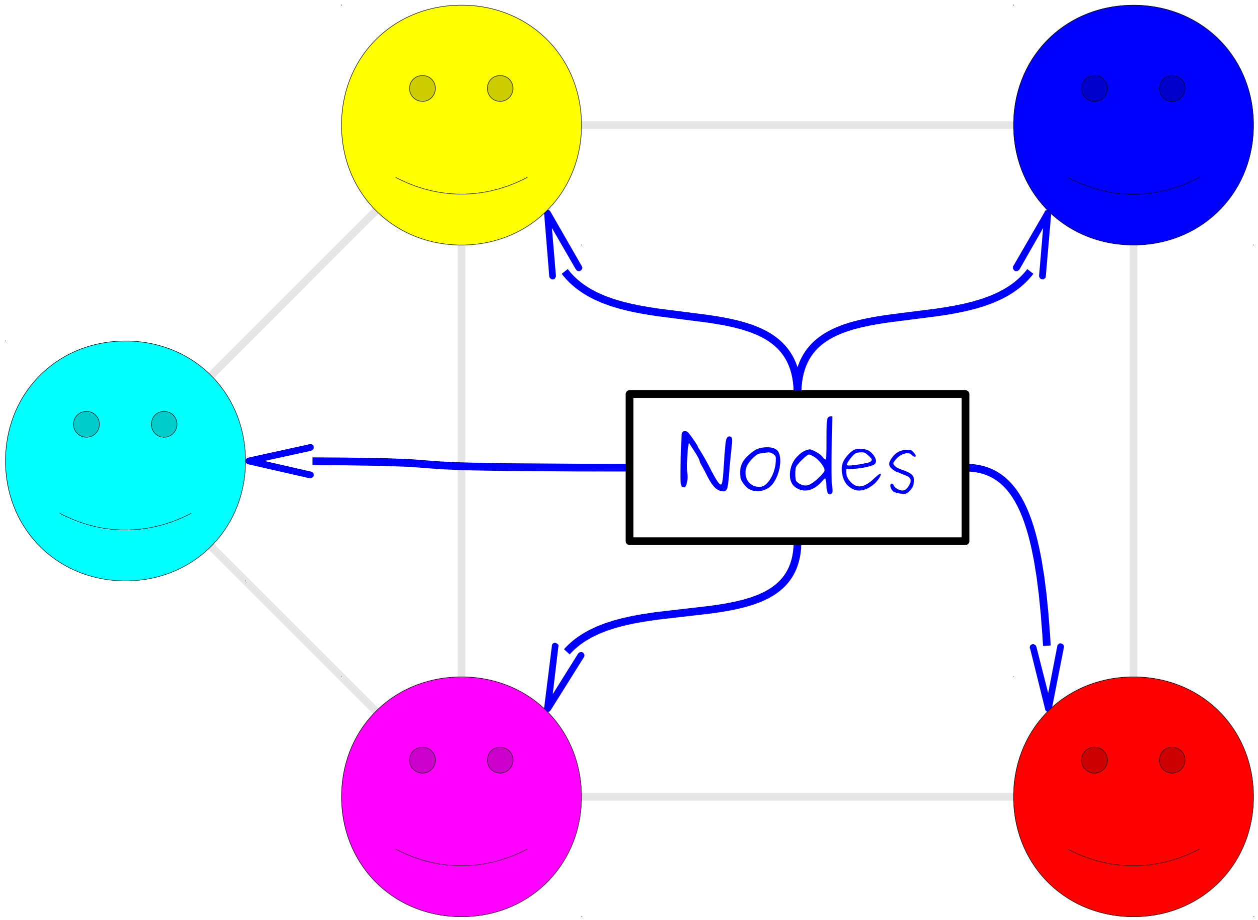

A graph consists of a set of ***nodes*** (or ***vertices***) connected by ***edges*** (or ***arcs***)

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

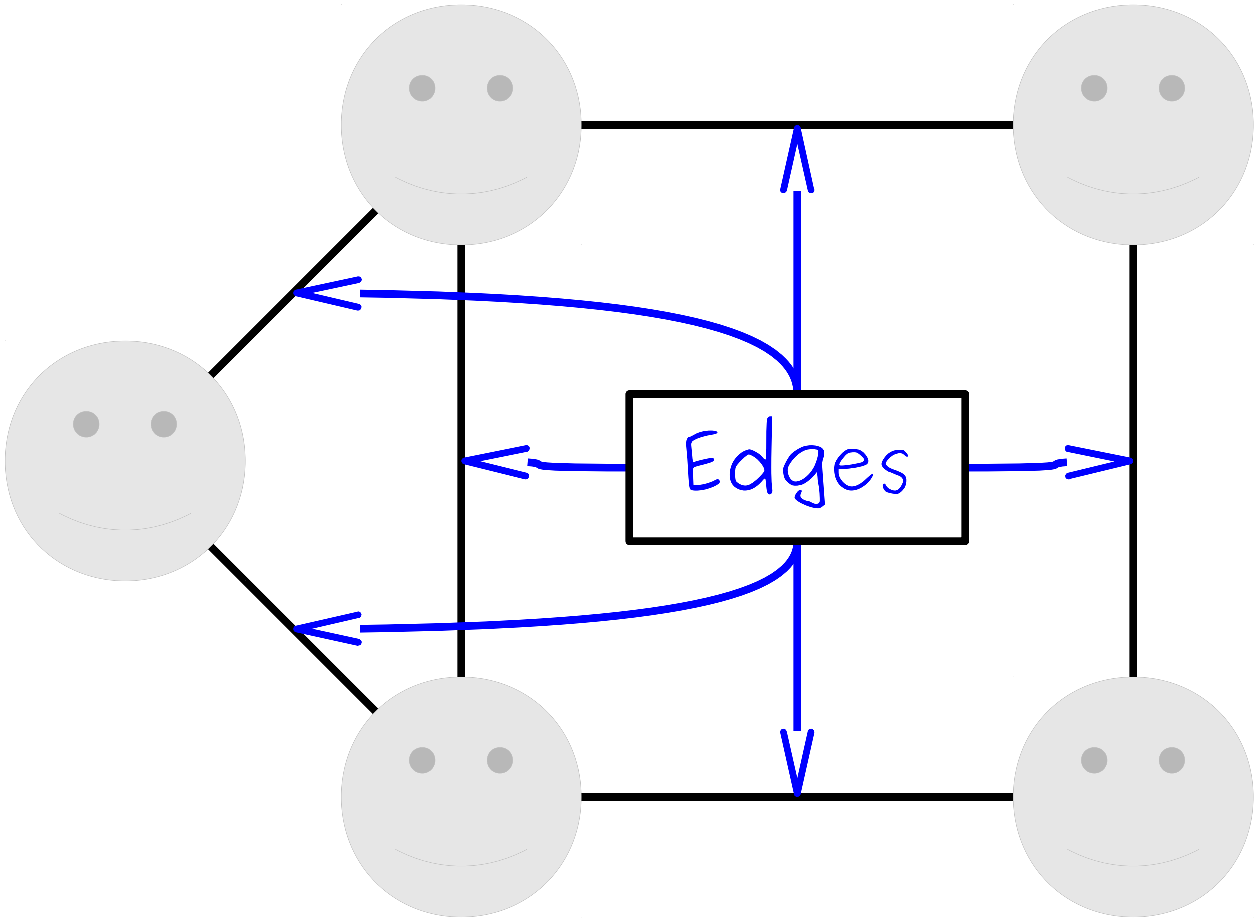

A graph consists of a set of ***nodes*** (or ***vertices***) connected by ***edges*** (or ***arcs***)

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-Vertex-Colorings

- If $G = (V, E)$ is a graph, a ***k-vertex-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.

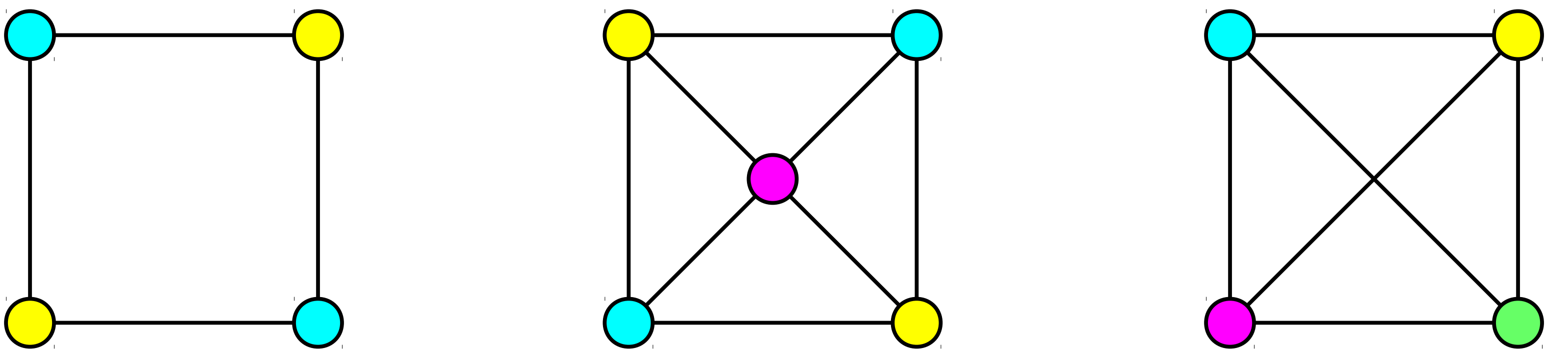

- The ***chromatic number*** of G, denoted $\boldsymbol{\chi(G)}$, is the minimum number of colors needed in any k-coloring of G.
- Today, we're going to see several results involving coloring parts of graphs. They don't necessarily involve *k-vertex-colorings* of graphs, so feel free to ask for clarifications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

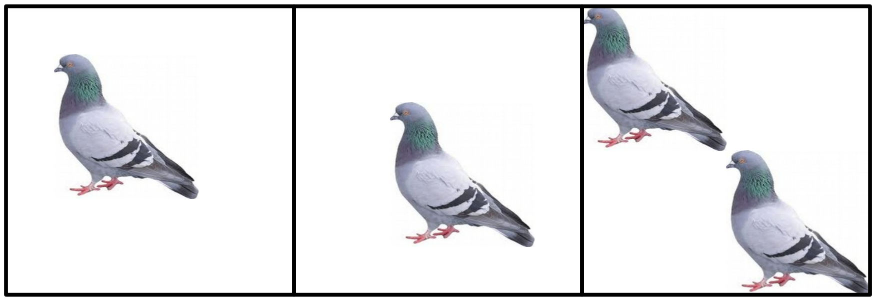

The Pigeonhole Principle

- ***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.

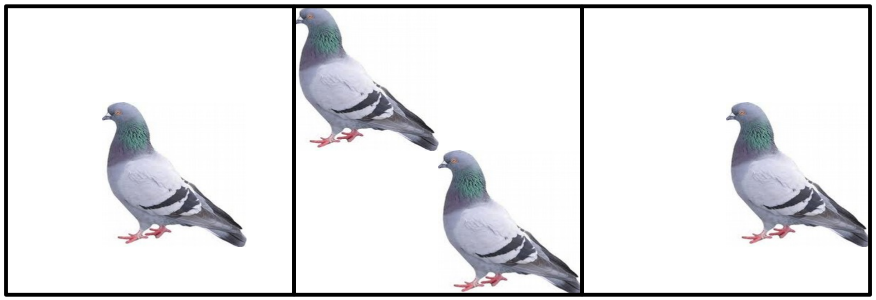

The Pigeonhole Principle

- ***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.

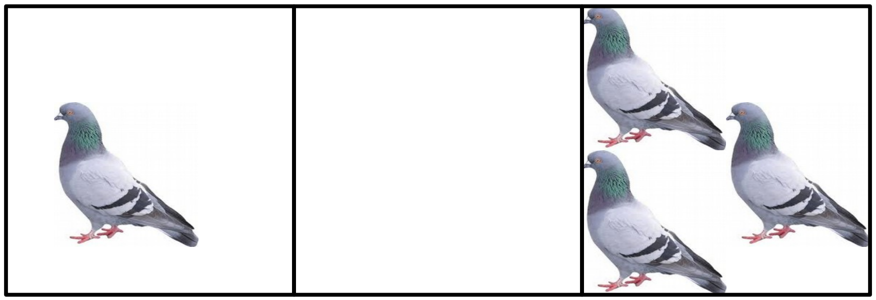

The Pigeonhole Principle

- ***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.

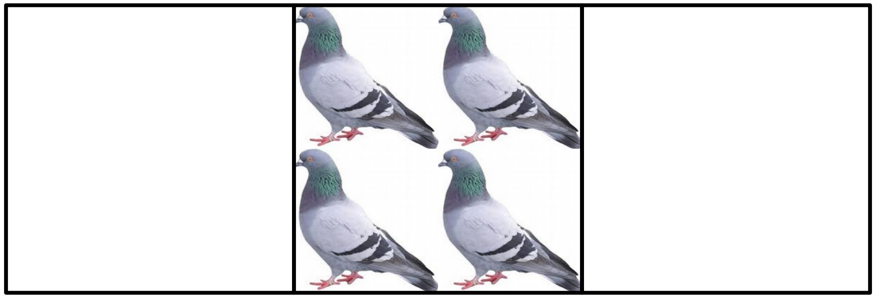

Some Simple Applications

- 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

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

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

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

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

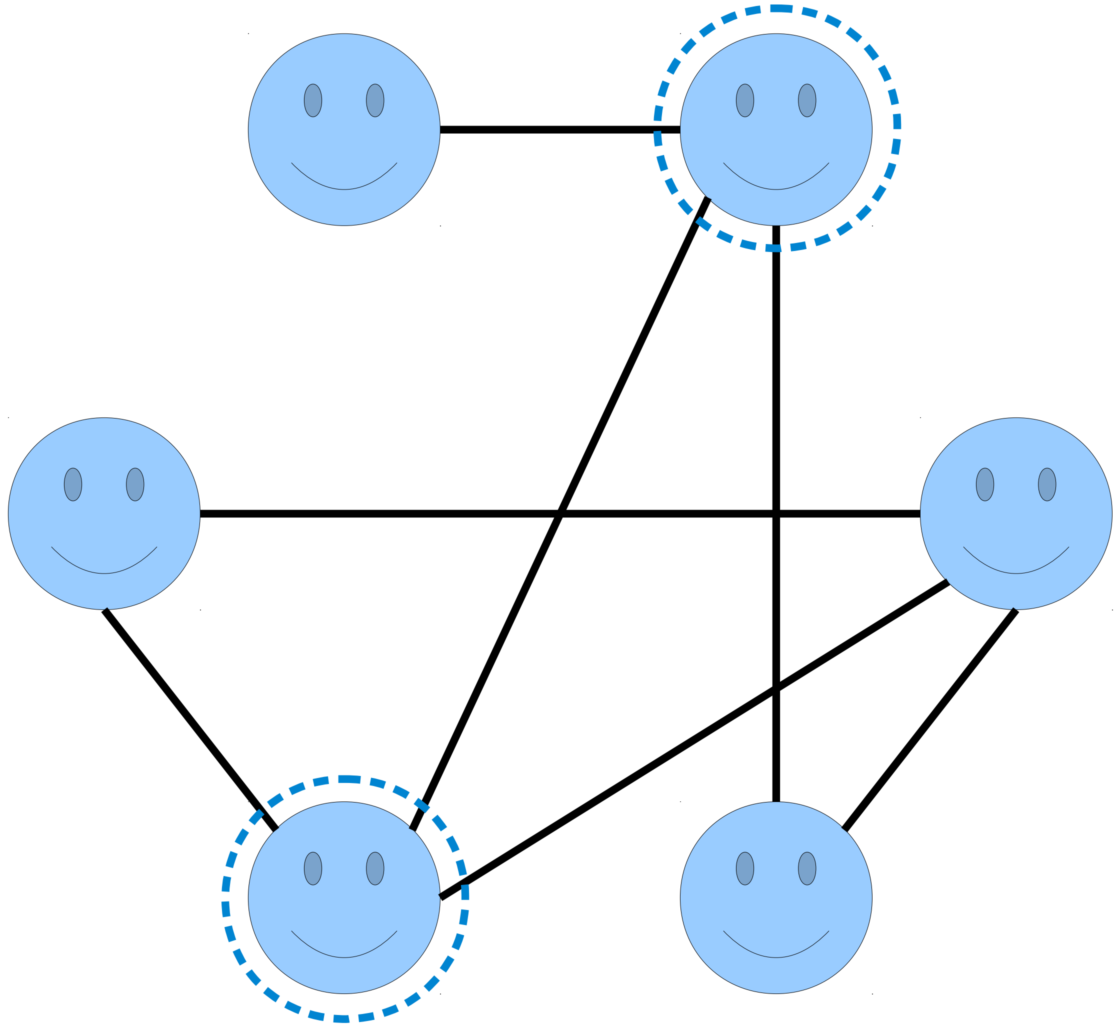

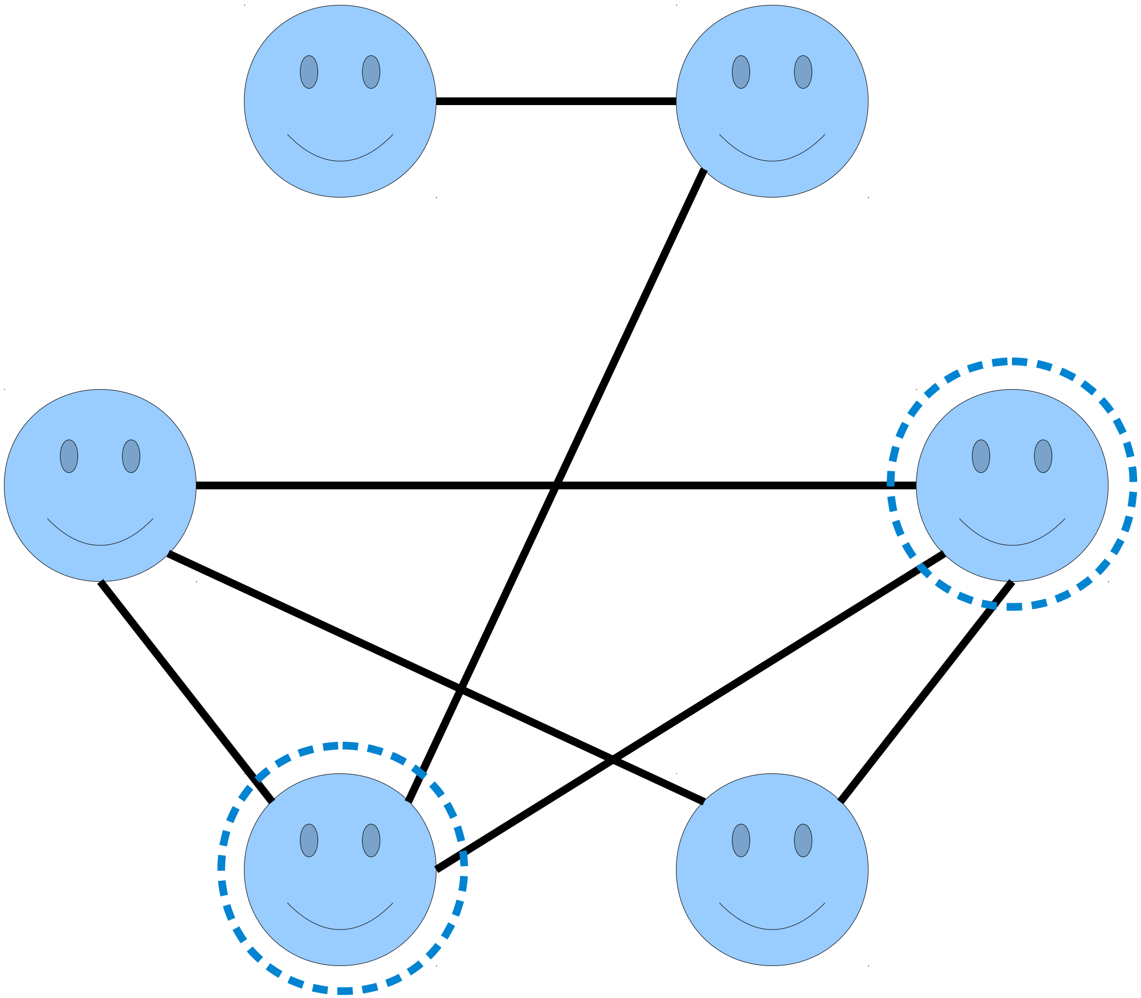

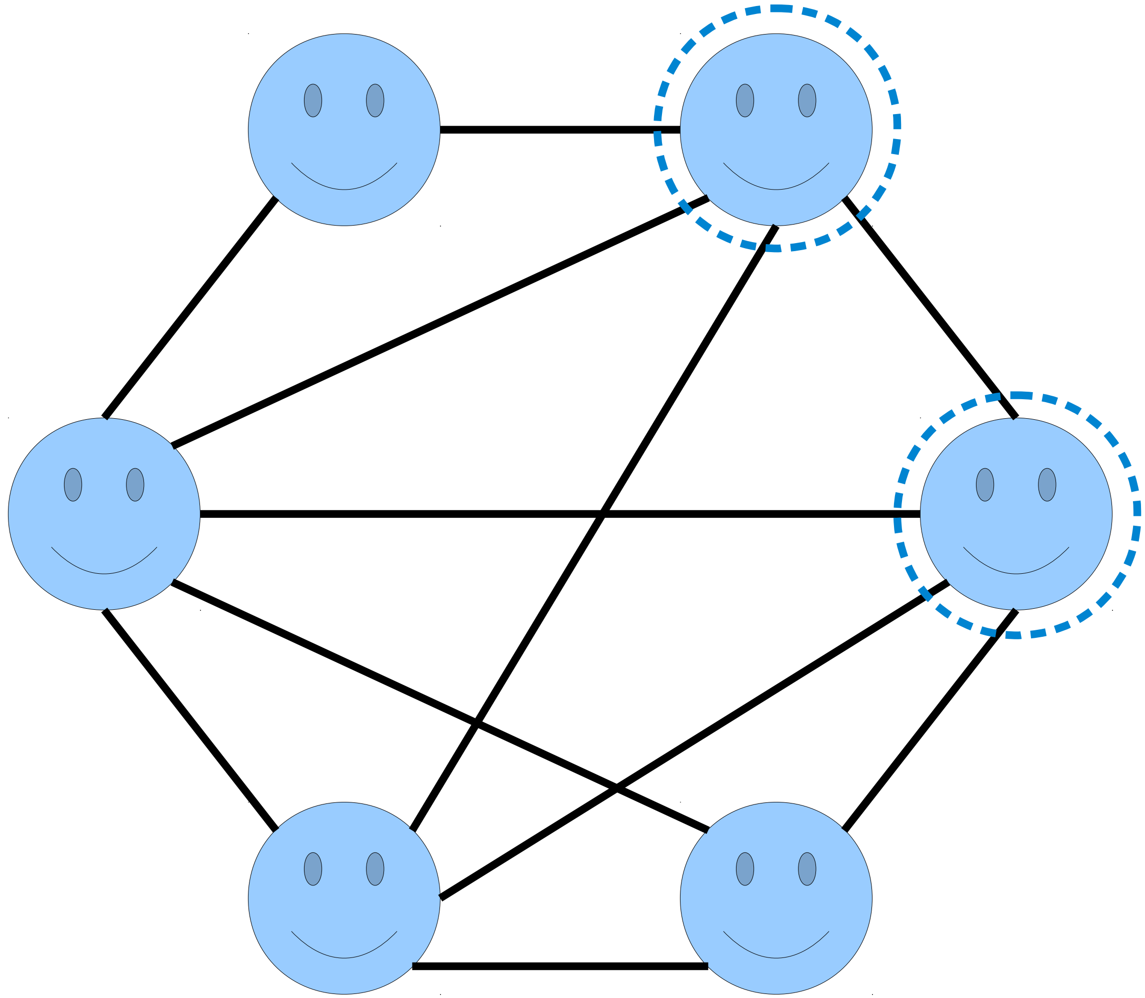

Degrees

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

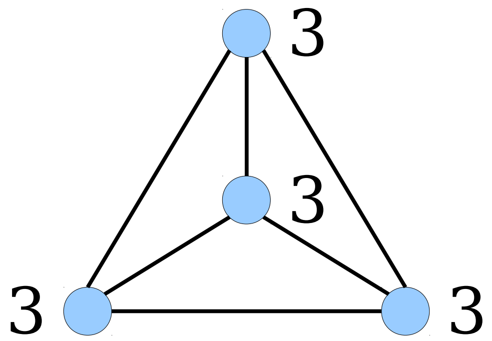

- ***Theorem:*** Every graph with at least two nodes has at least two nodes with the same degree.
  - Equivalently: at any party with at least two people, there are at least two people with the same number of Facebook friends at the party.

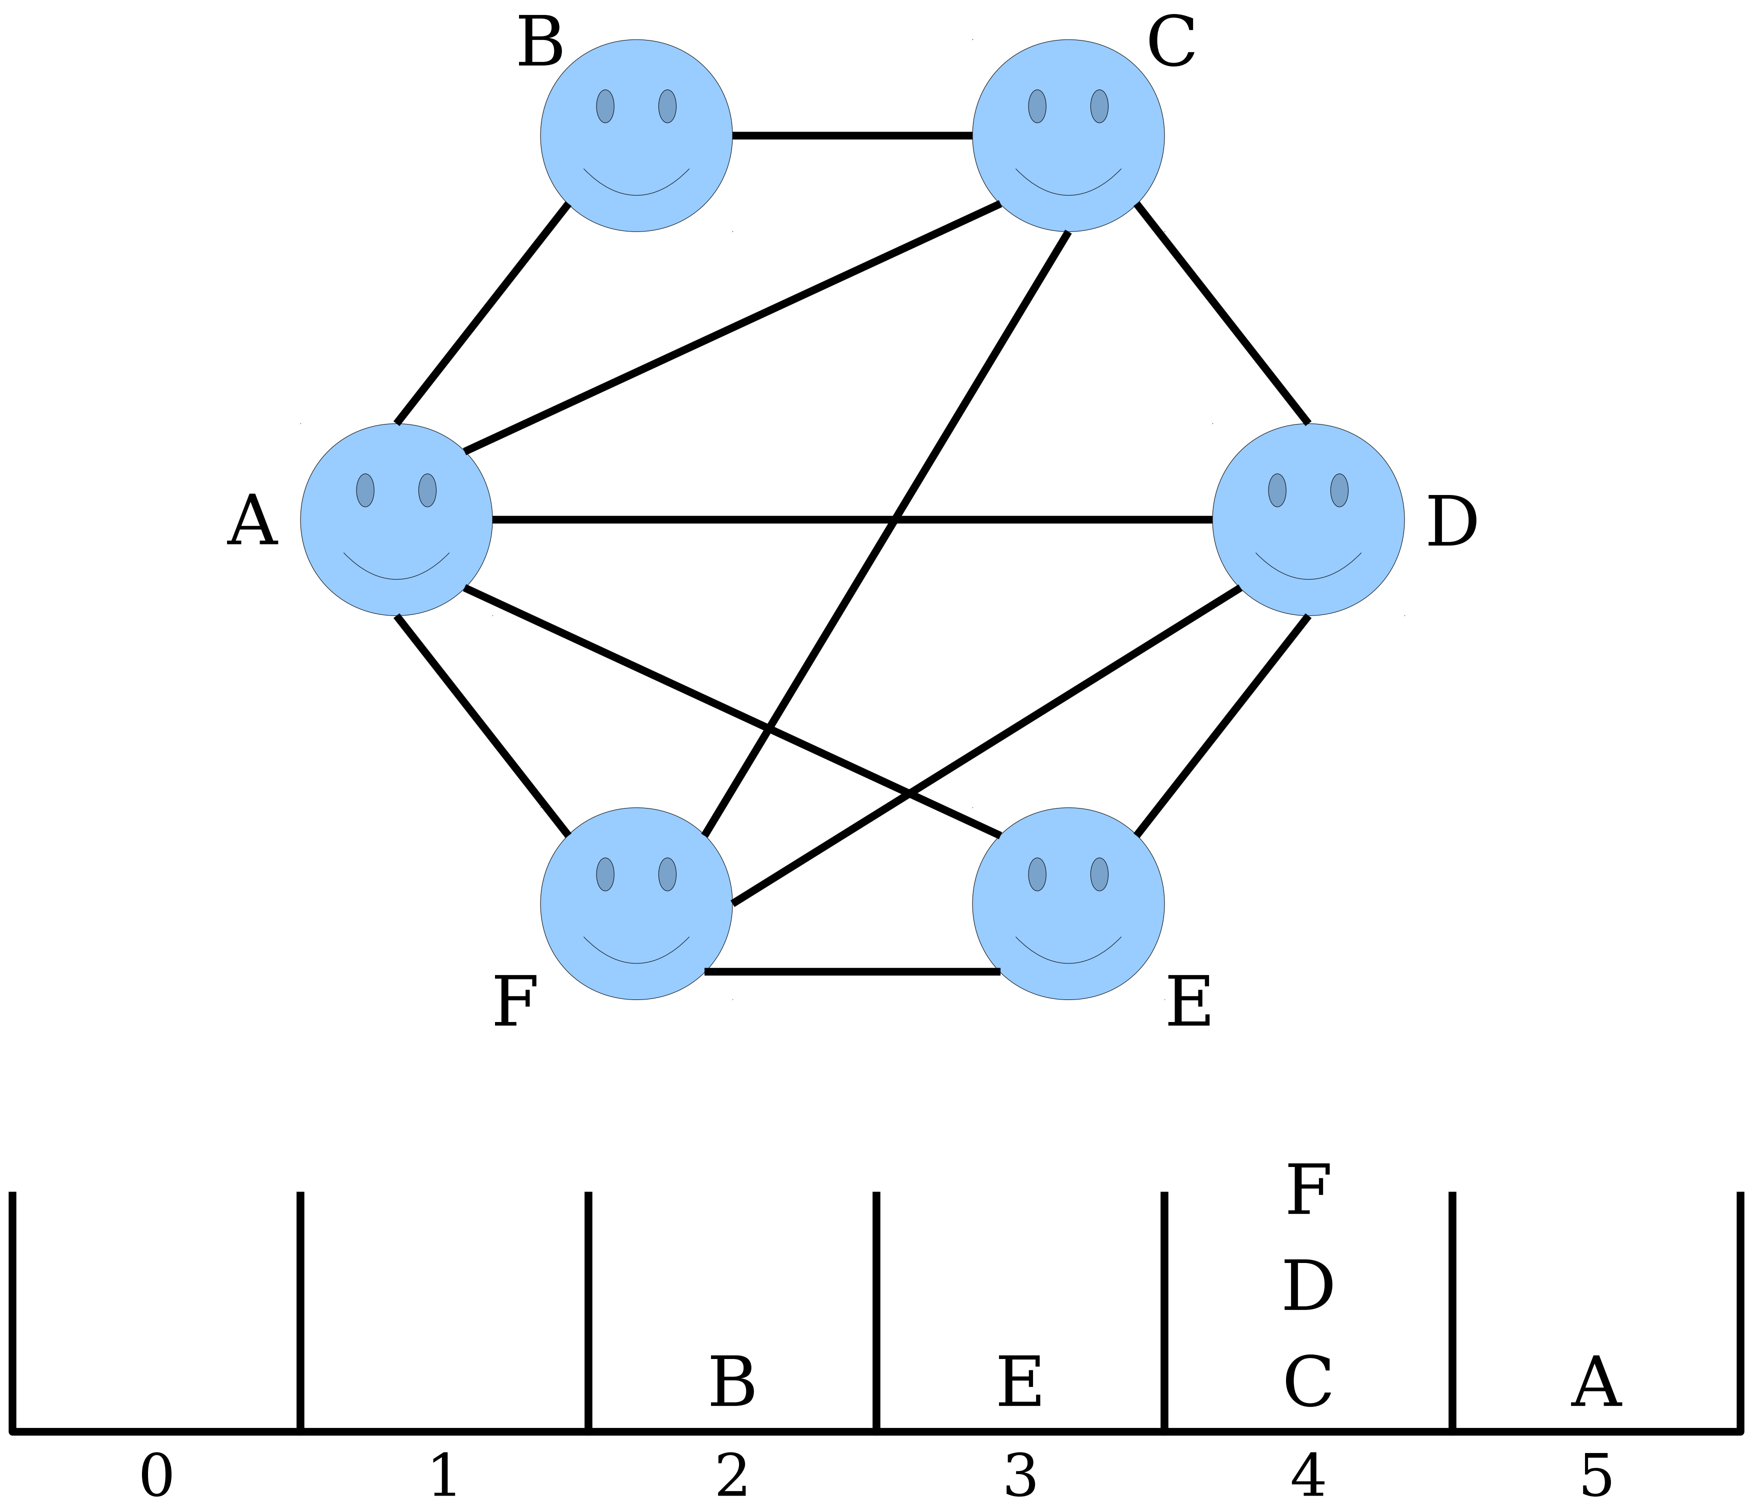
B
C
A
D
F
E
F
D
B
E
C
A
0
1
2
3
4
5

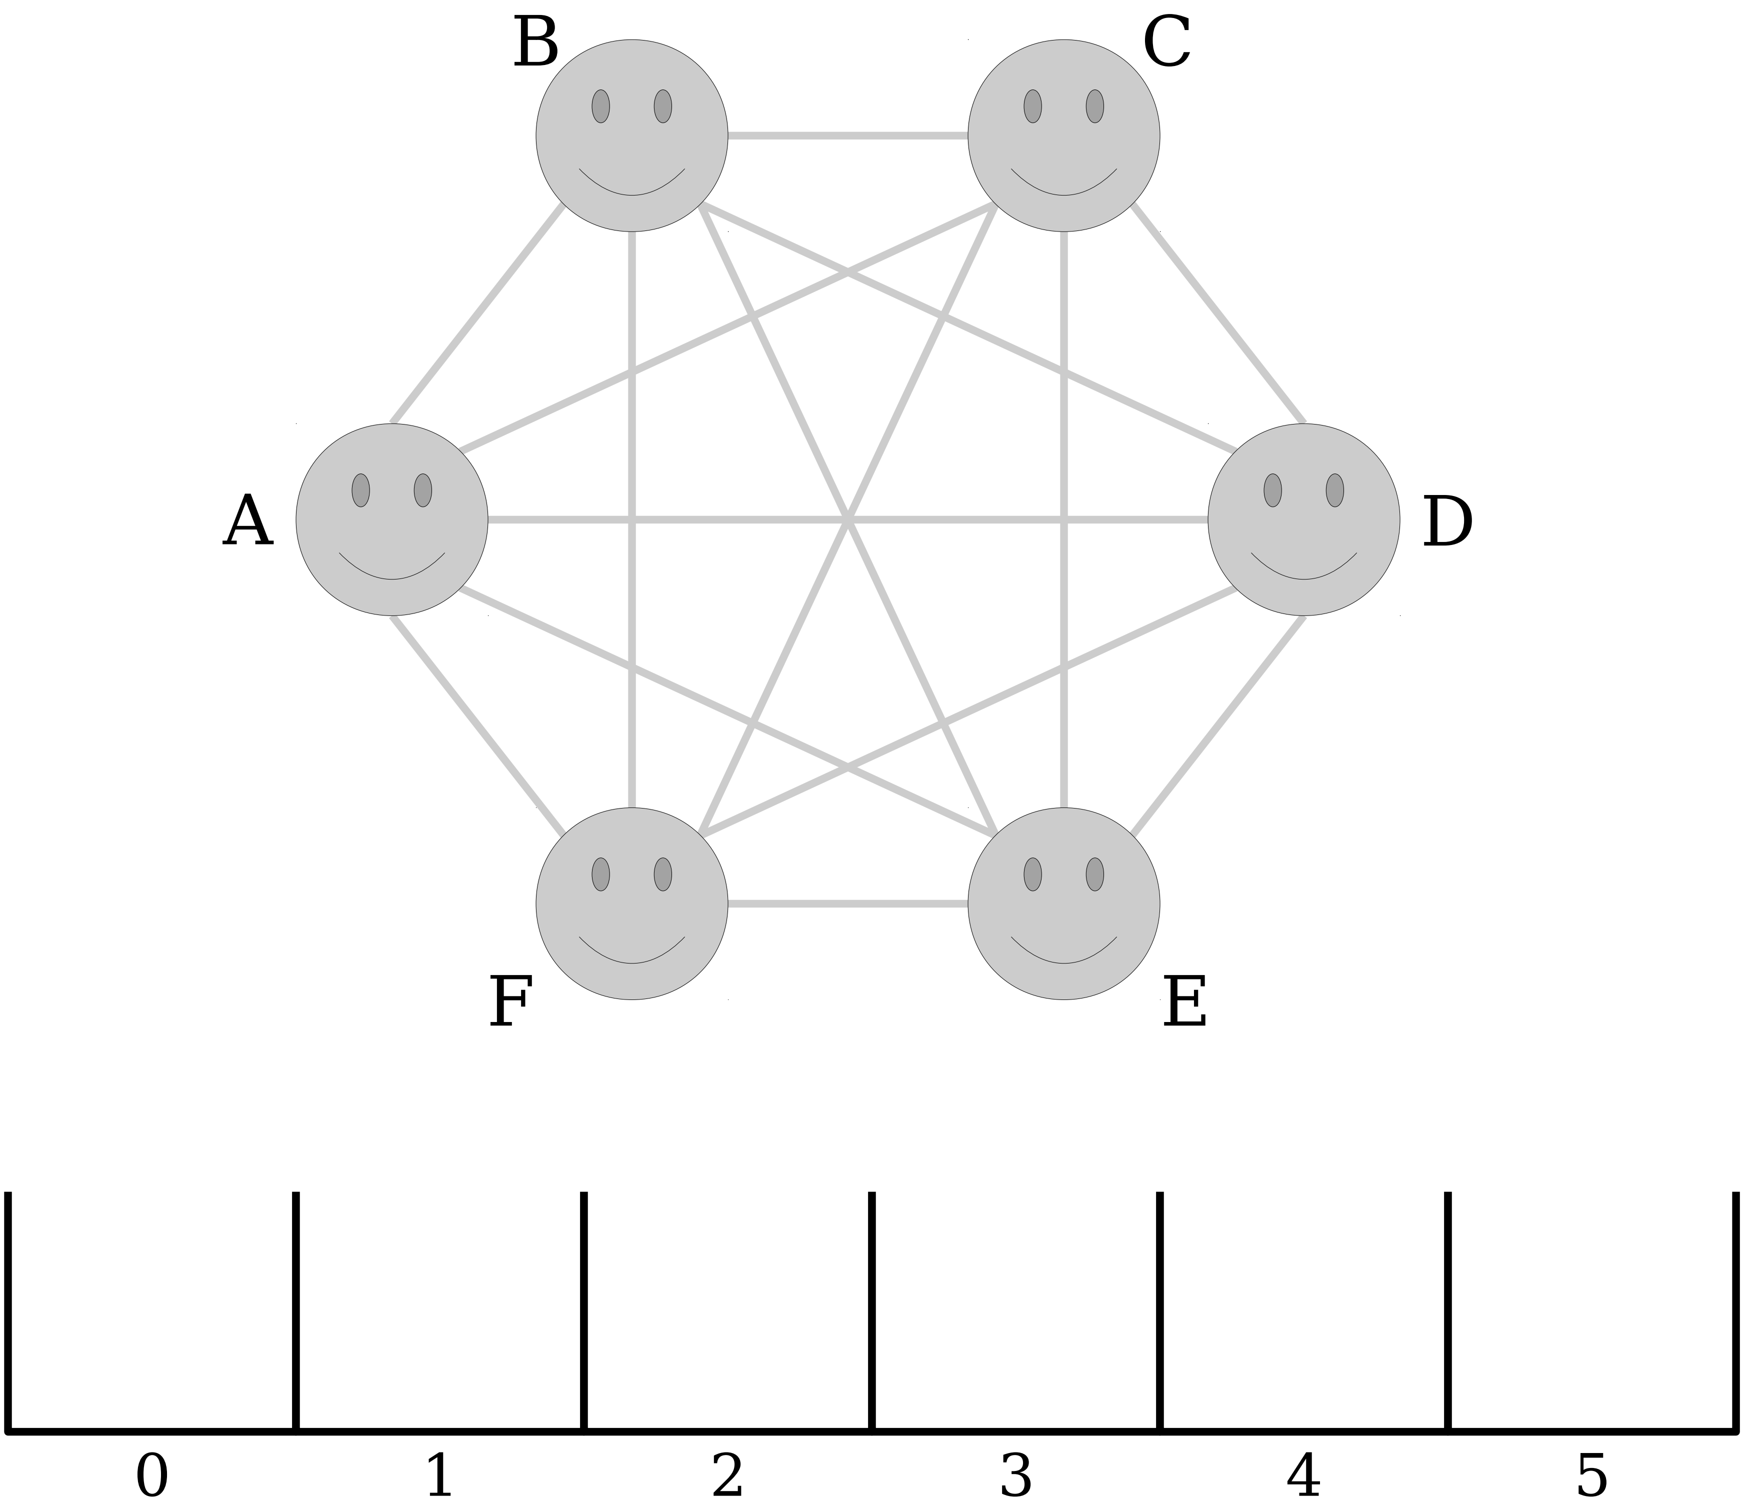
B
C
A
D
F
E
0
1
2
3
4
5

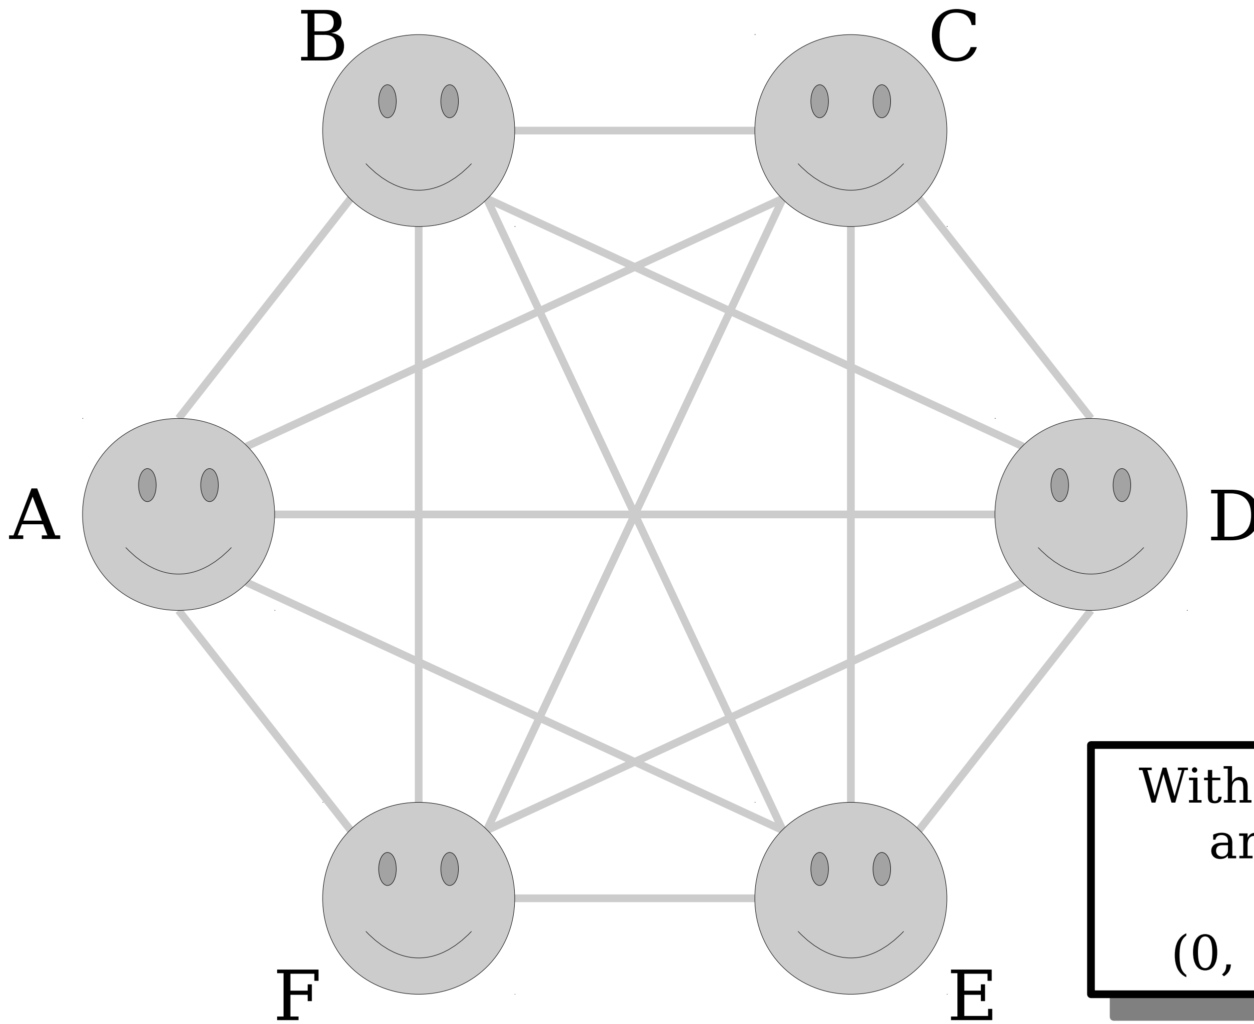

With n nodes, there are n possible degrees ($0, 1, 2, \ldots, n - 1$)

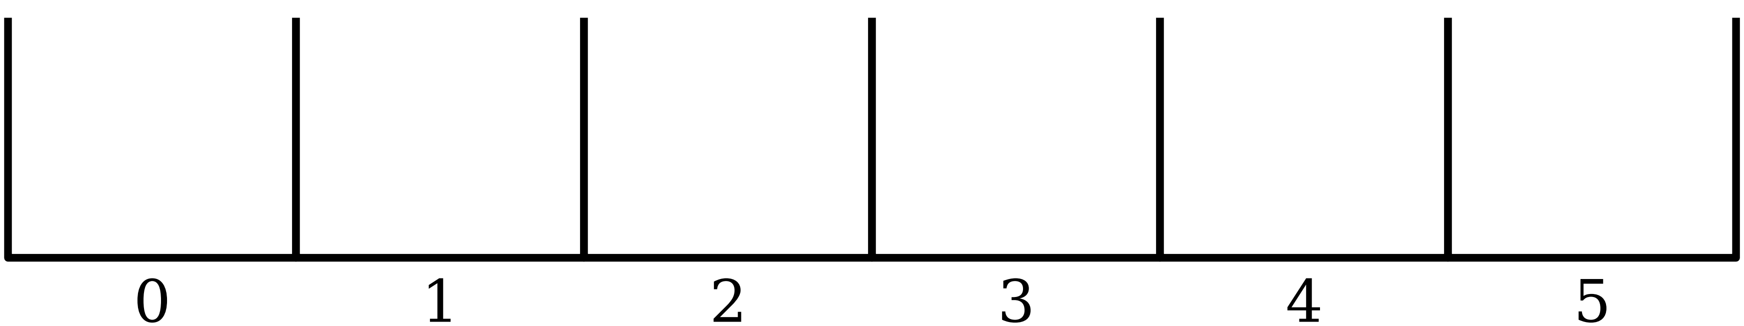

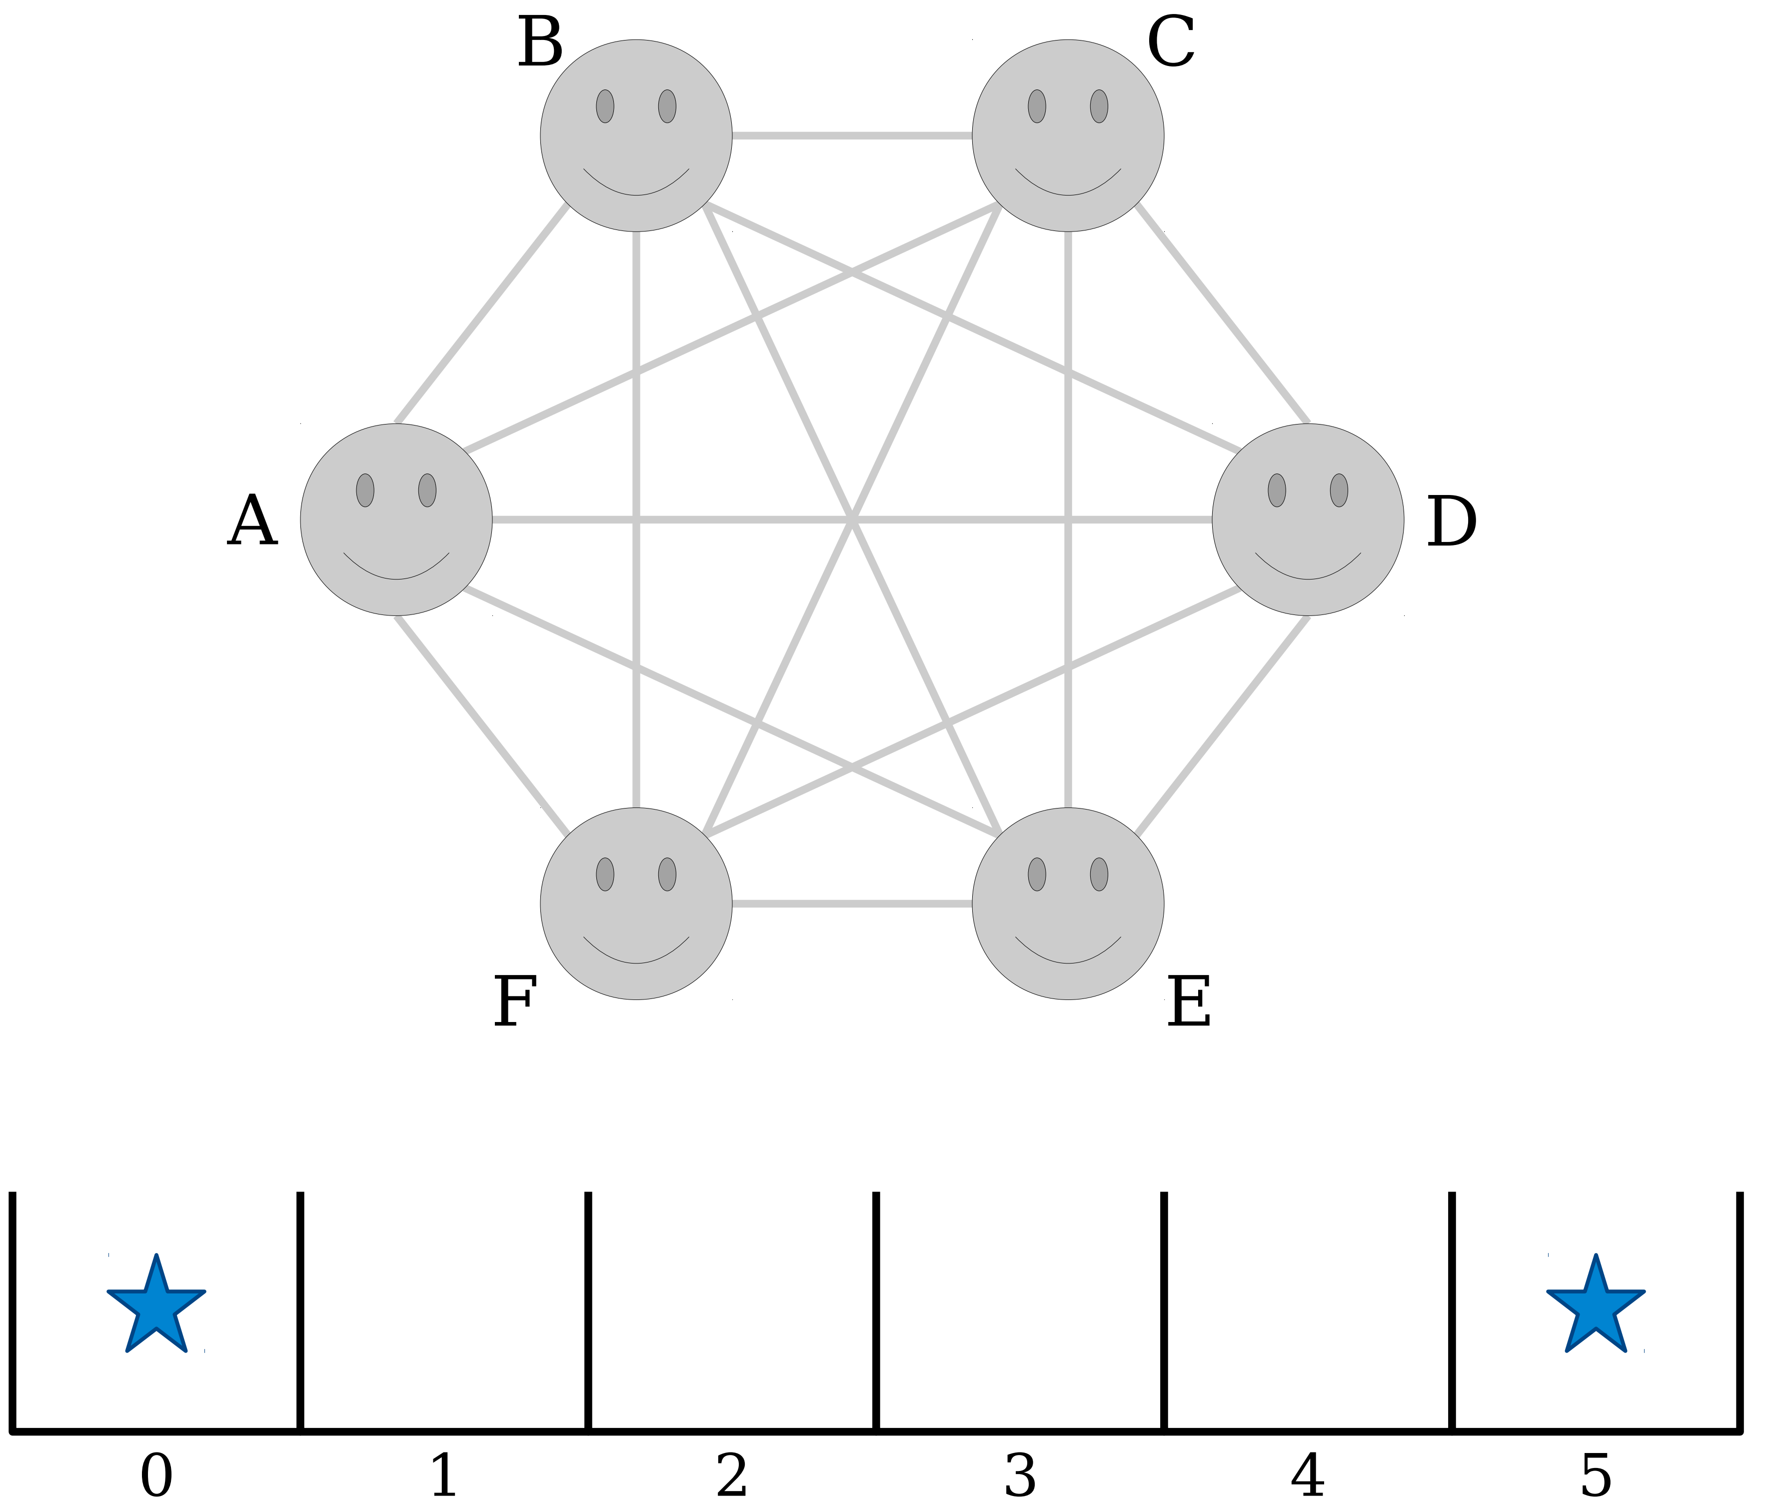
B
C
A
D
F
E
0
1
2
3
4
5

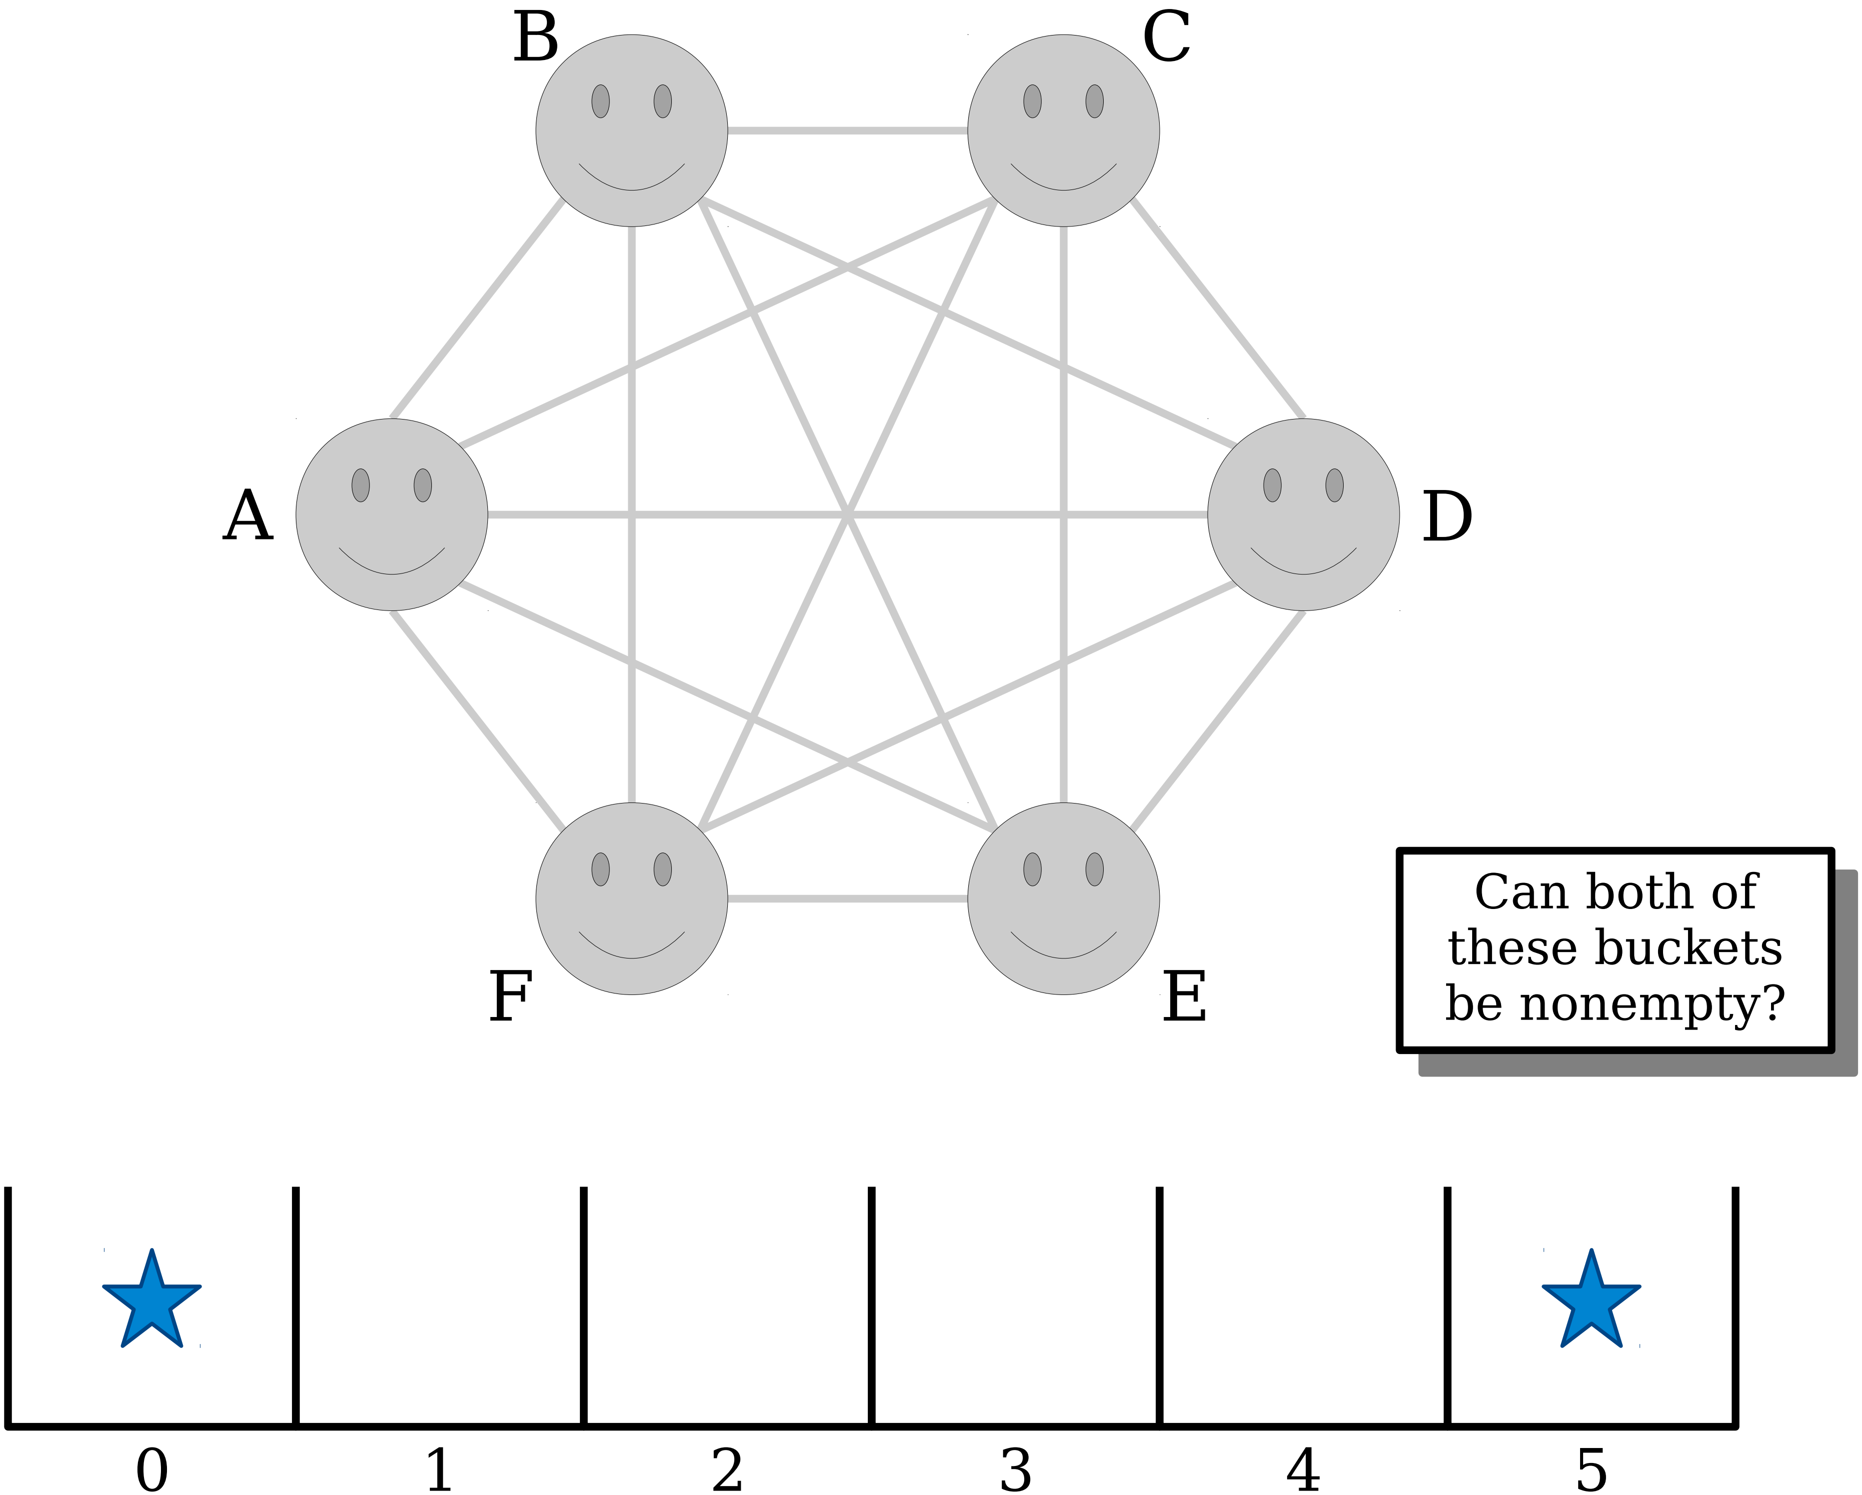
B
C
A
D
F
E
Can both of these buckets be nonempty?
0
1
2
3
4
5

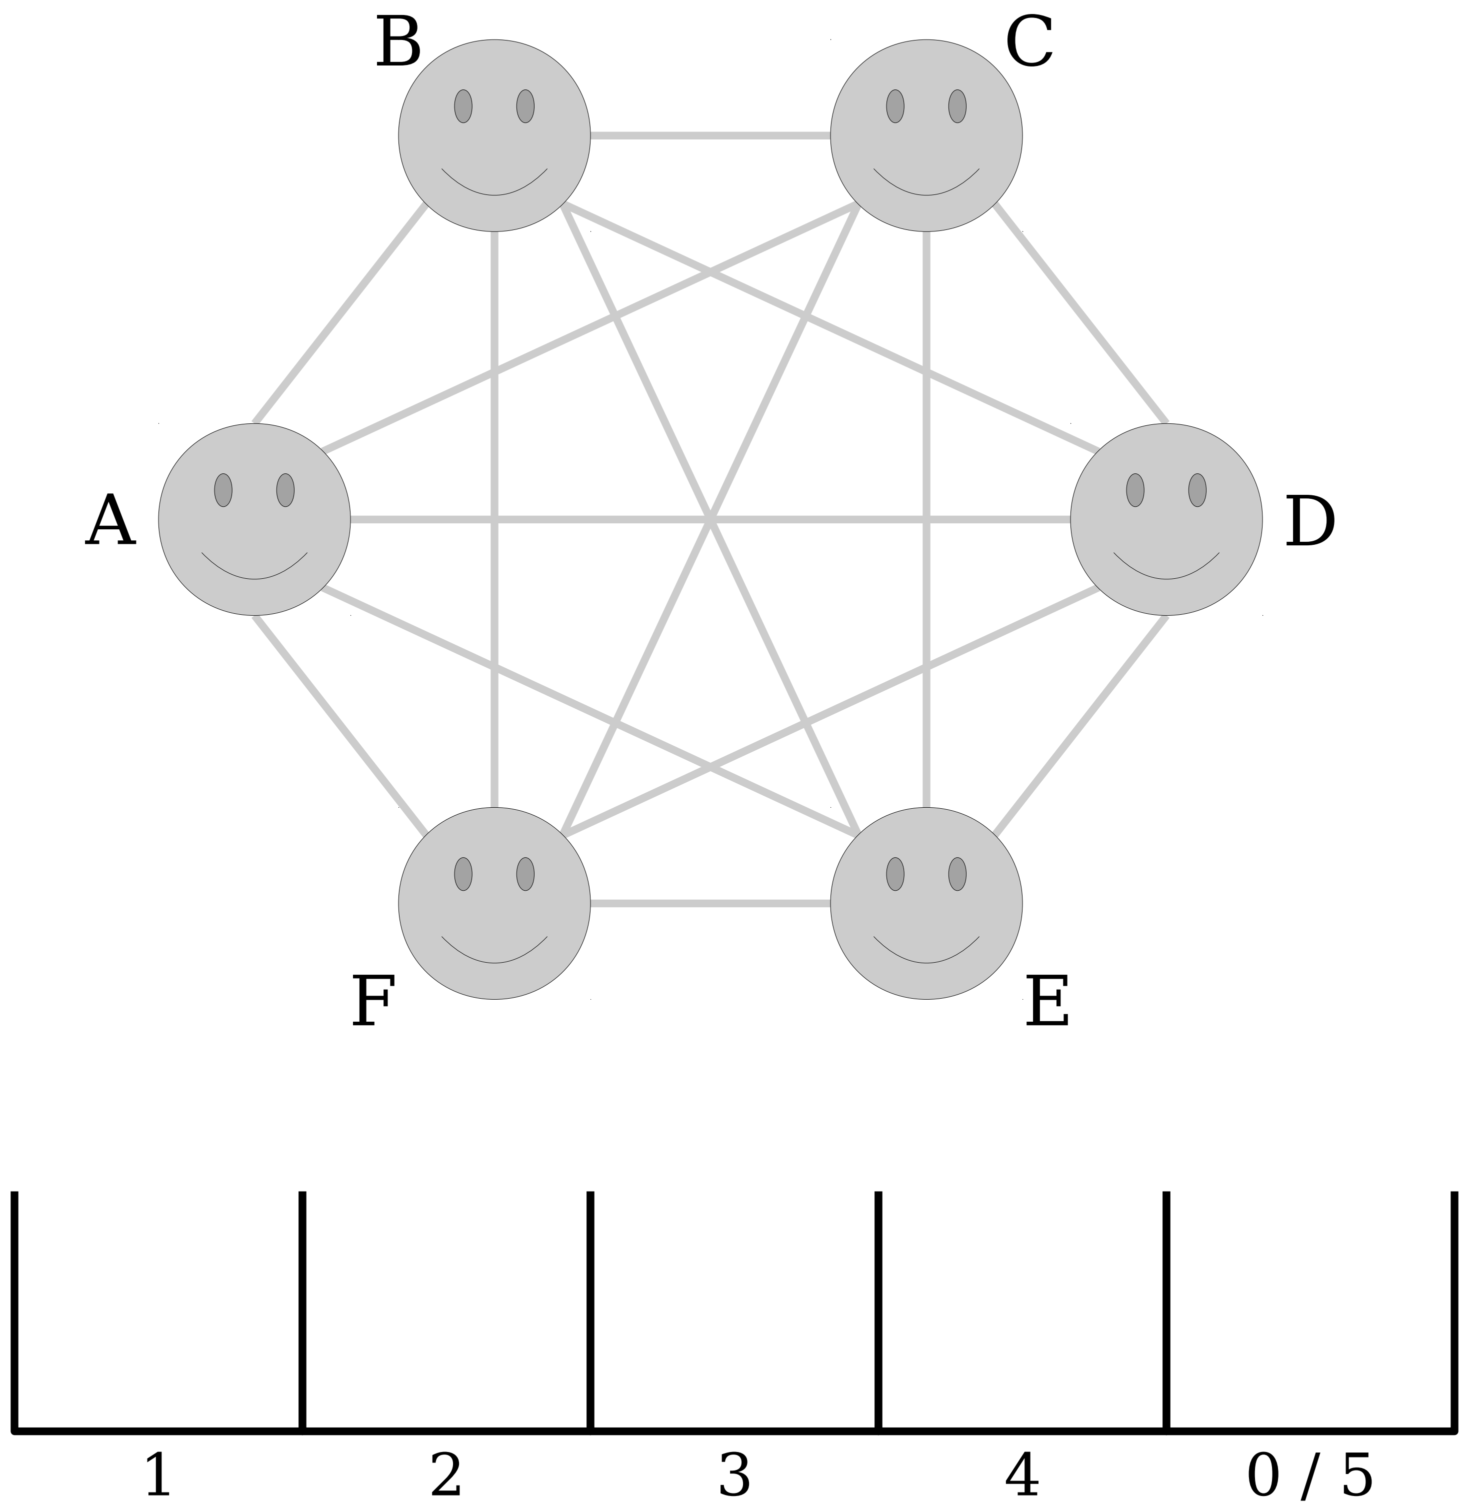
B
C
A
D
F
E
1
2
3
4
0 / 5

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

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

Proof 1:

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.

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, $0, 1, 2, \ldots,$ and $n - 1$.

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 $n - 1$:

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*.

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.)

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, …, $n - 1$.

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 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

The Pigeonhole Principle

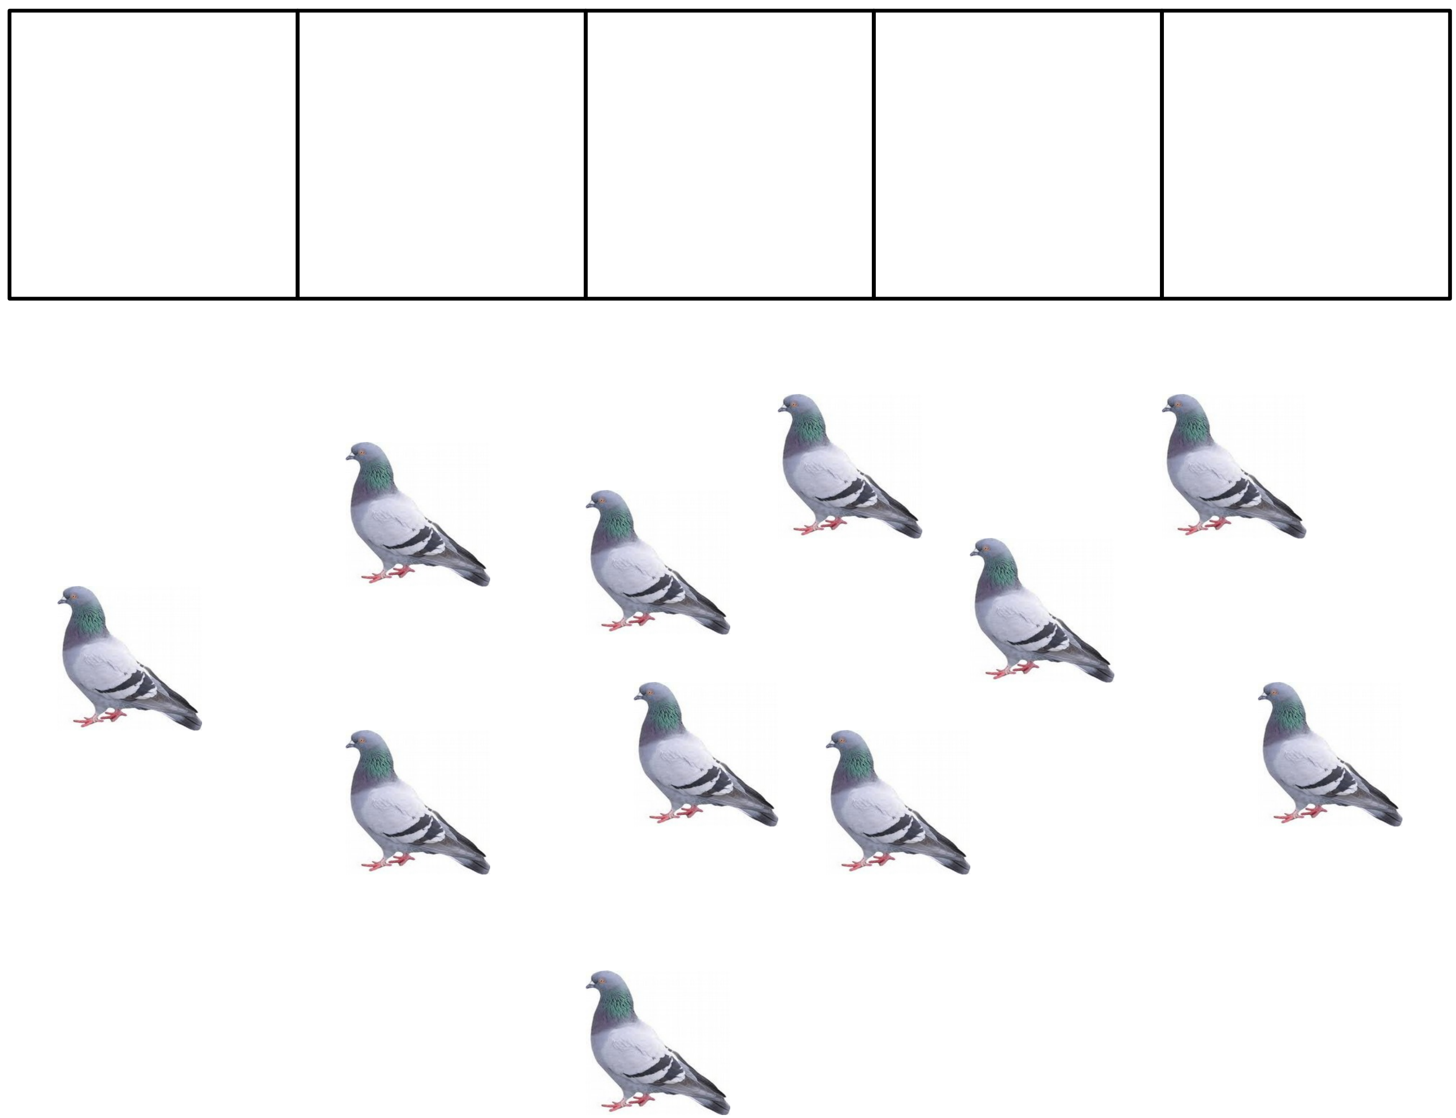

The Pigeonhole Principle

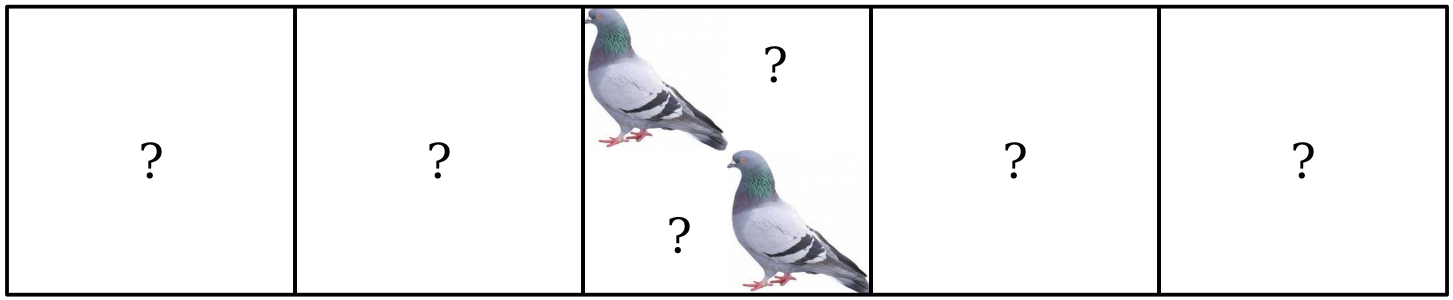

The Pigeonhole Principle

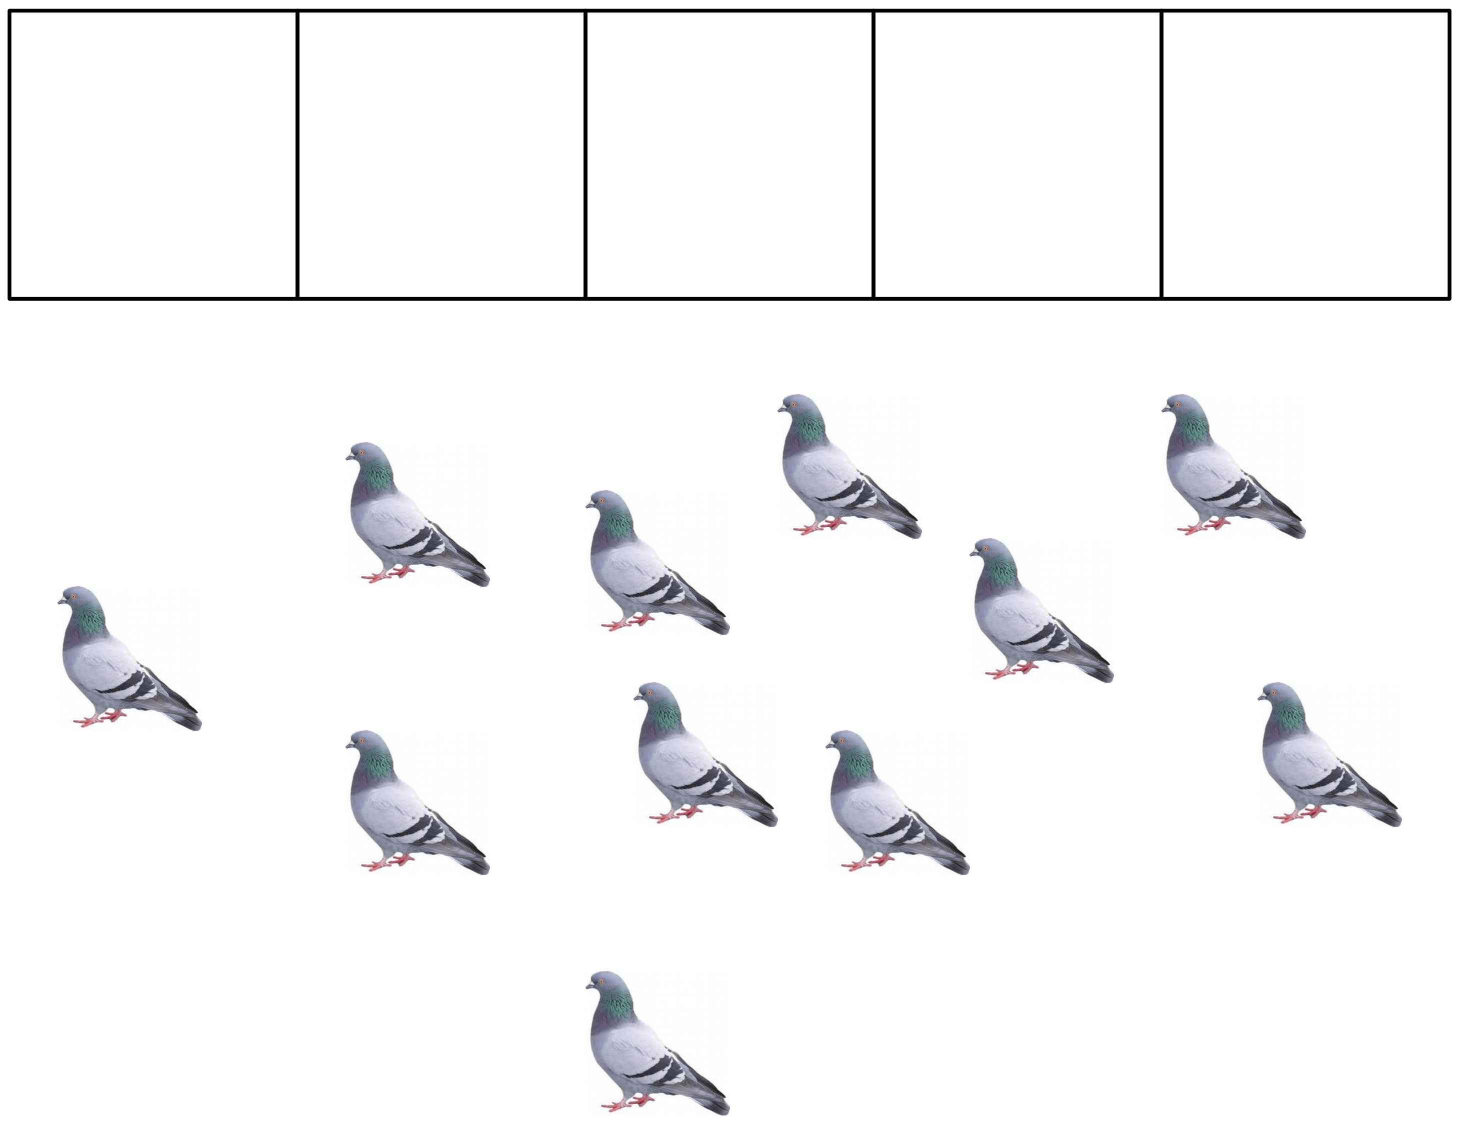

The Pigeonhole Principle

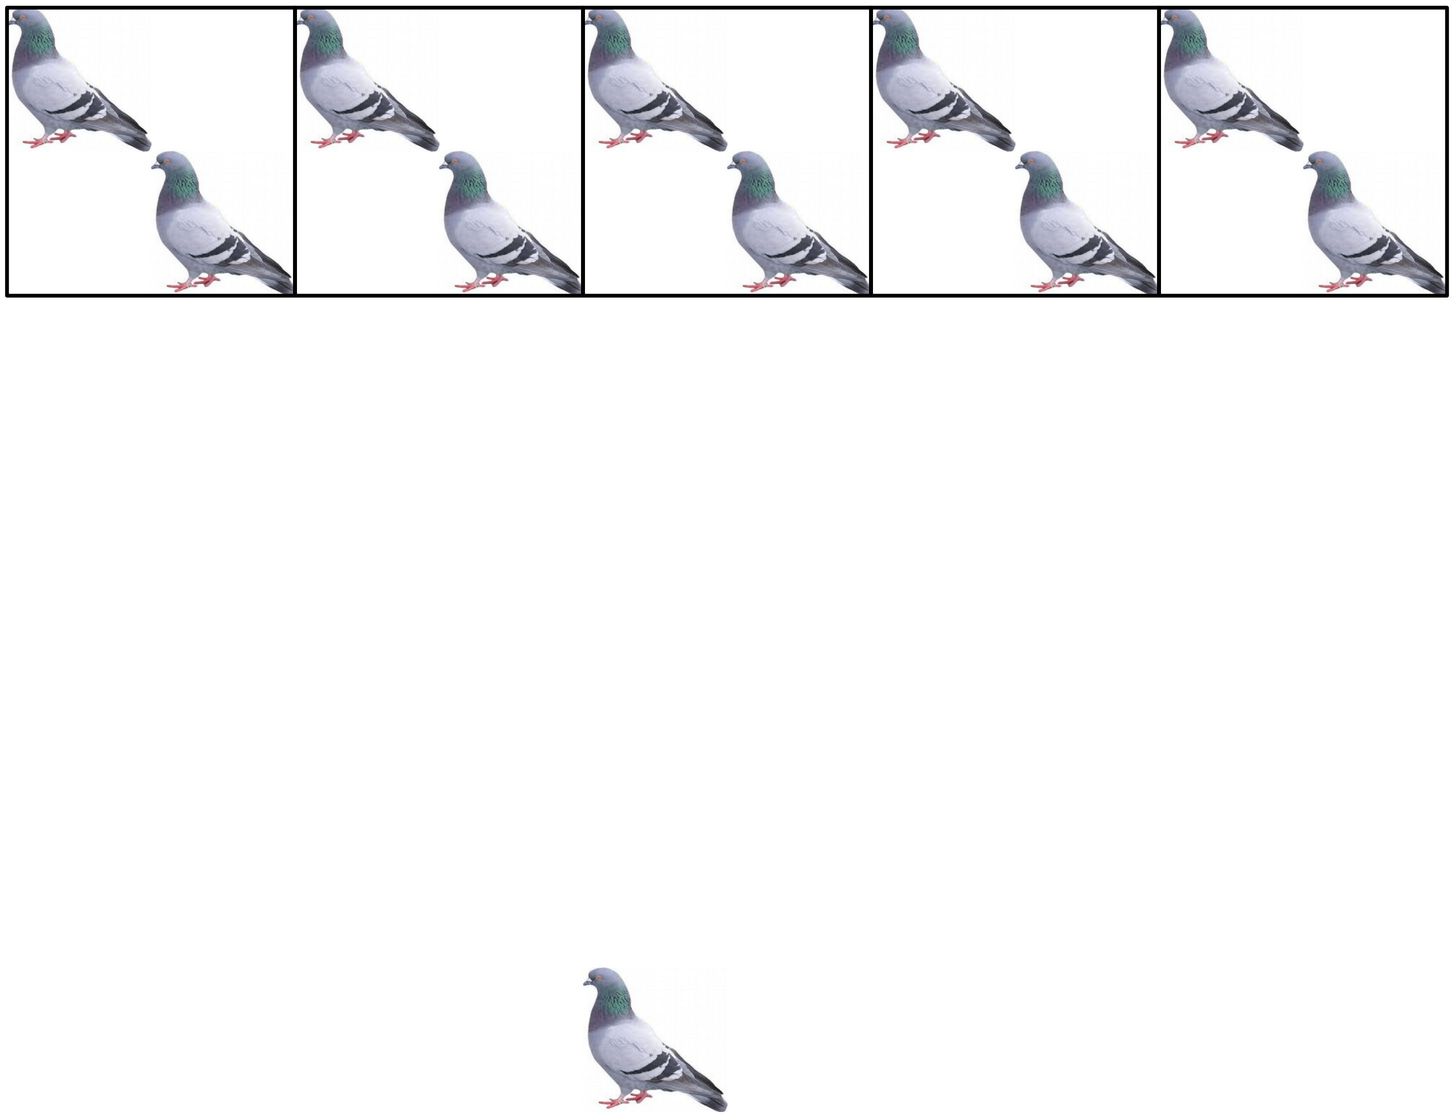

The Pigeonhole Principle

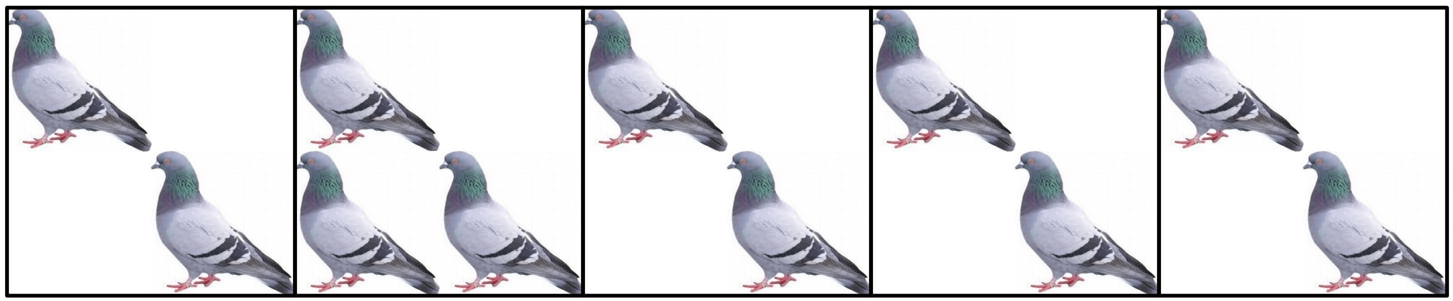

The Pigeonhole Principle

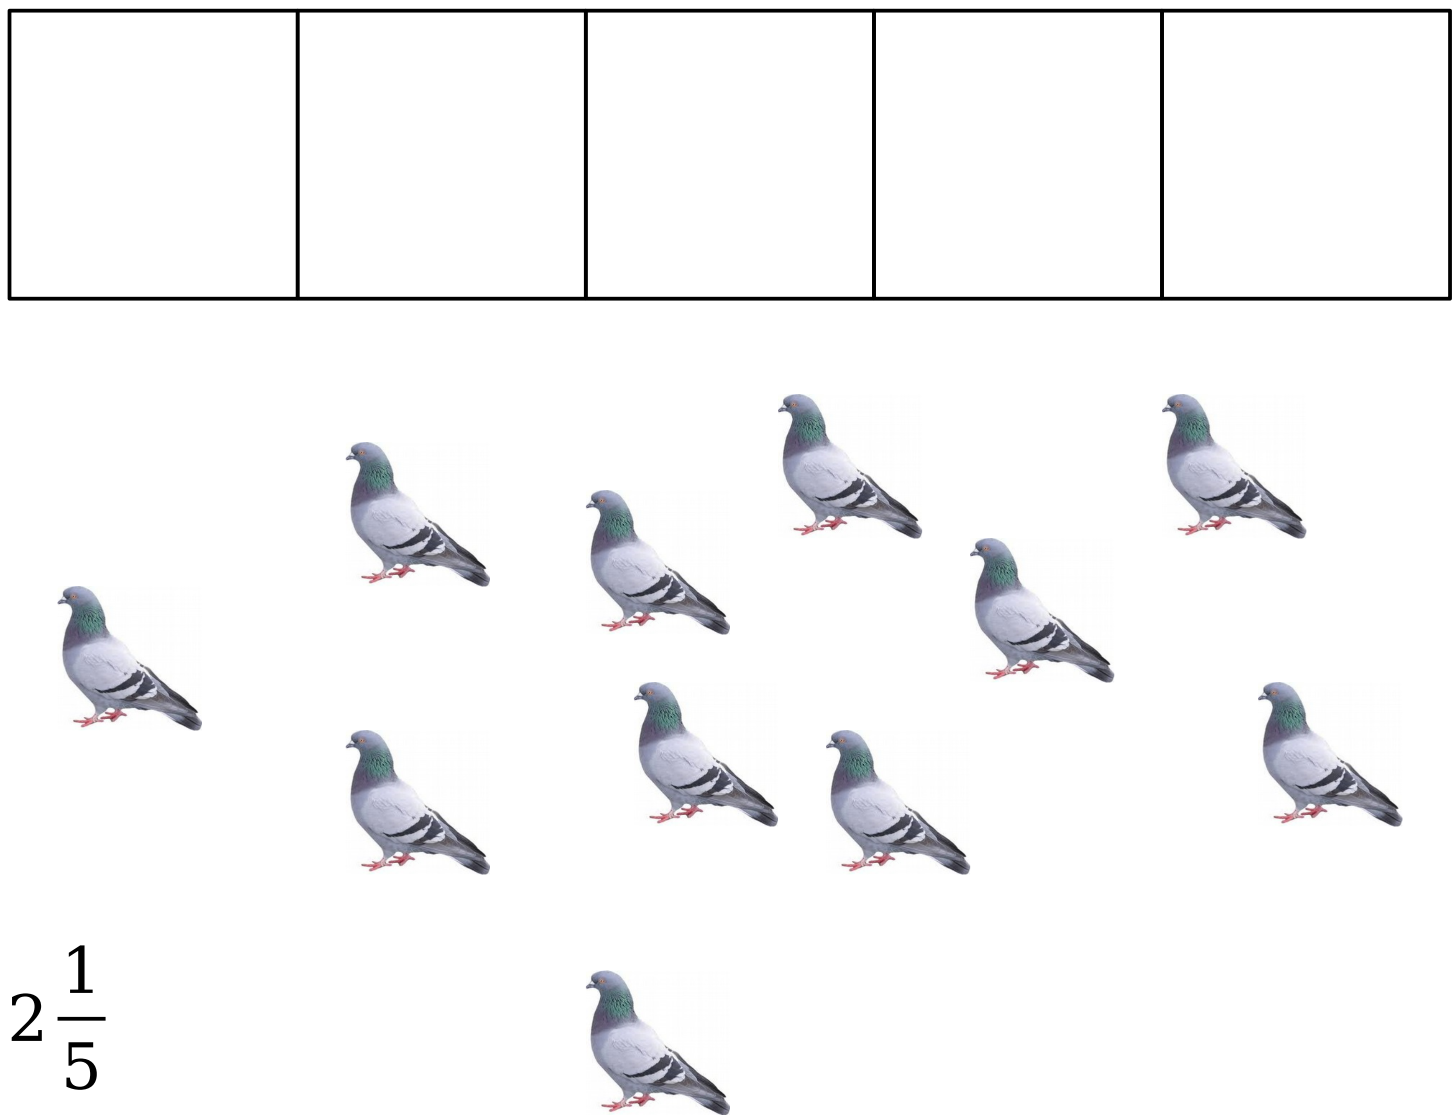

$$\frac{11}{5} = 2\frac{1}{5}$$

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."

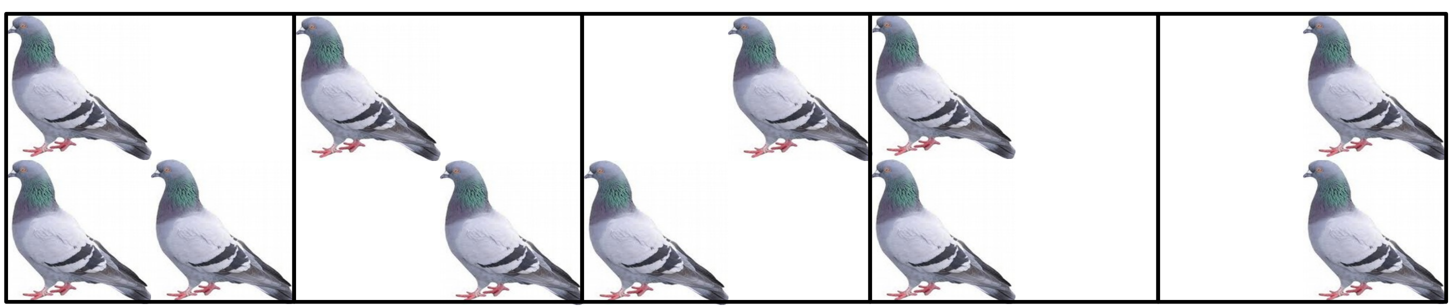

$m = 11$
$n = 5$

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

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."

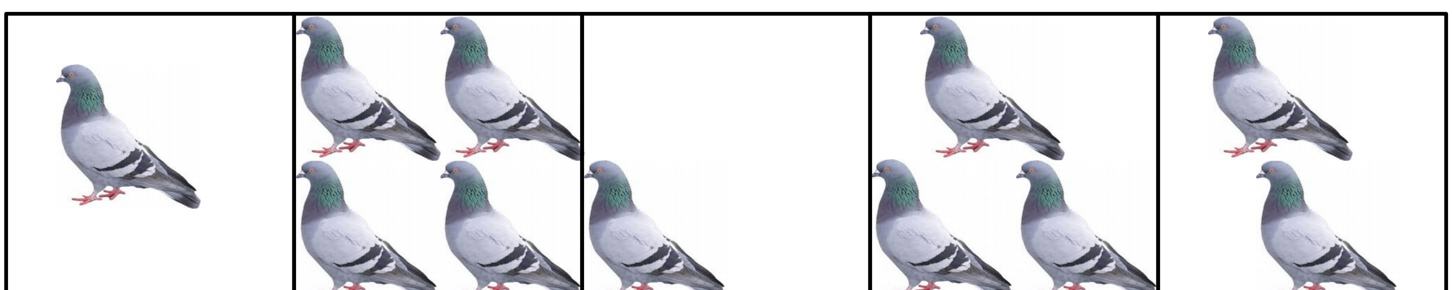

$m = 11$
$n = 5$

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

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."

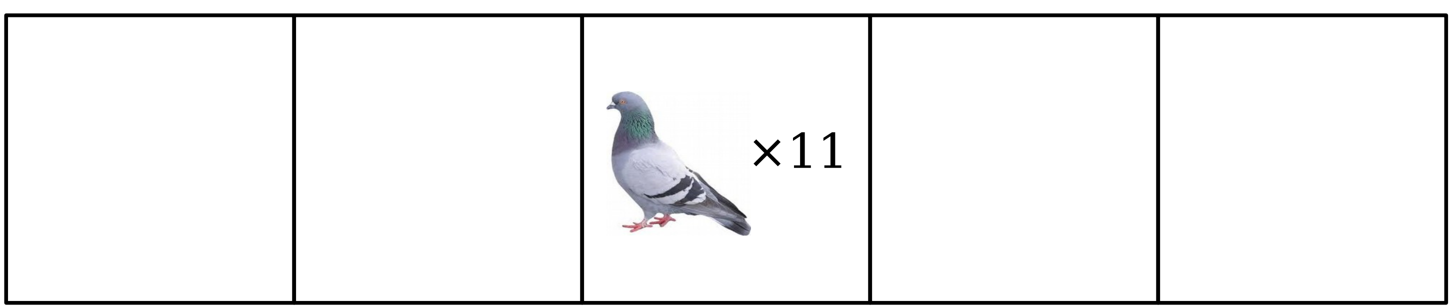

$m = 11$
$n = 5$

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

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).

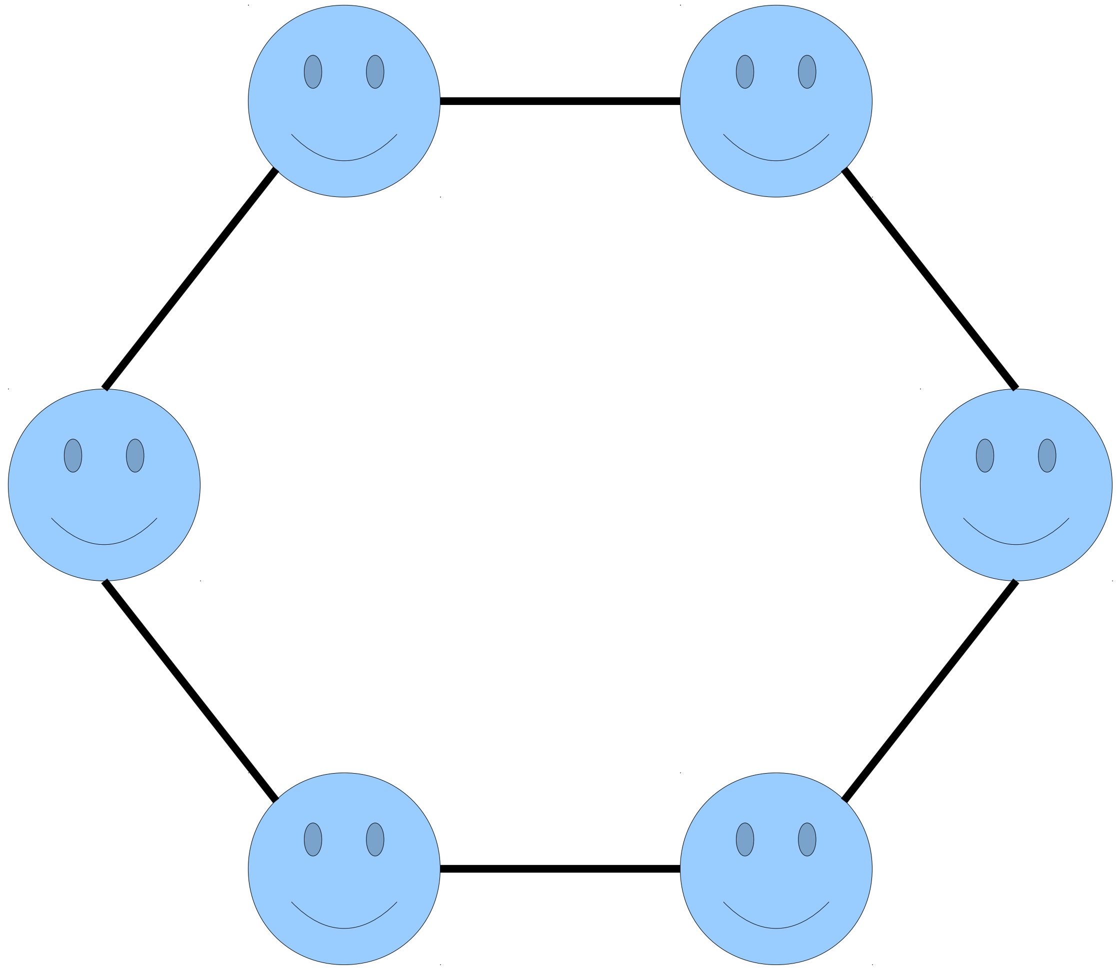

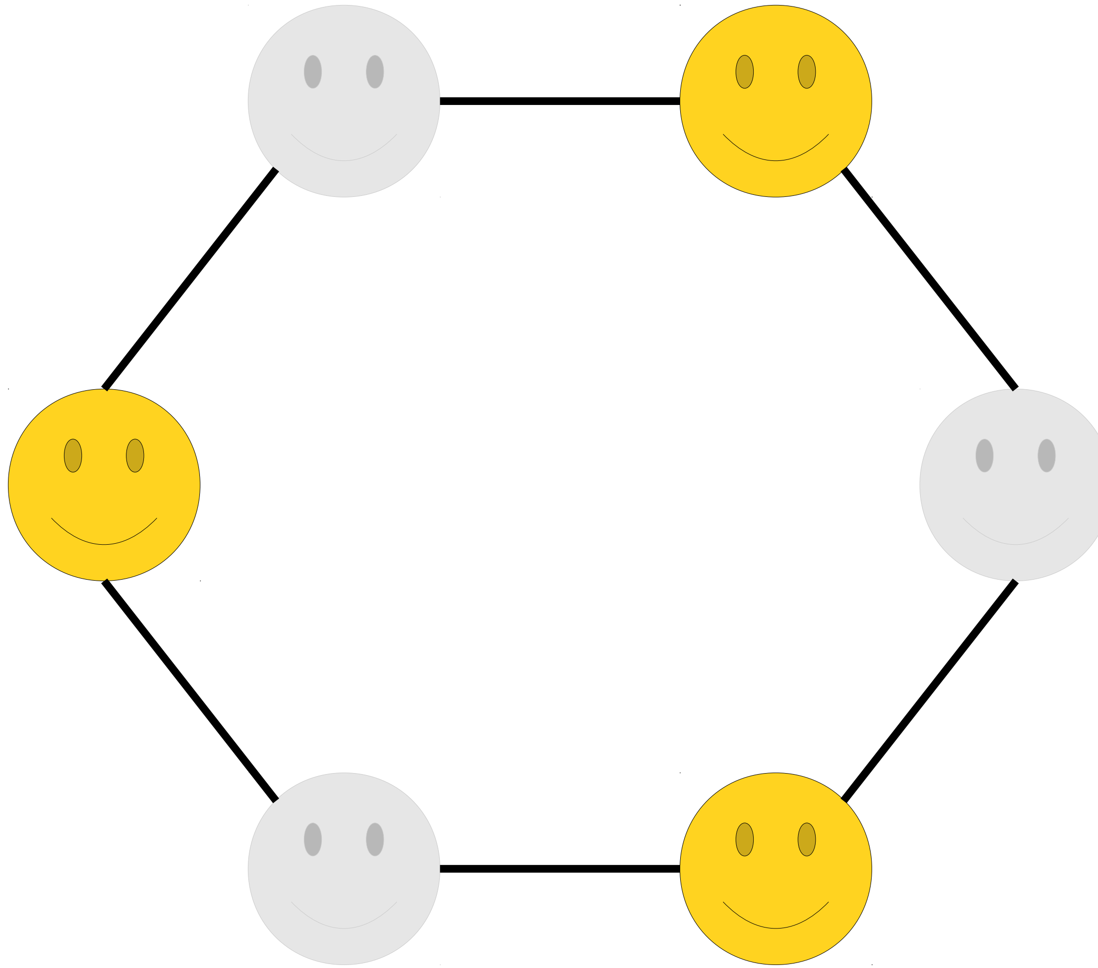

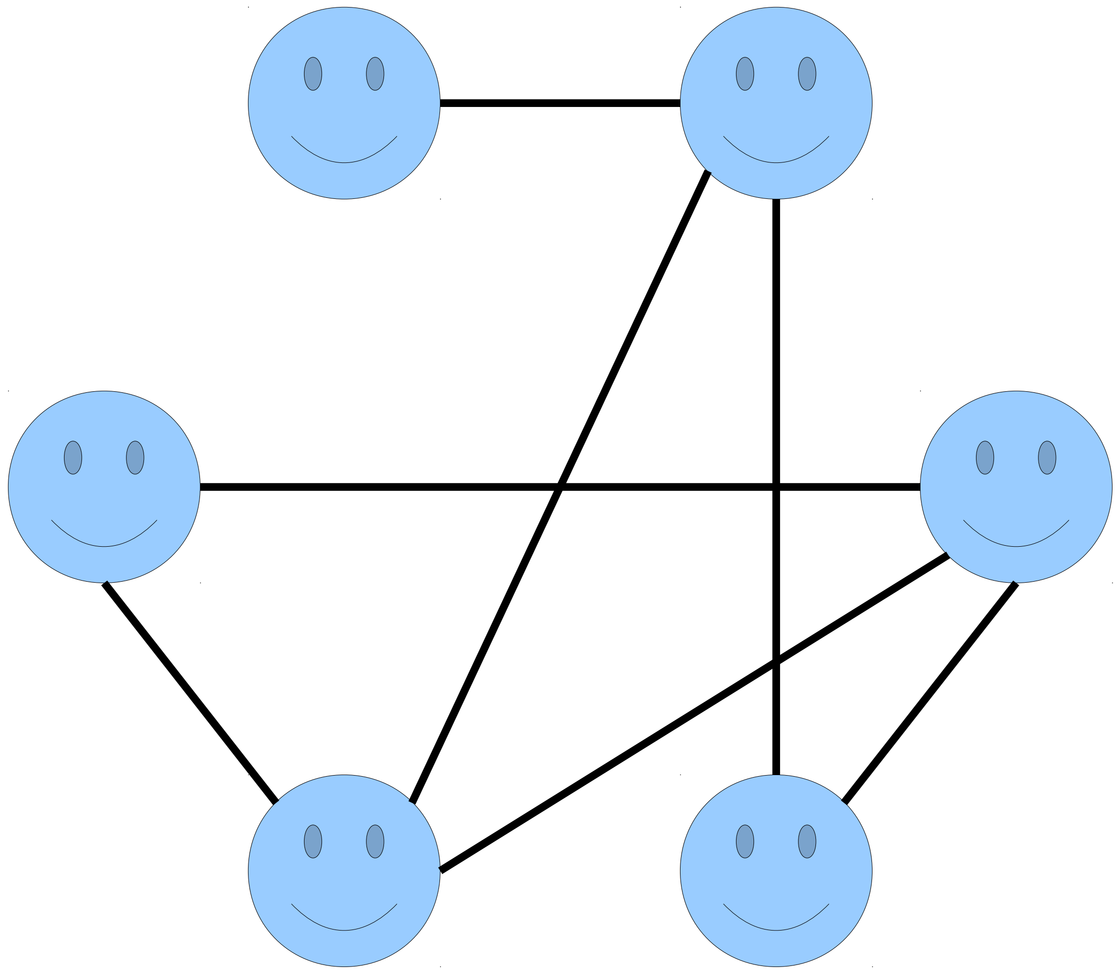

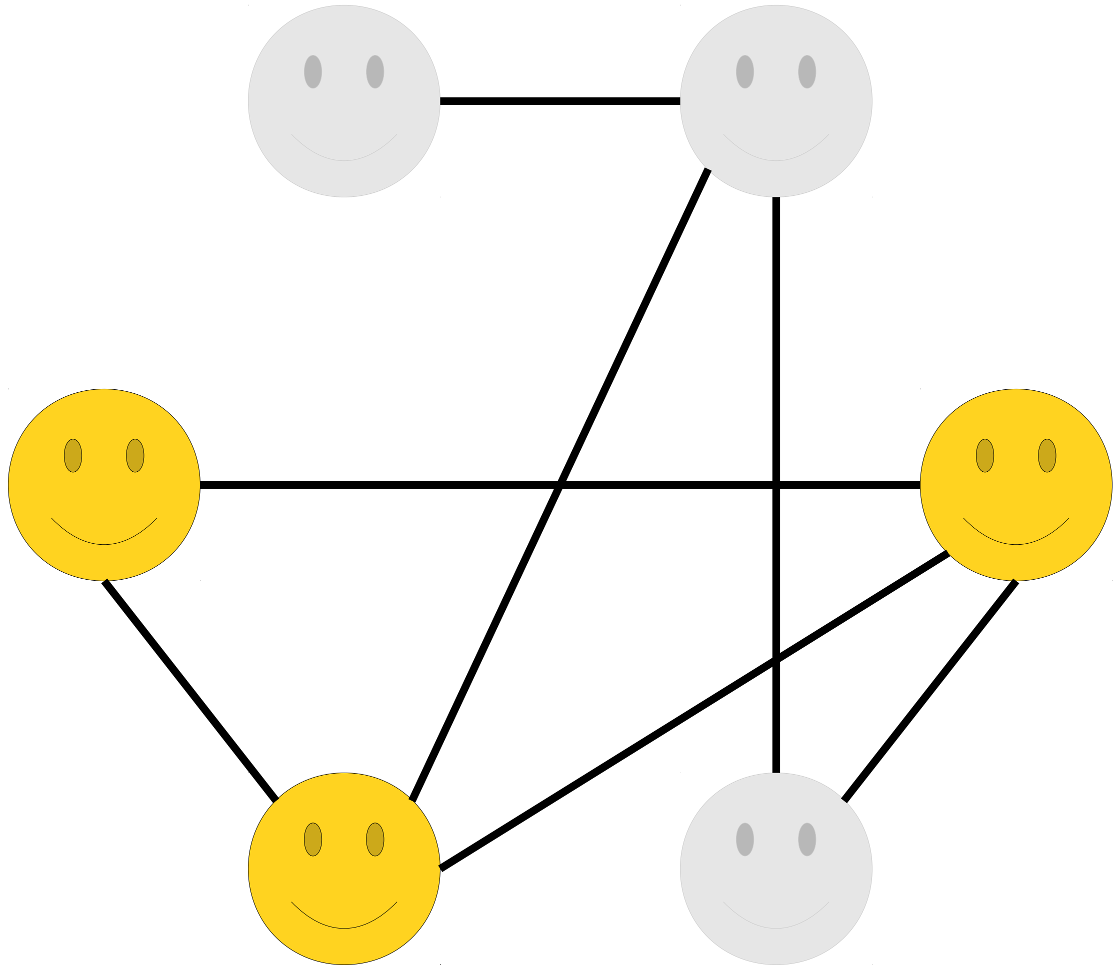

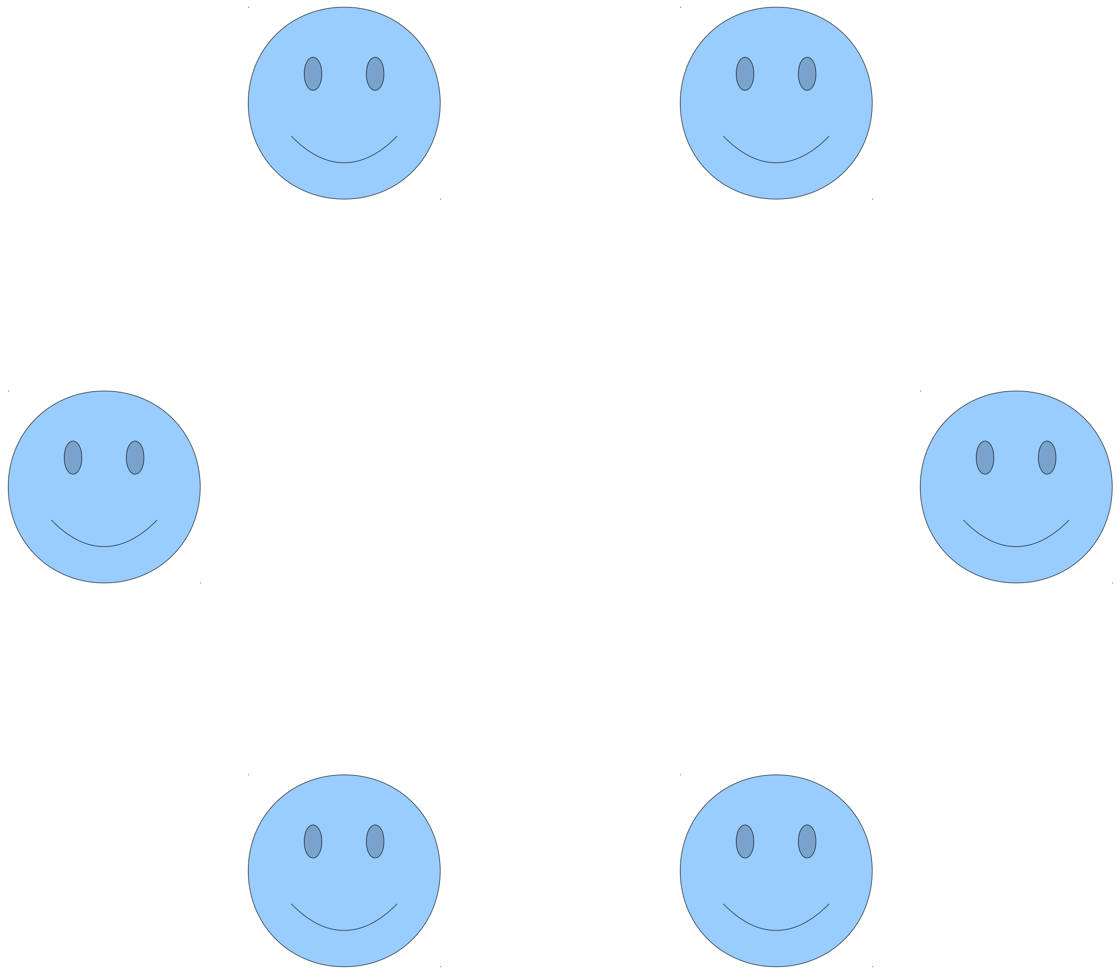

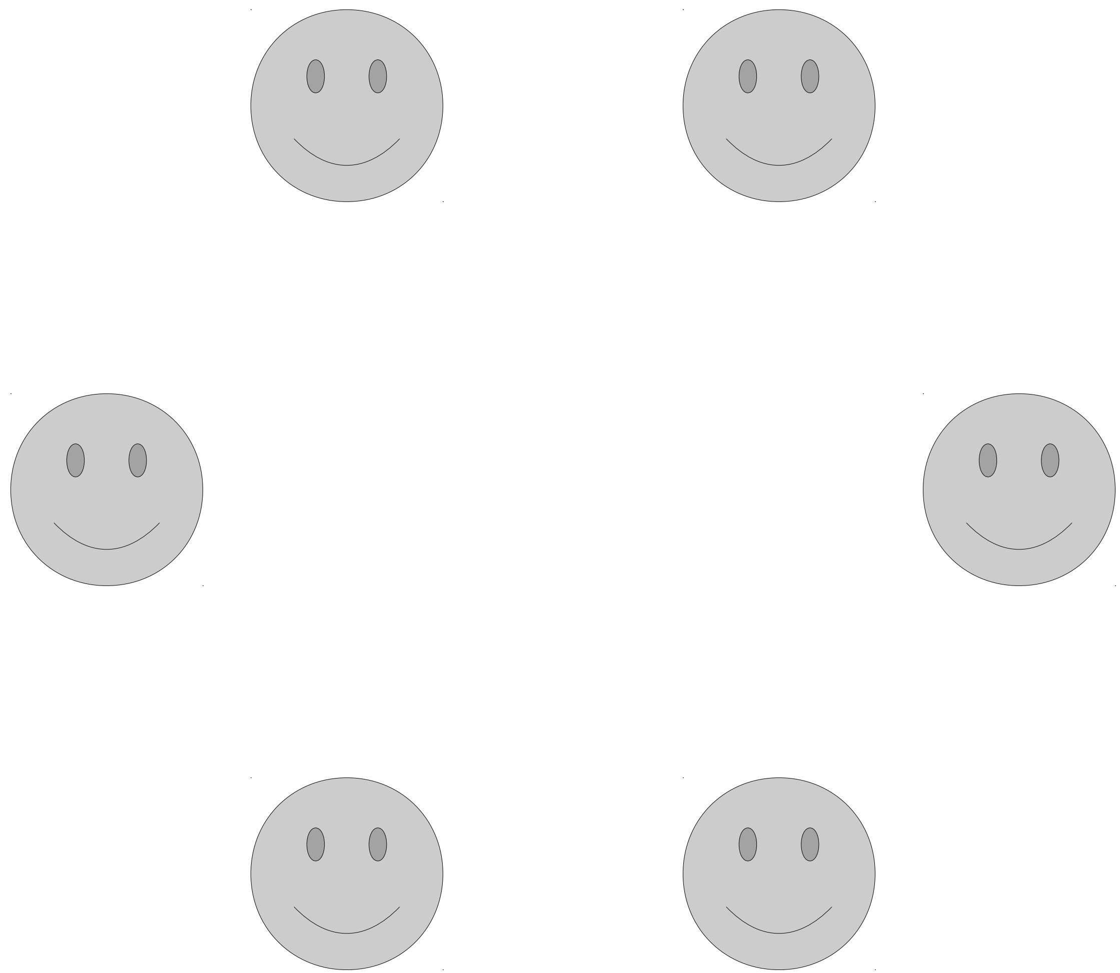

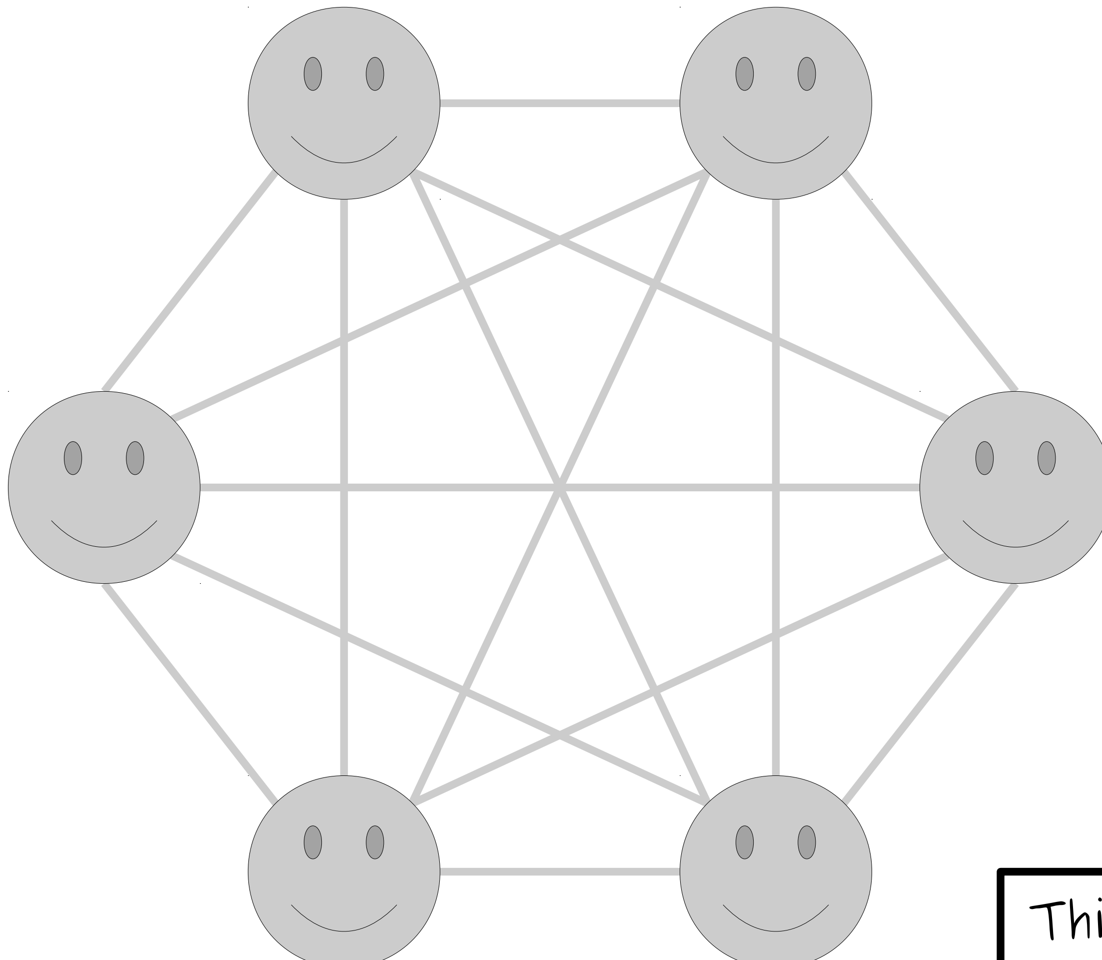

This graph is called a *6-clique*, by the way.

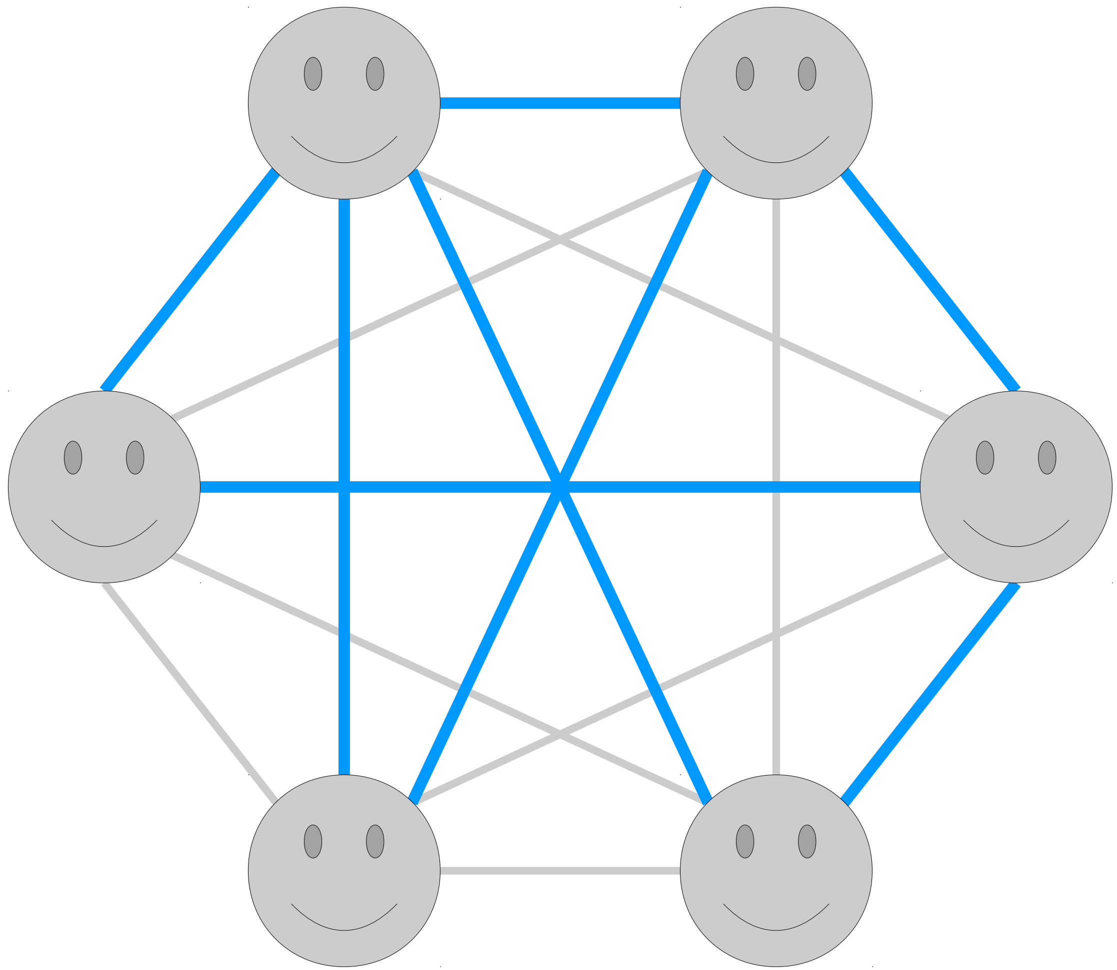

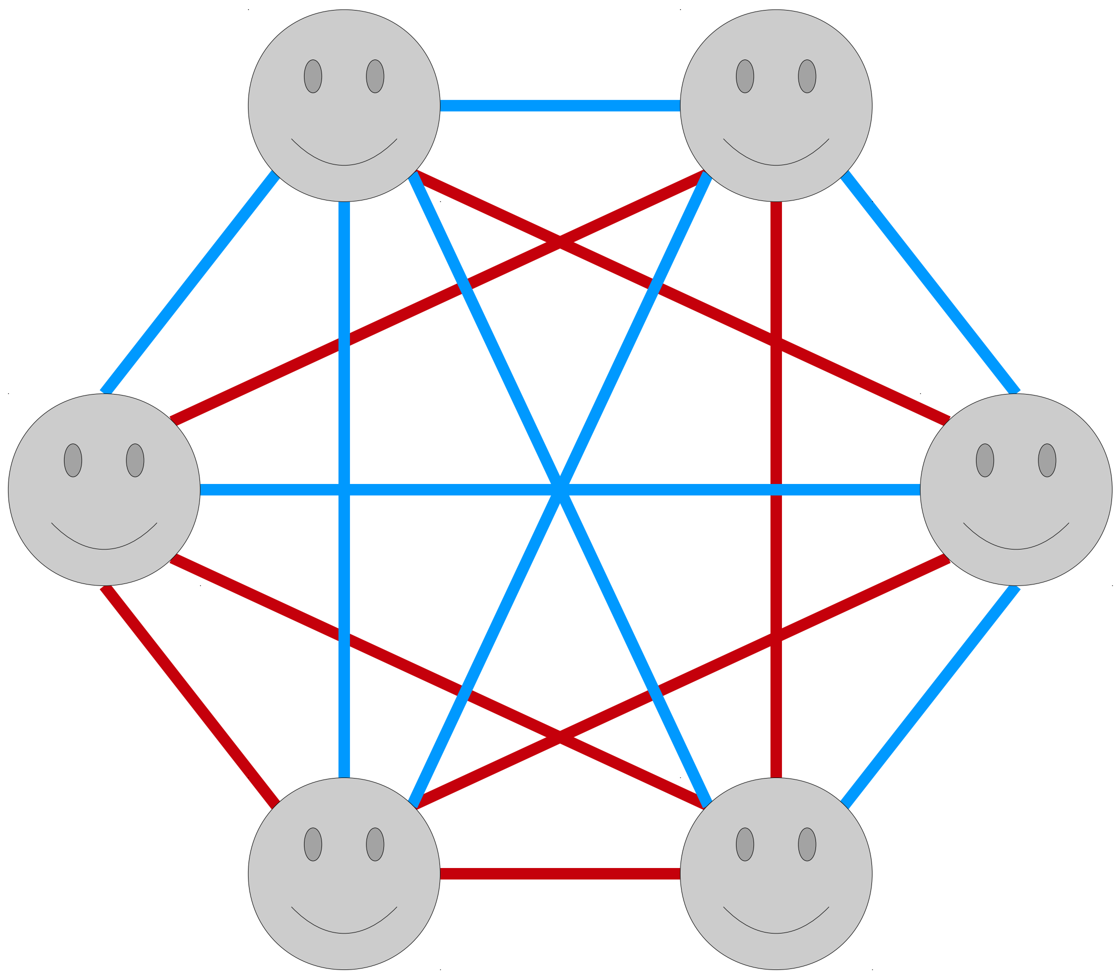

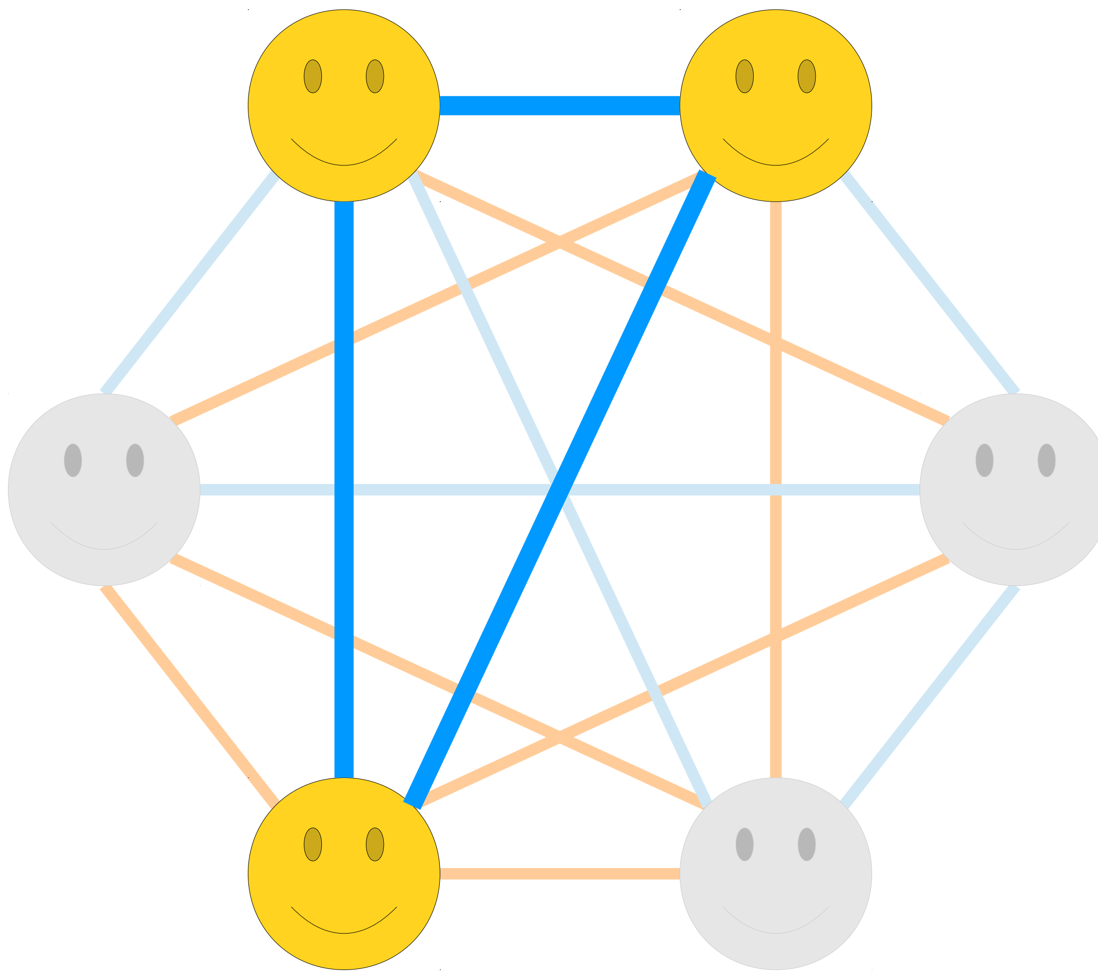

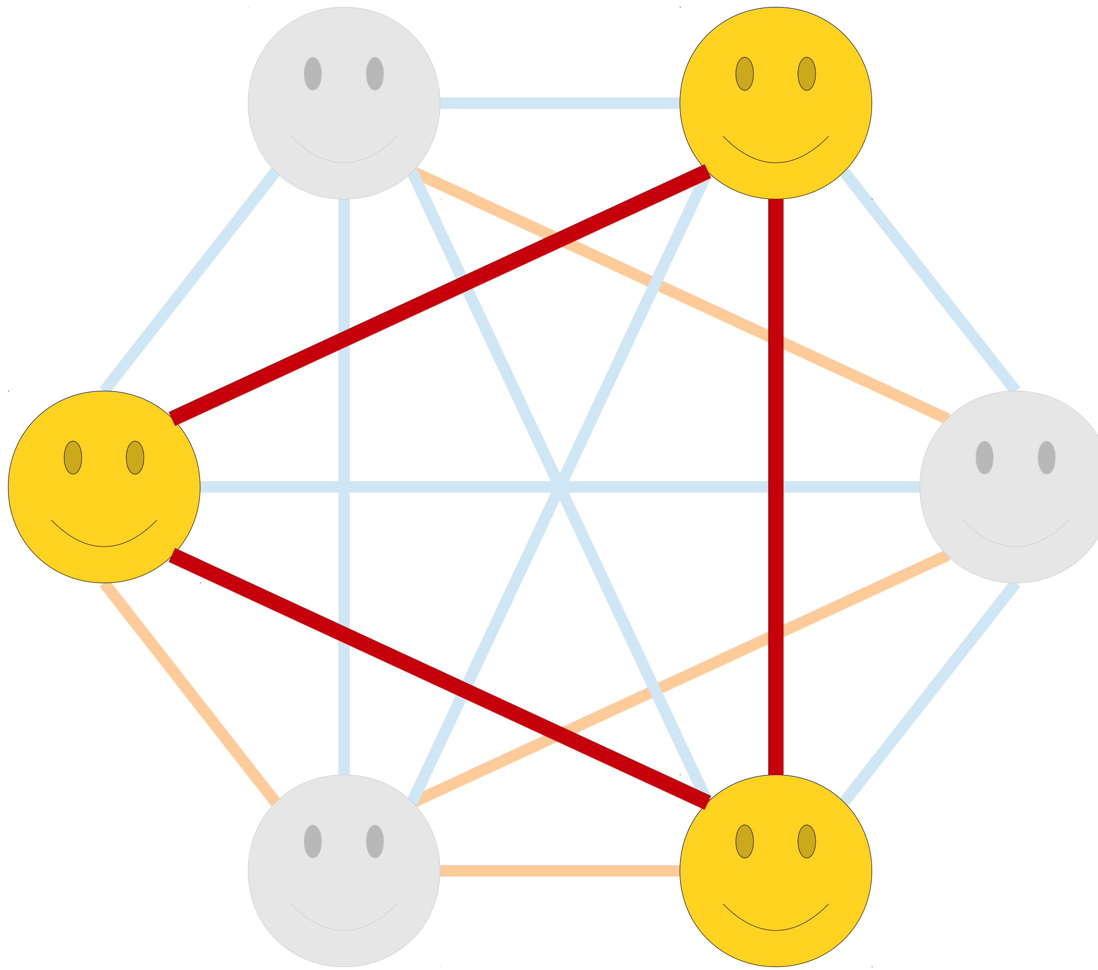

Friends and Strangers Restated

- 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?

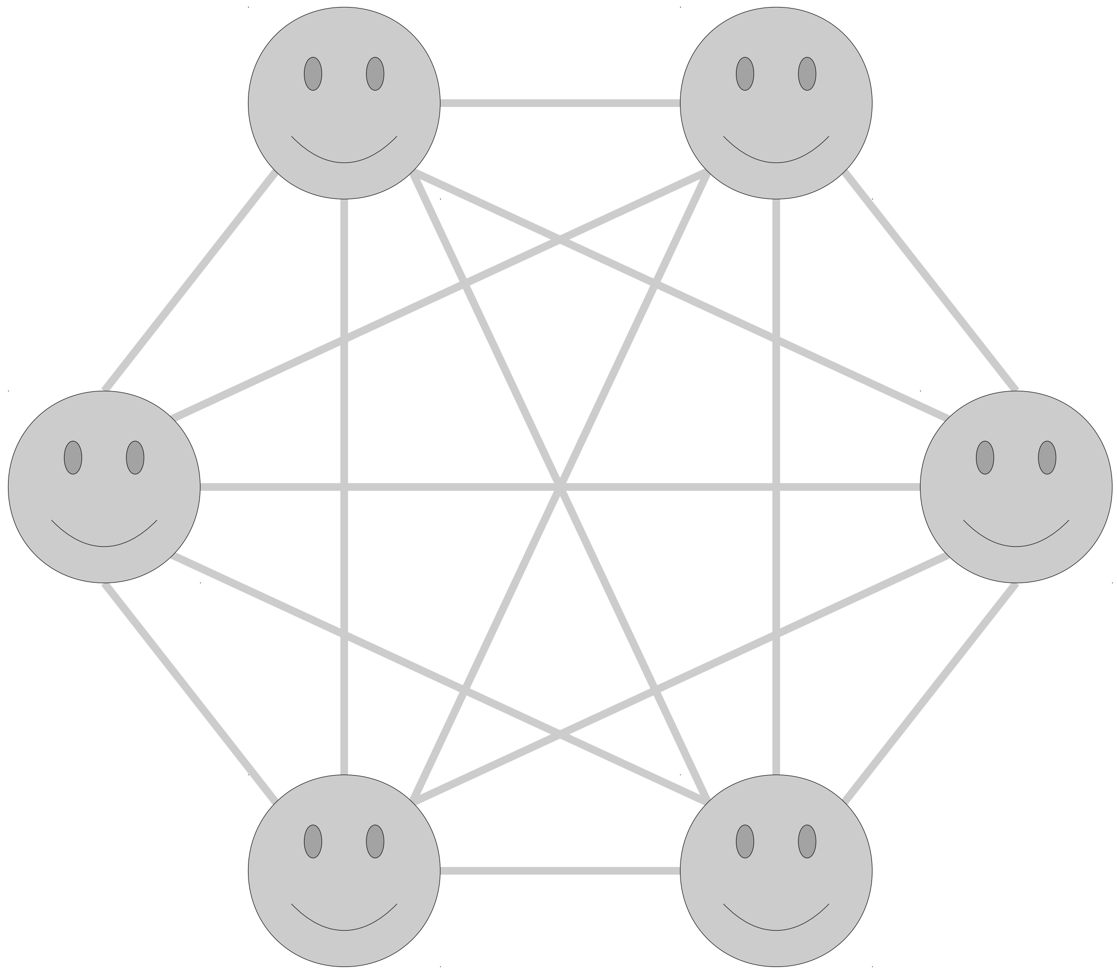

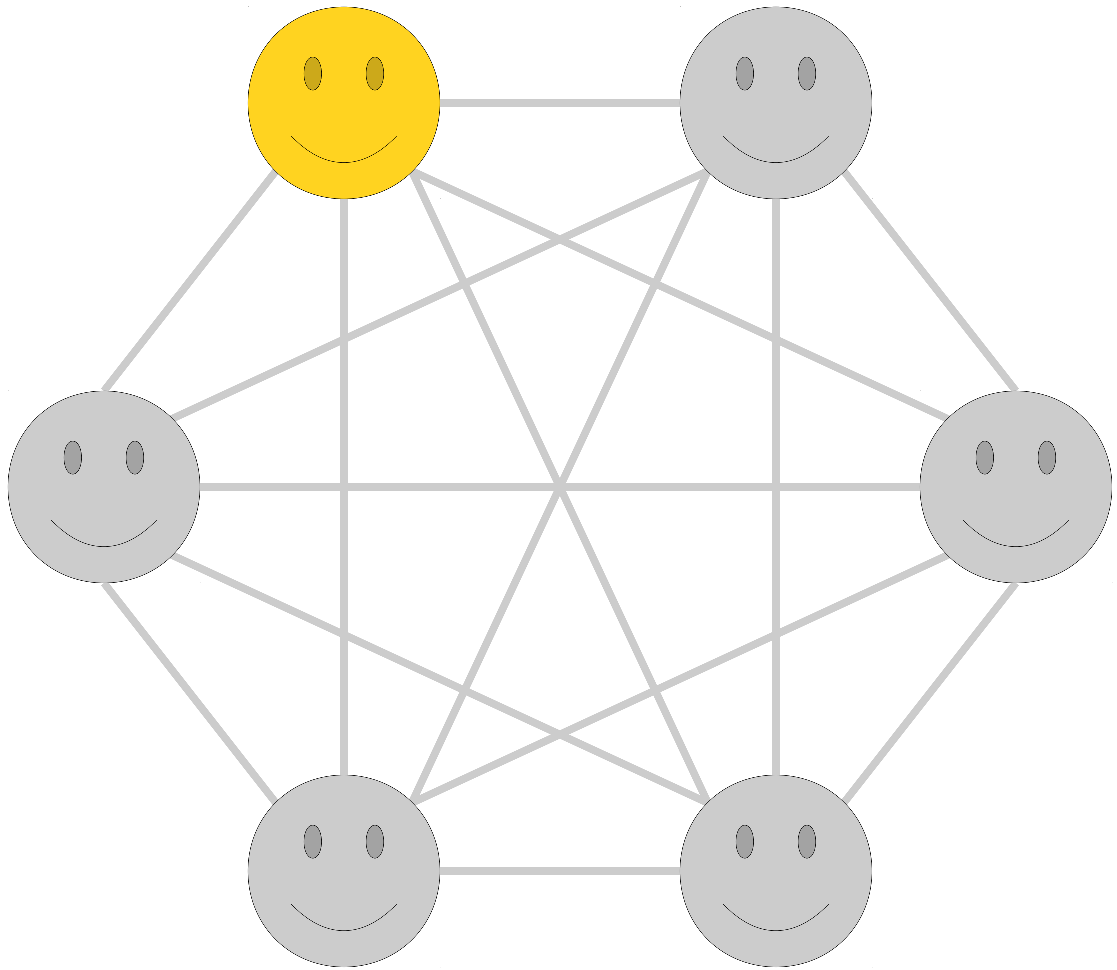

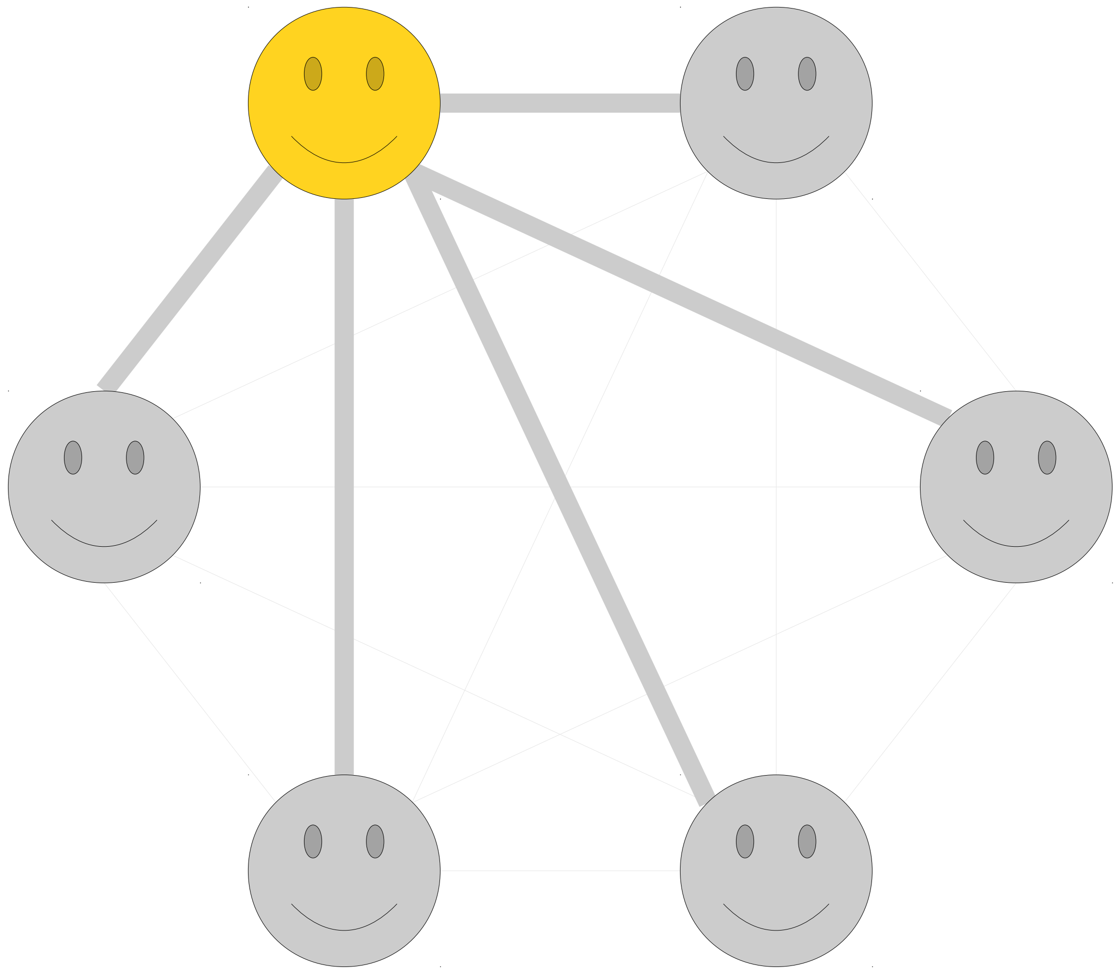

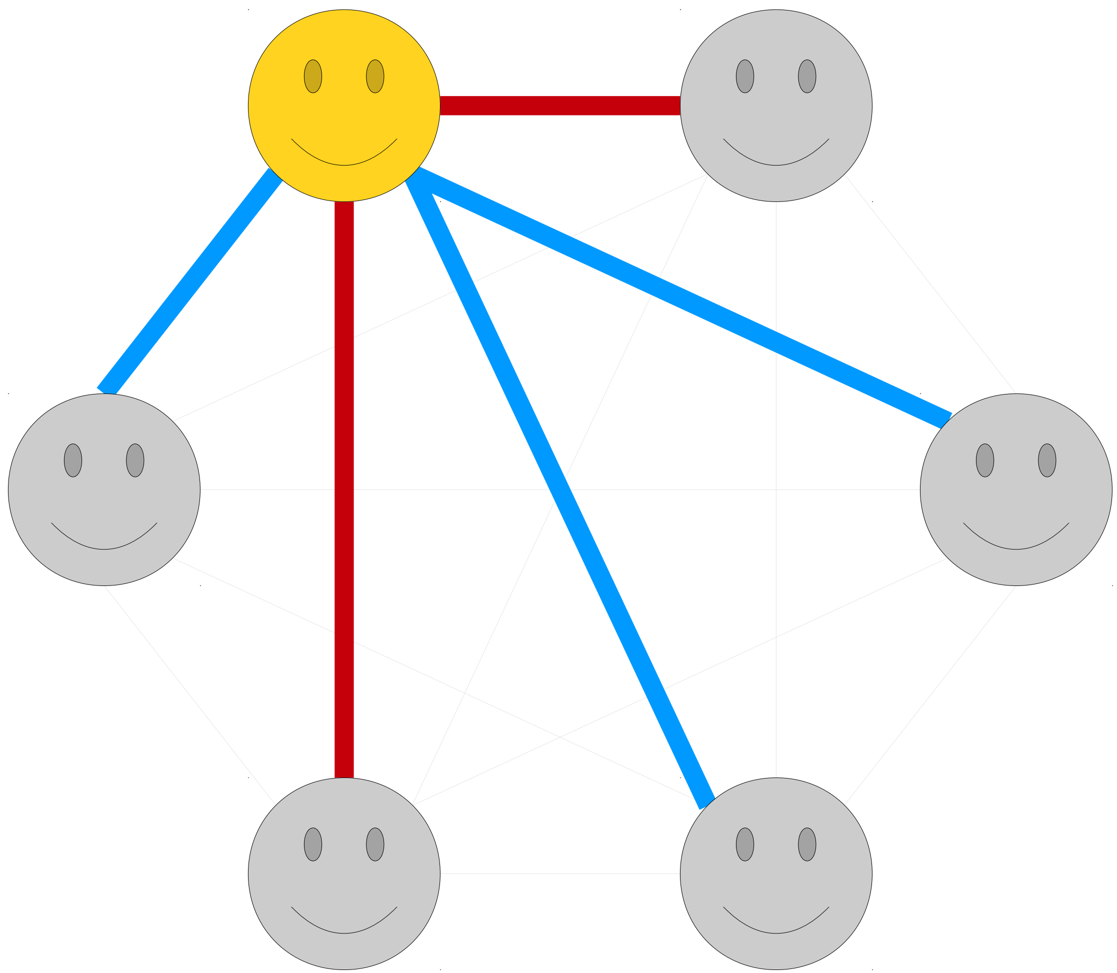

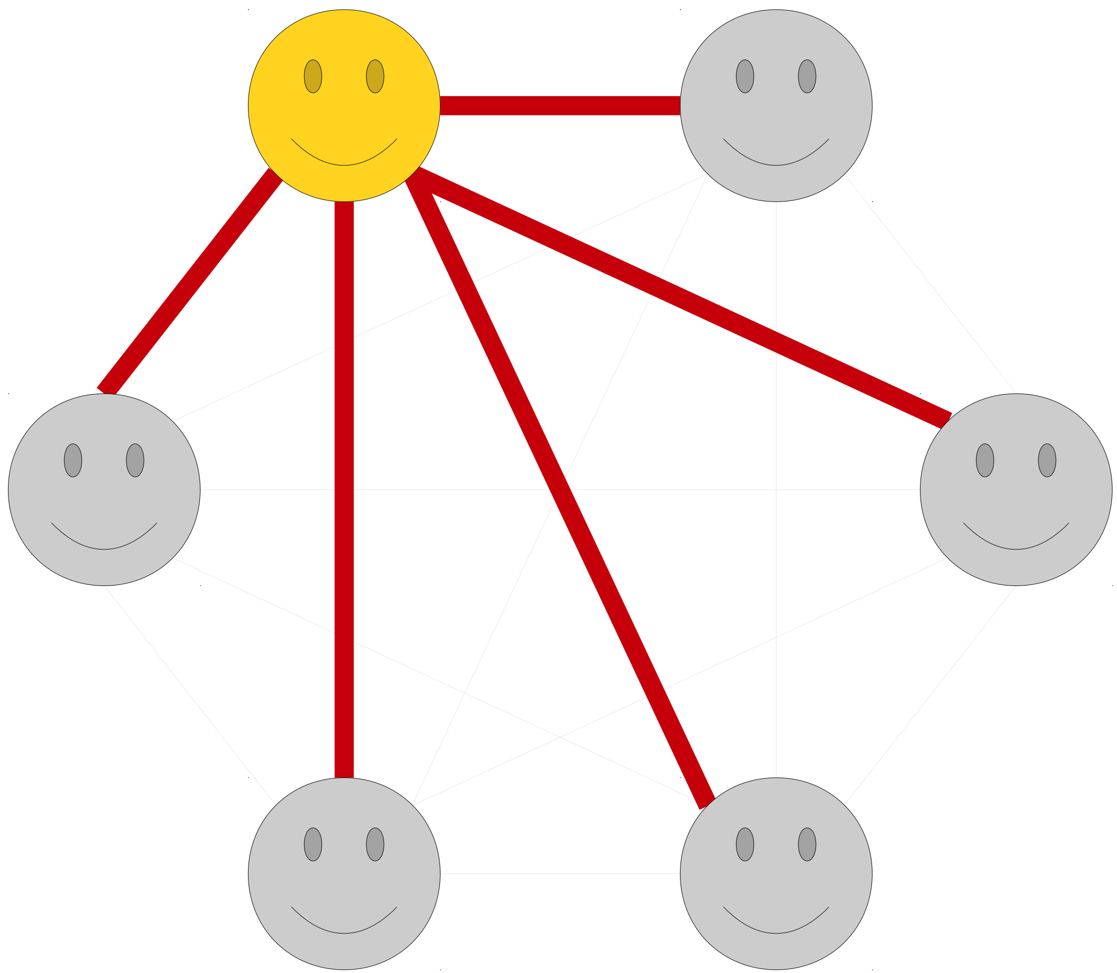

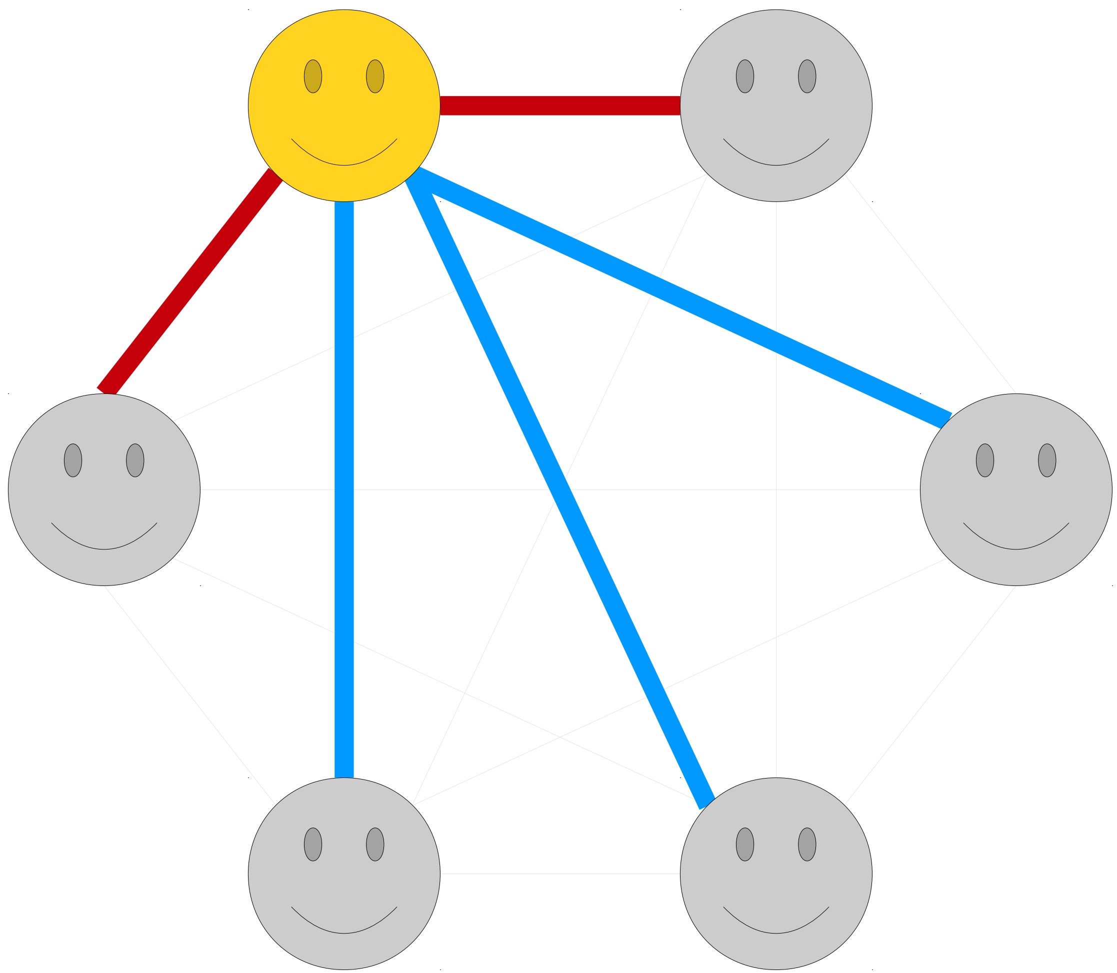

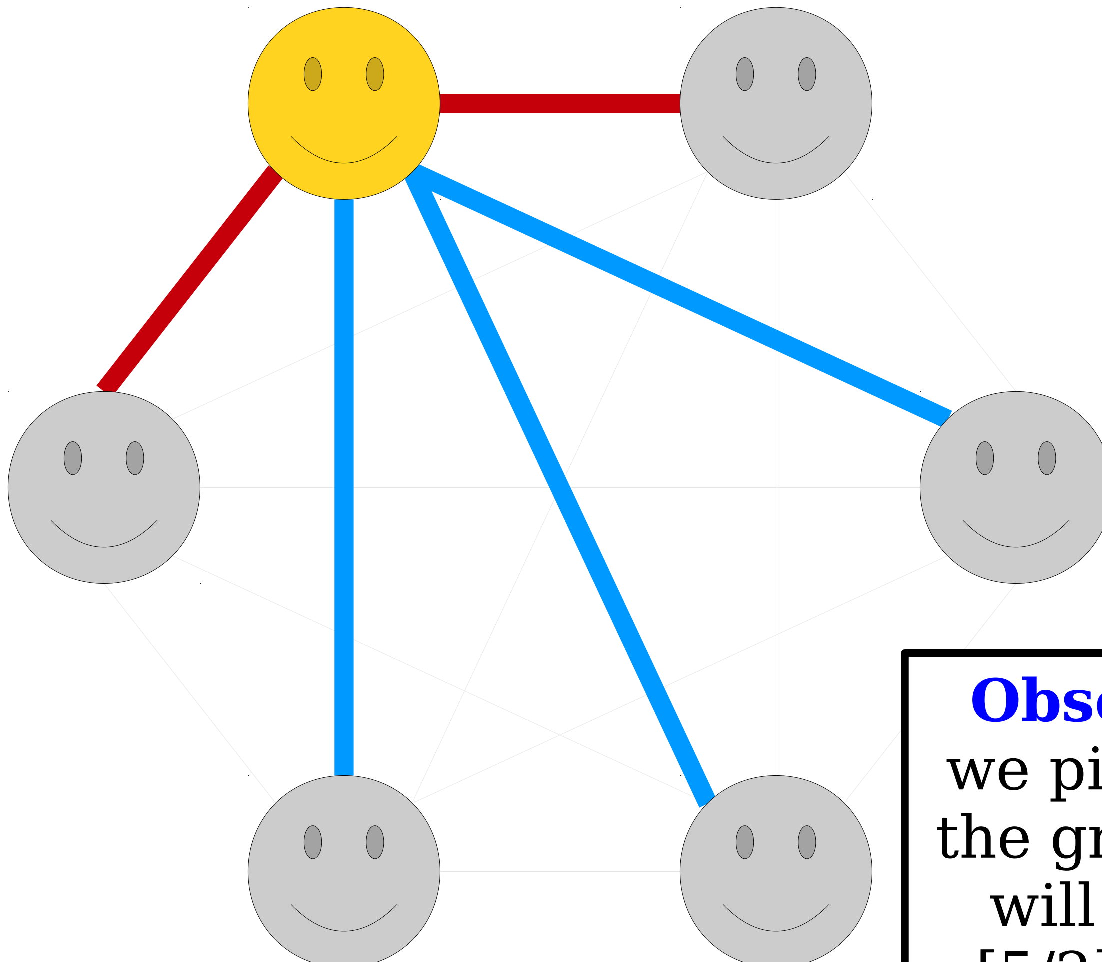

Observation 1: If we pick any node in the graph, that node will have at least ⌈5/2⌉ = 3 edges of the same color incident to it.

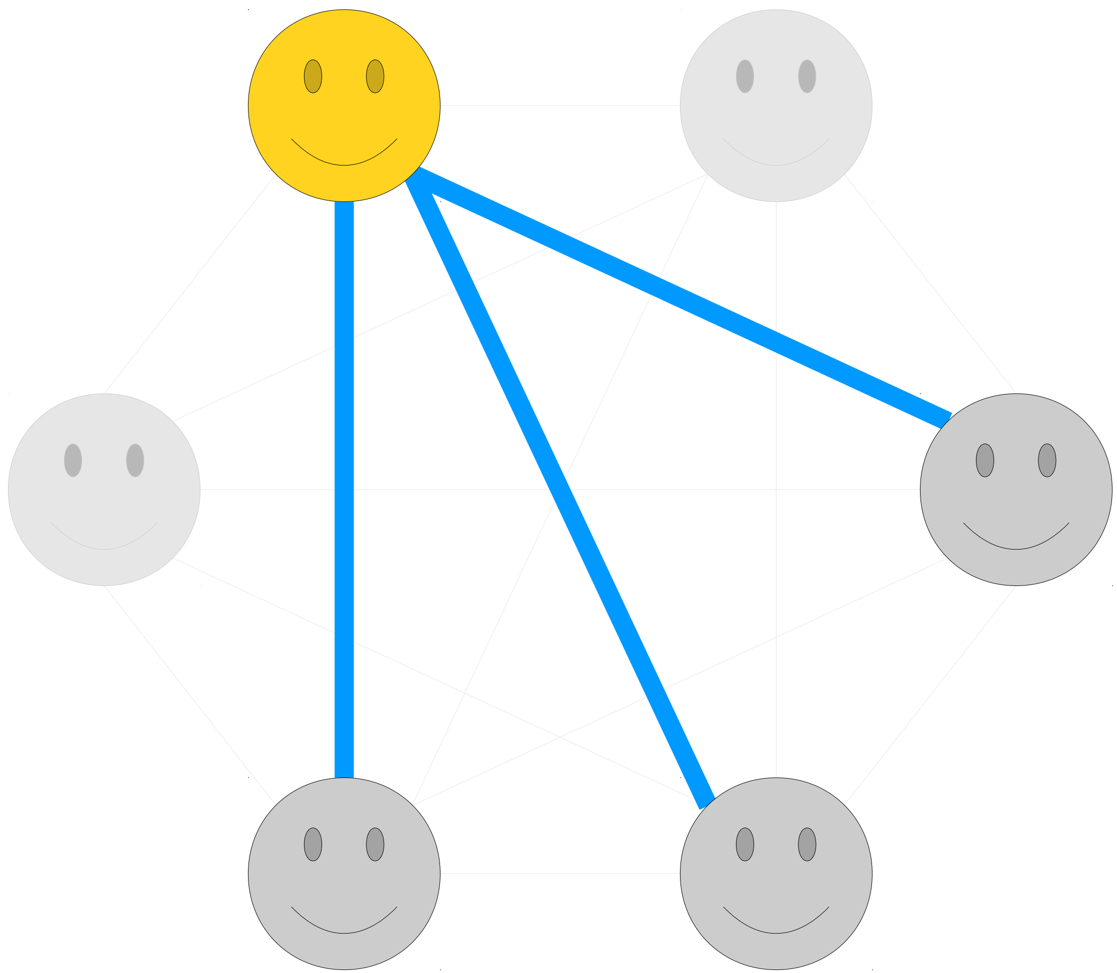

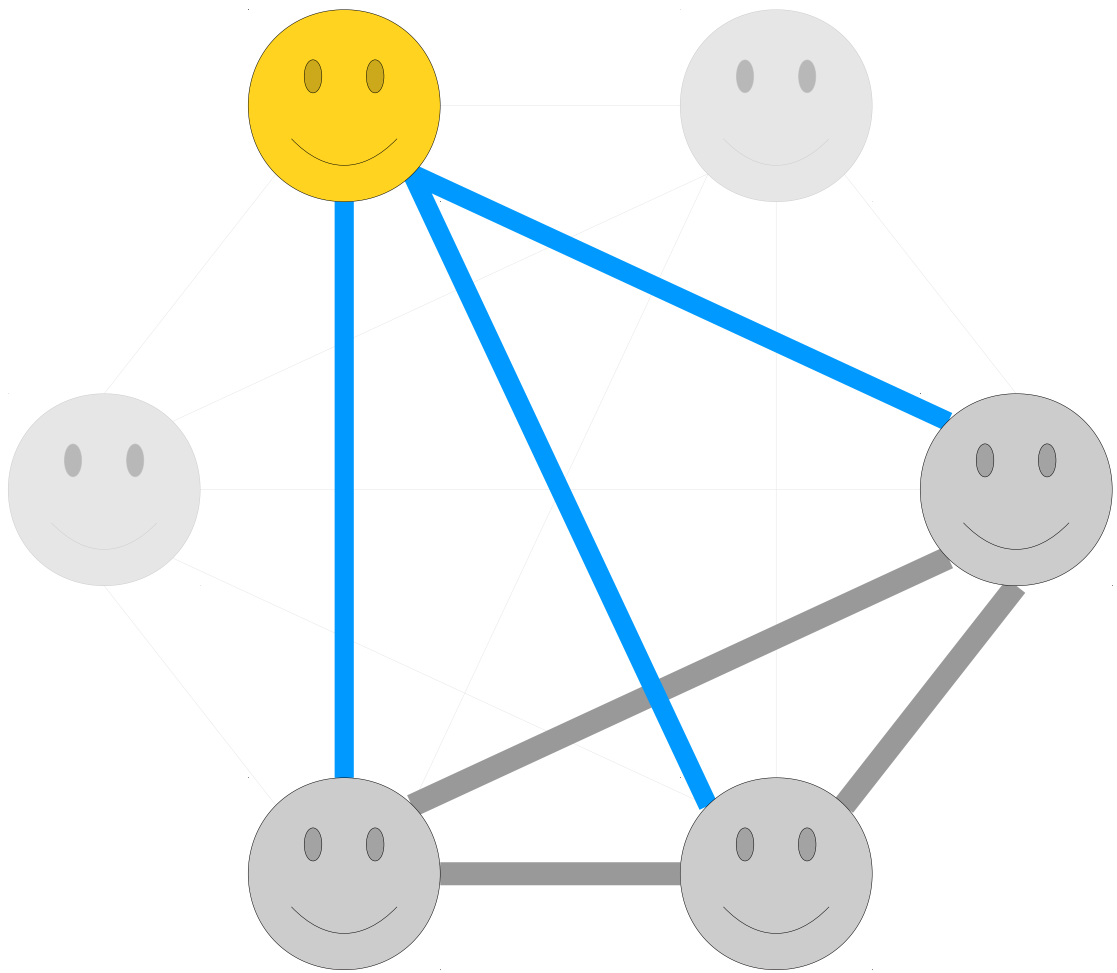

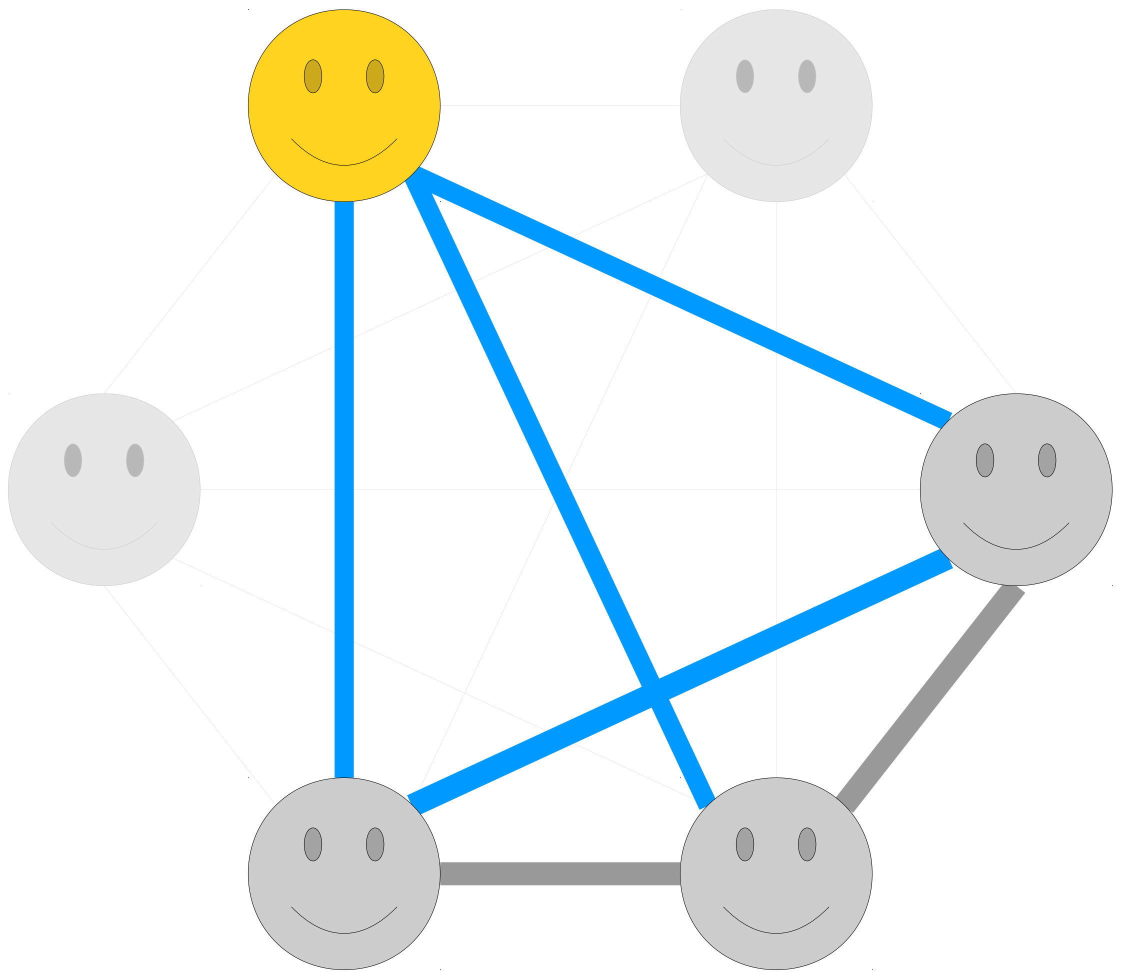

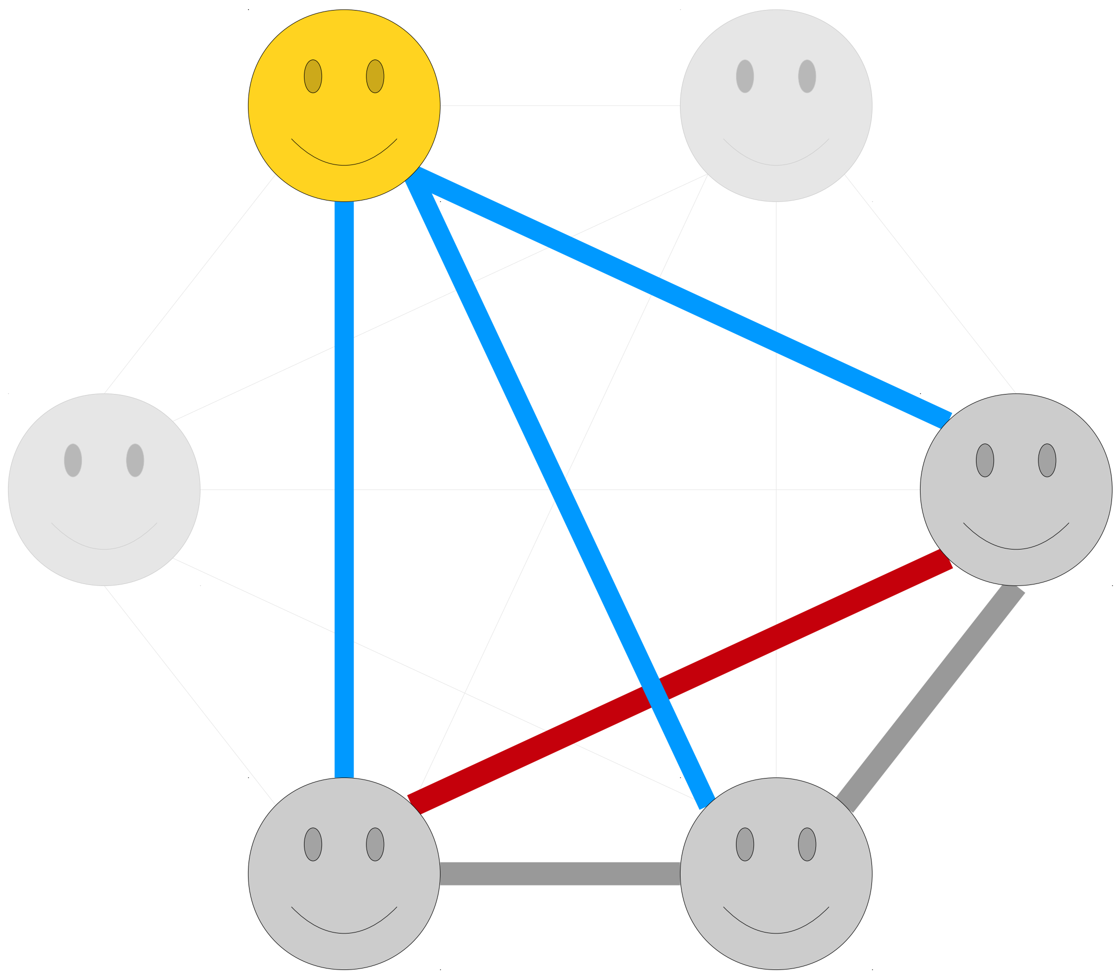

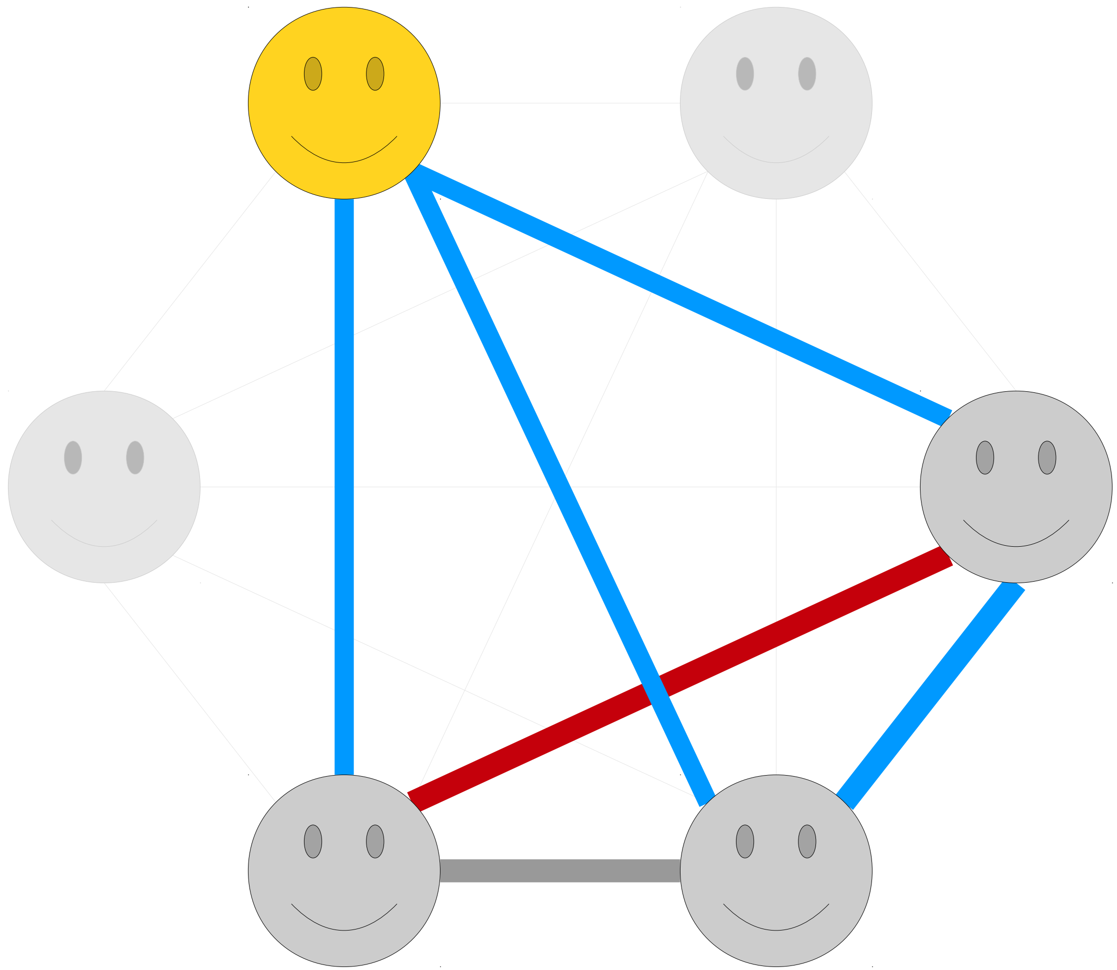

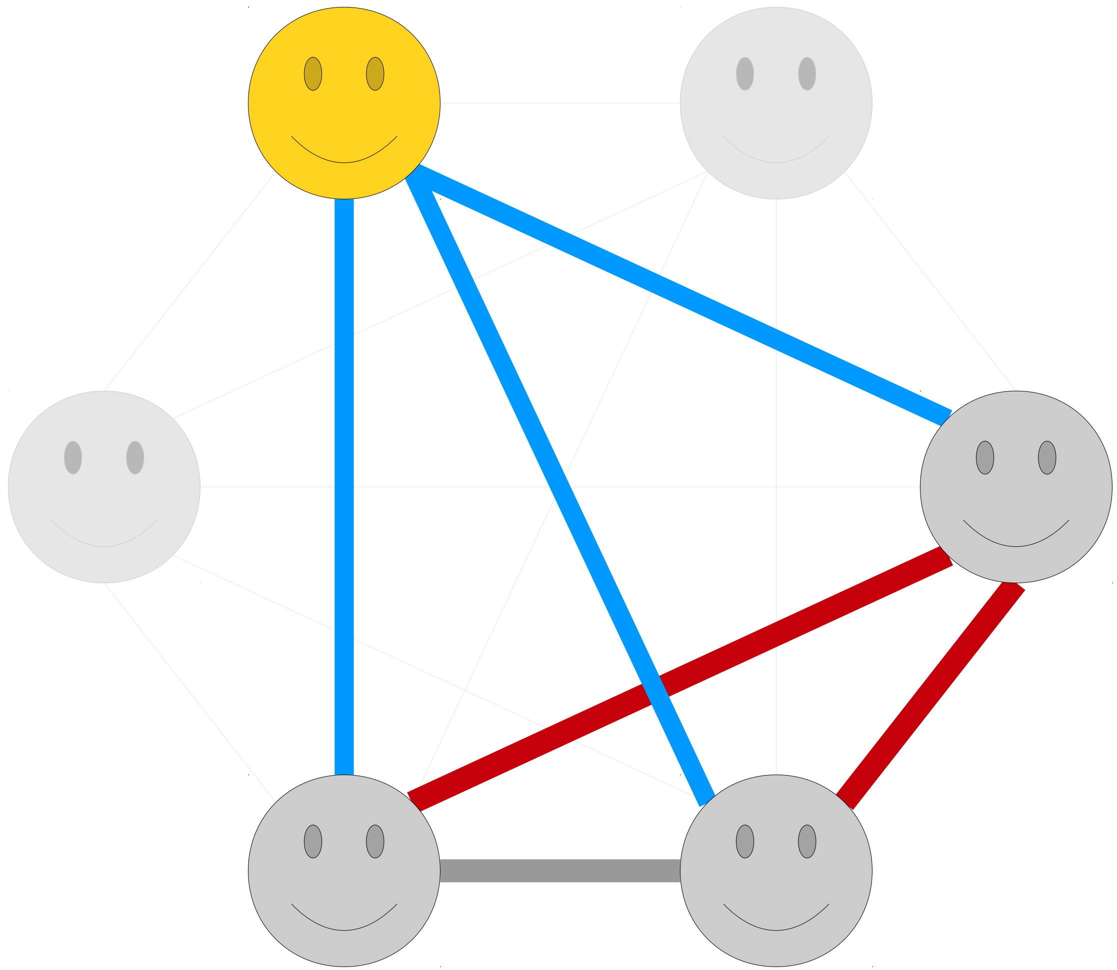

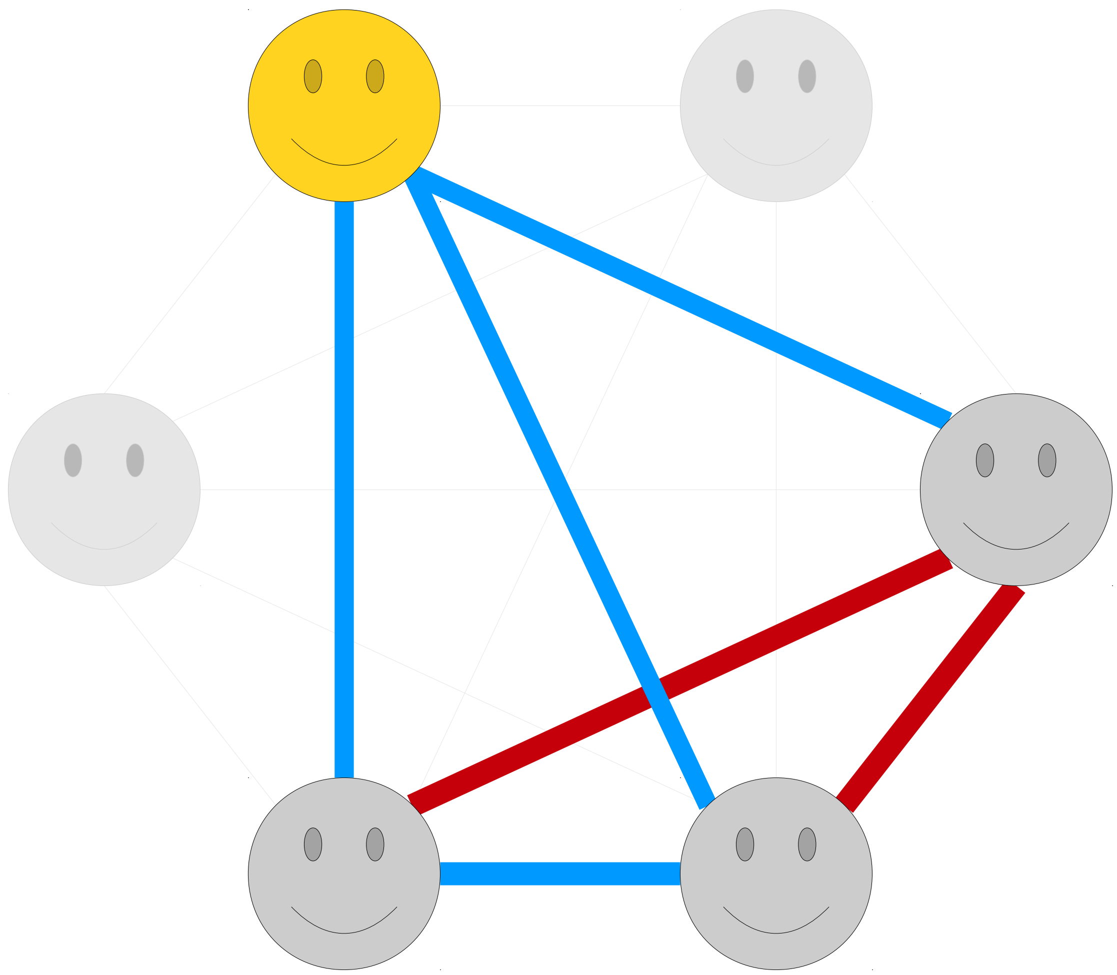

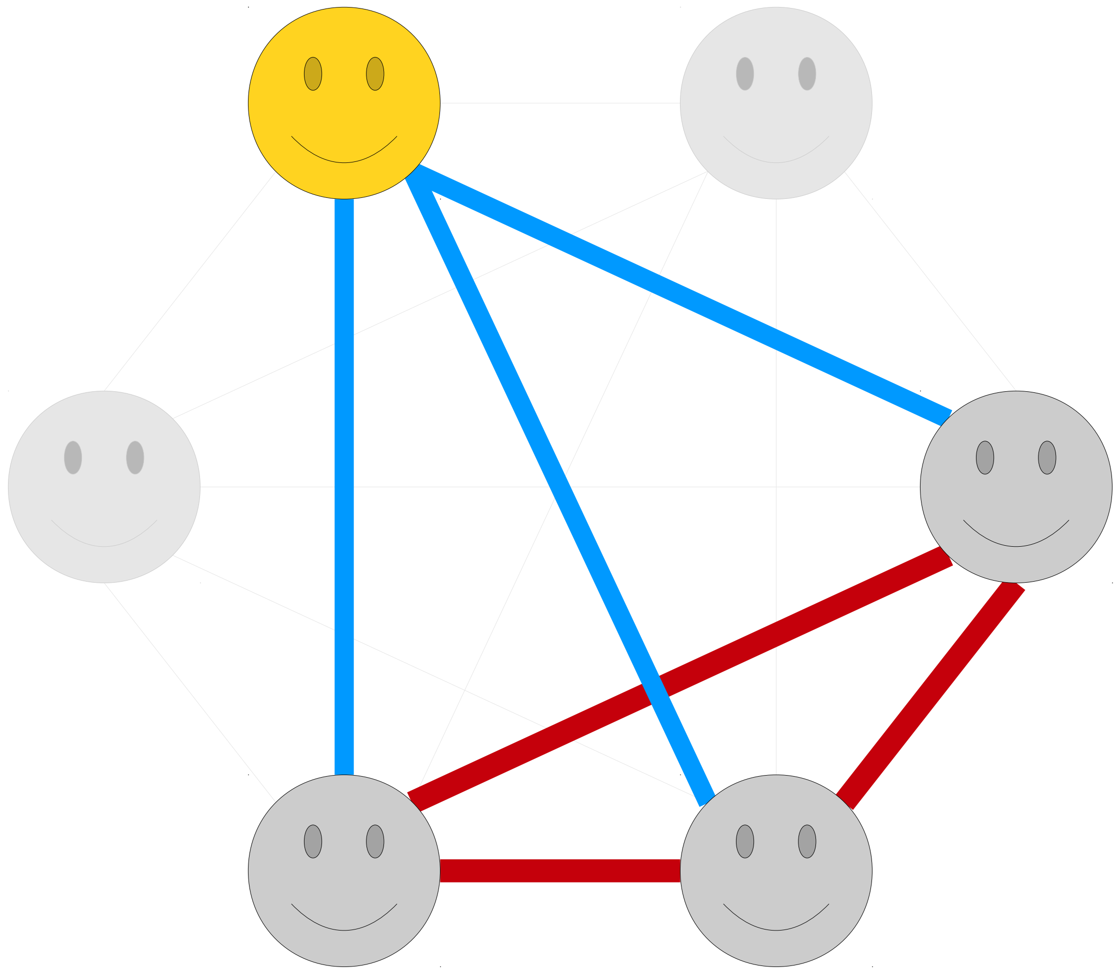

Theorem: Consider a 6-clique in which every edge is colored either red or blue. Then there must be a triangle of red edges, a triangle of blue edges, or both.

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. ■

Ramsey Theory

- The proof we did is a special case of a broader result.
- ***Theorem (Ramsey's Theorem):*** For any natural number n, there is a natural number $R(n)$ such that if the edges of an $R(n)$-clique are colored red or blue, the resulting graph will contain either a red n-clique or a blue n-clique.
  - Our proof was that $R(3) \leq 6$.
- A more philosophical take on this theorem: true disorder is impossible at a large scale, since no matter how you organize things, you're guaranteed to find some interesting substructure.

Time-Out for Announcements!

Problem Sets

- The Problem Set Three was due at 2:30PM today.
  - Take late days if you need them. Thanks for putting in such a great effort.
- Problem Set Four goes out today.
  - There's a checkpoint due on Monday at 2:30PM. It's designed to be very, very short.
  - Remaining problems are due on 2:30PM next Friday.

Midterm Exam Logistics

- The first midterm exam is ***Monday, October 23rd***, from ***7:00PM – 10:00PM***. Locations are divvied up by last (family) name:
  - `Abb – Lop`: Go to ***Cubberly Auditorium***.
  - `Mac – Zwa`: Go to ***Hewlett 200***.
- You're responsible for Lectures 00 – 05 and topics covered in PS1 – PS2. Later lectures (relations forward) and problem sets (PS3 onward) won't be tested here.
- 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.

Midterm Practice

- On the course website, you'll find
  - Extra Practice Problems 1, with solutions;
  - Practice Midterm 1, with solutions;
  - Practice Midterm 2, with solutions;
  - a link to the CS103A website, which has problems with solutions; and
  - Practice Midterm 3, with solutions.
- We'll have extra staffing on Piazza over the weekend. We've also added in some extra office hours shifts this weekend; check the calendar!

Your Questions

“Why 137?”

There’s something called the *fine structure constant* that determines the strength of electromagnetic interactions. You find it by throwing a ton of mathematical and physical constants into a formula:

$$\alpha = \frac{1}{4\pi\varepsilon_0}\frac{e^2}{\hbar c} = \frac{\mu_0}{4\pi}\frac{e^2 c}{\hbar} = \frac{k_e e^2}{\hbar c} = \frac{c\mu_0}{2R_K} = \frac{e^2}{4\pi}\frac{Z_0}{\hbar}$$

The value is approximately 1 / 137, and for a while people thought it was <u>exactly</u> 1 / 137, leading to some weird, misguided theological explanations. There’s also evidence that it might <u>not</u> be constant, indicating that something we think never changes actually does change!

Also, it’s prime, and it’s a great “nothing-up-my-sleeve” number.

Back to CS103!

A Little Math Puzzle

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?

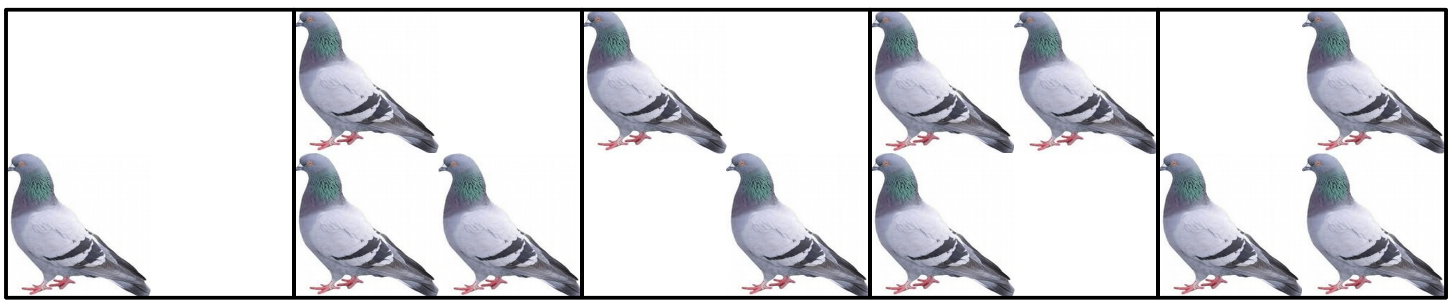

$m = 12$ pigeons
$n = 5$ boxes

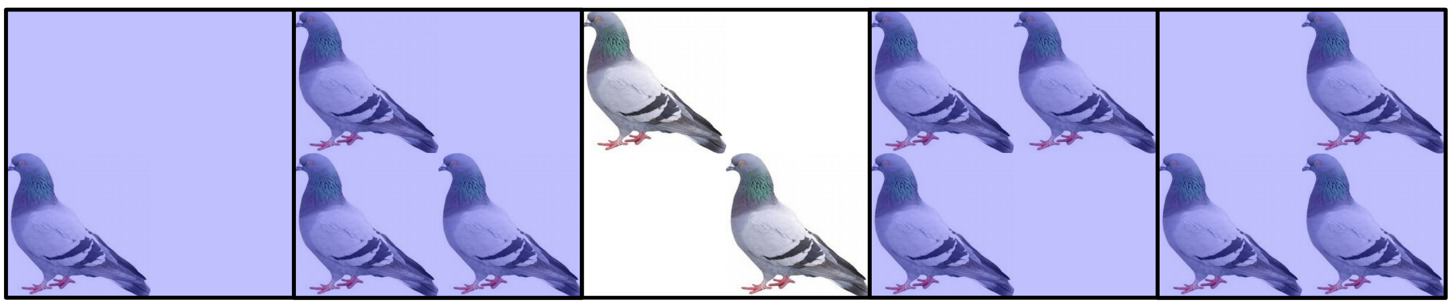

$m = 12$ pigeons
$n = 5$ boxes

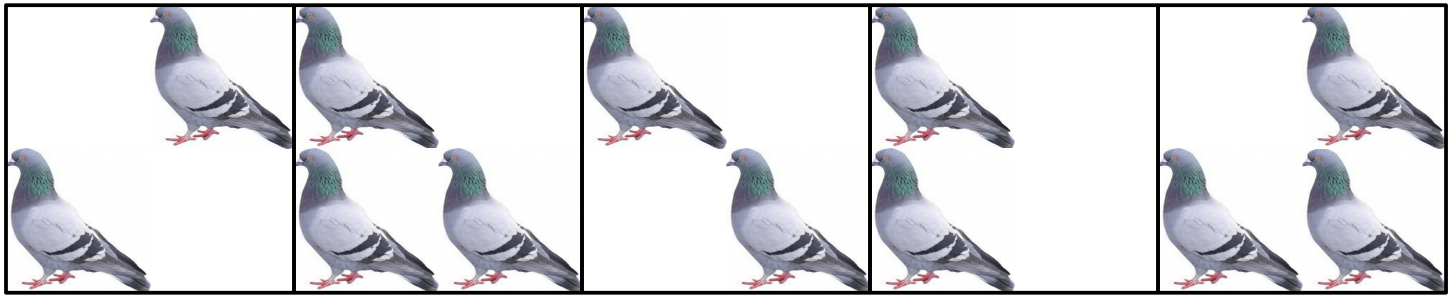

$m = 12$ pigeons
$n = 5$ boxes

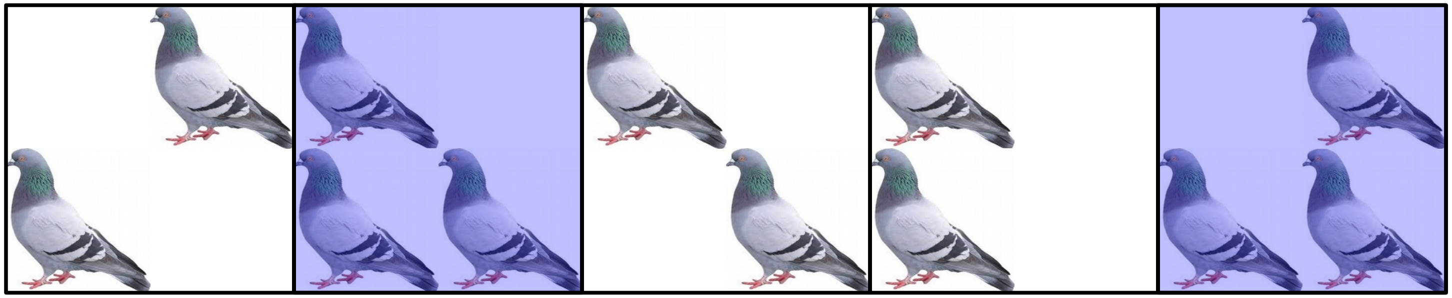

m = 12 pigeons
n = 5 boxes

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

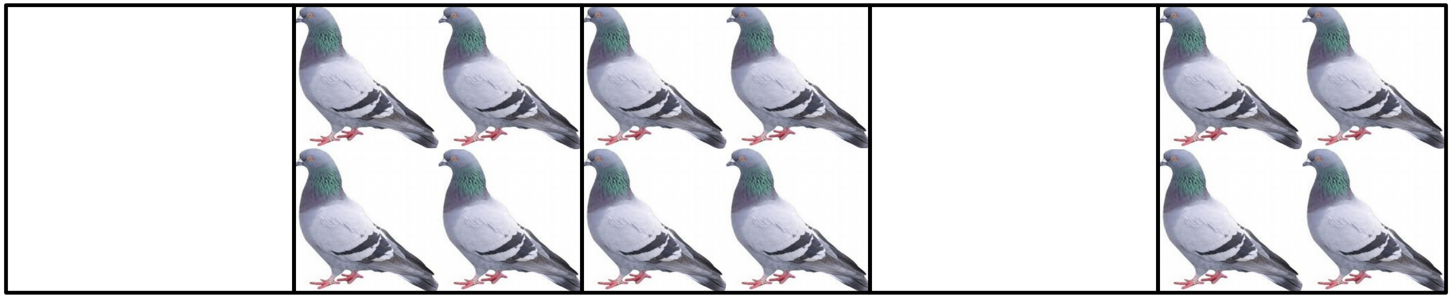

$m = 12$ pigeons
$n = 5$ boxes

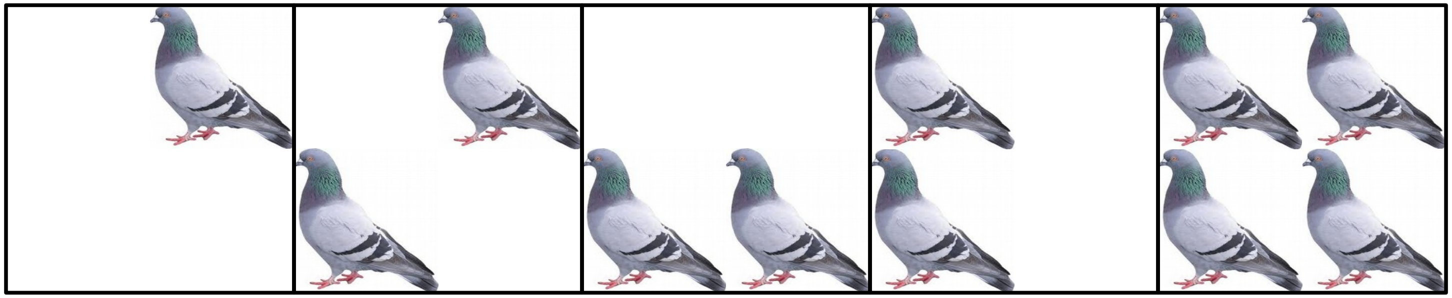

$m = 11$ pigeons
$n = 5$ boxes

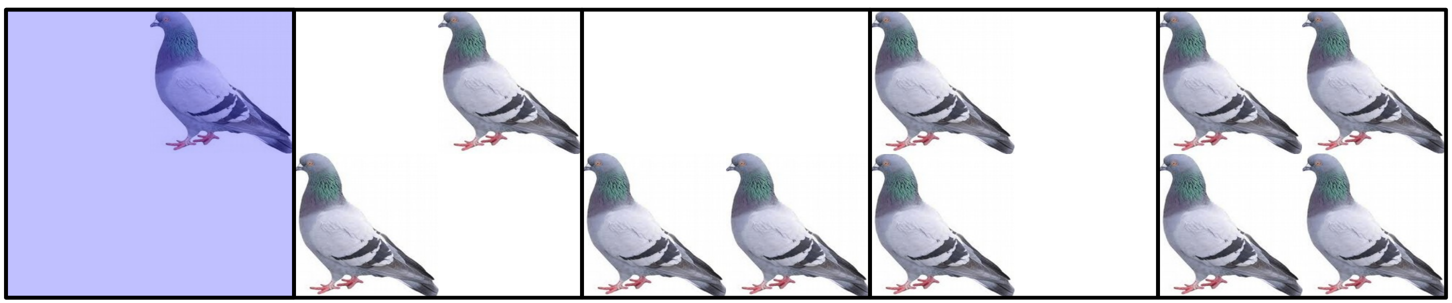

$m = 11$ pigeons
$n = 5$ boxes

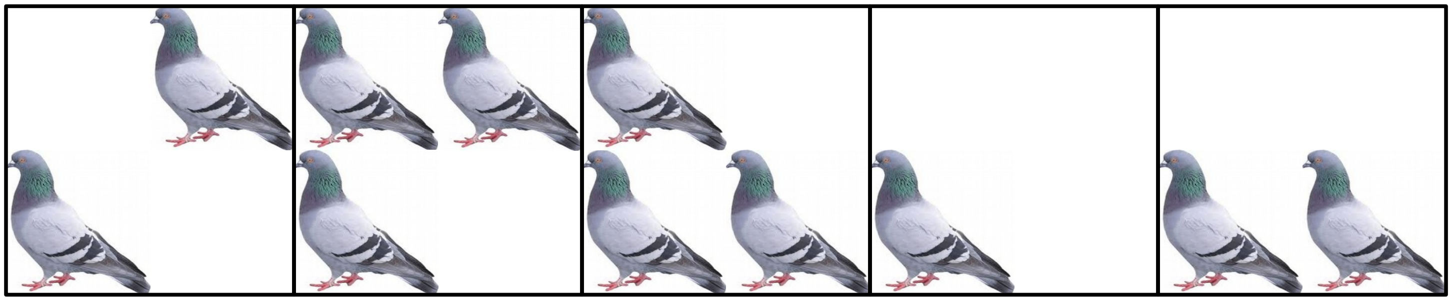

m = 11 pigeons
n = 5 boxes

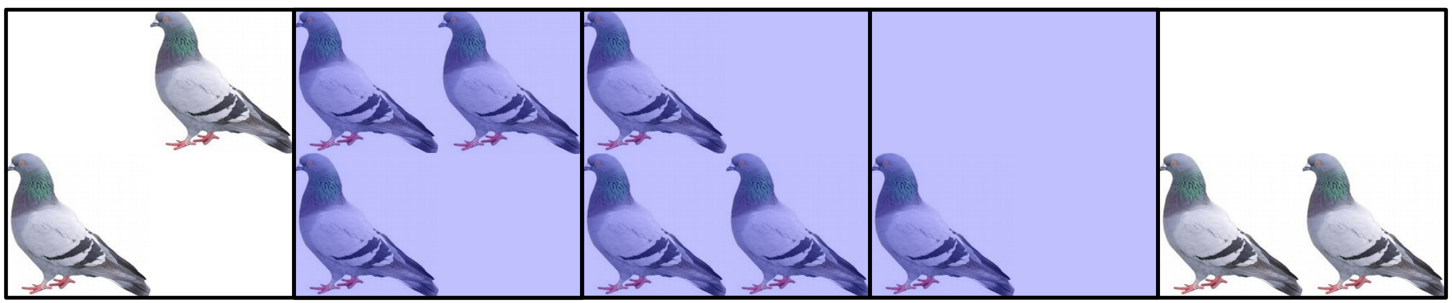

$m = 11$ pigeons
$n = 5$ boxes

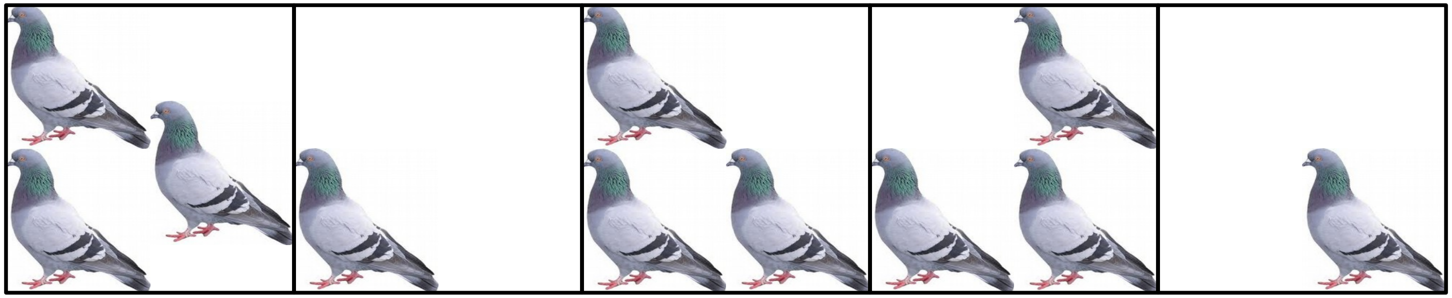

$m = 11$ pigeons
$n = 5$ boxes

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

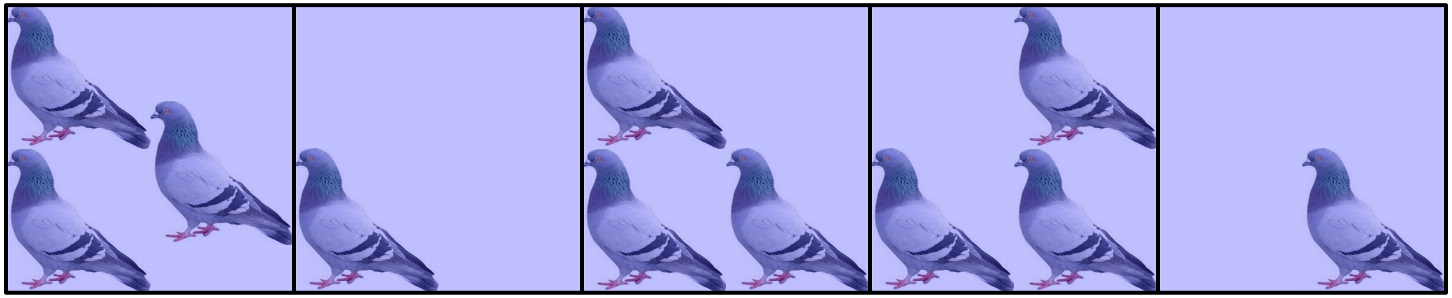

m = 11 pigeons
n = 5 boxes

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:

The Handshaking Lemma

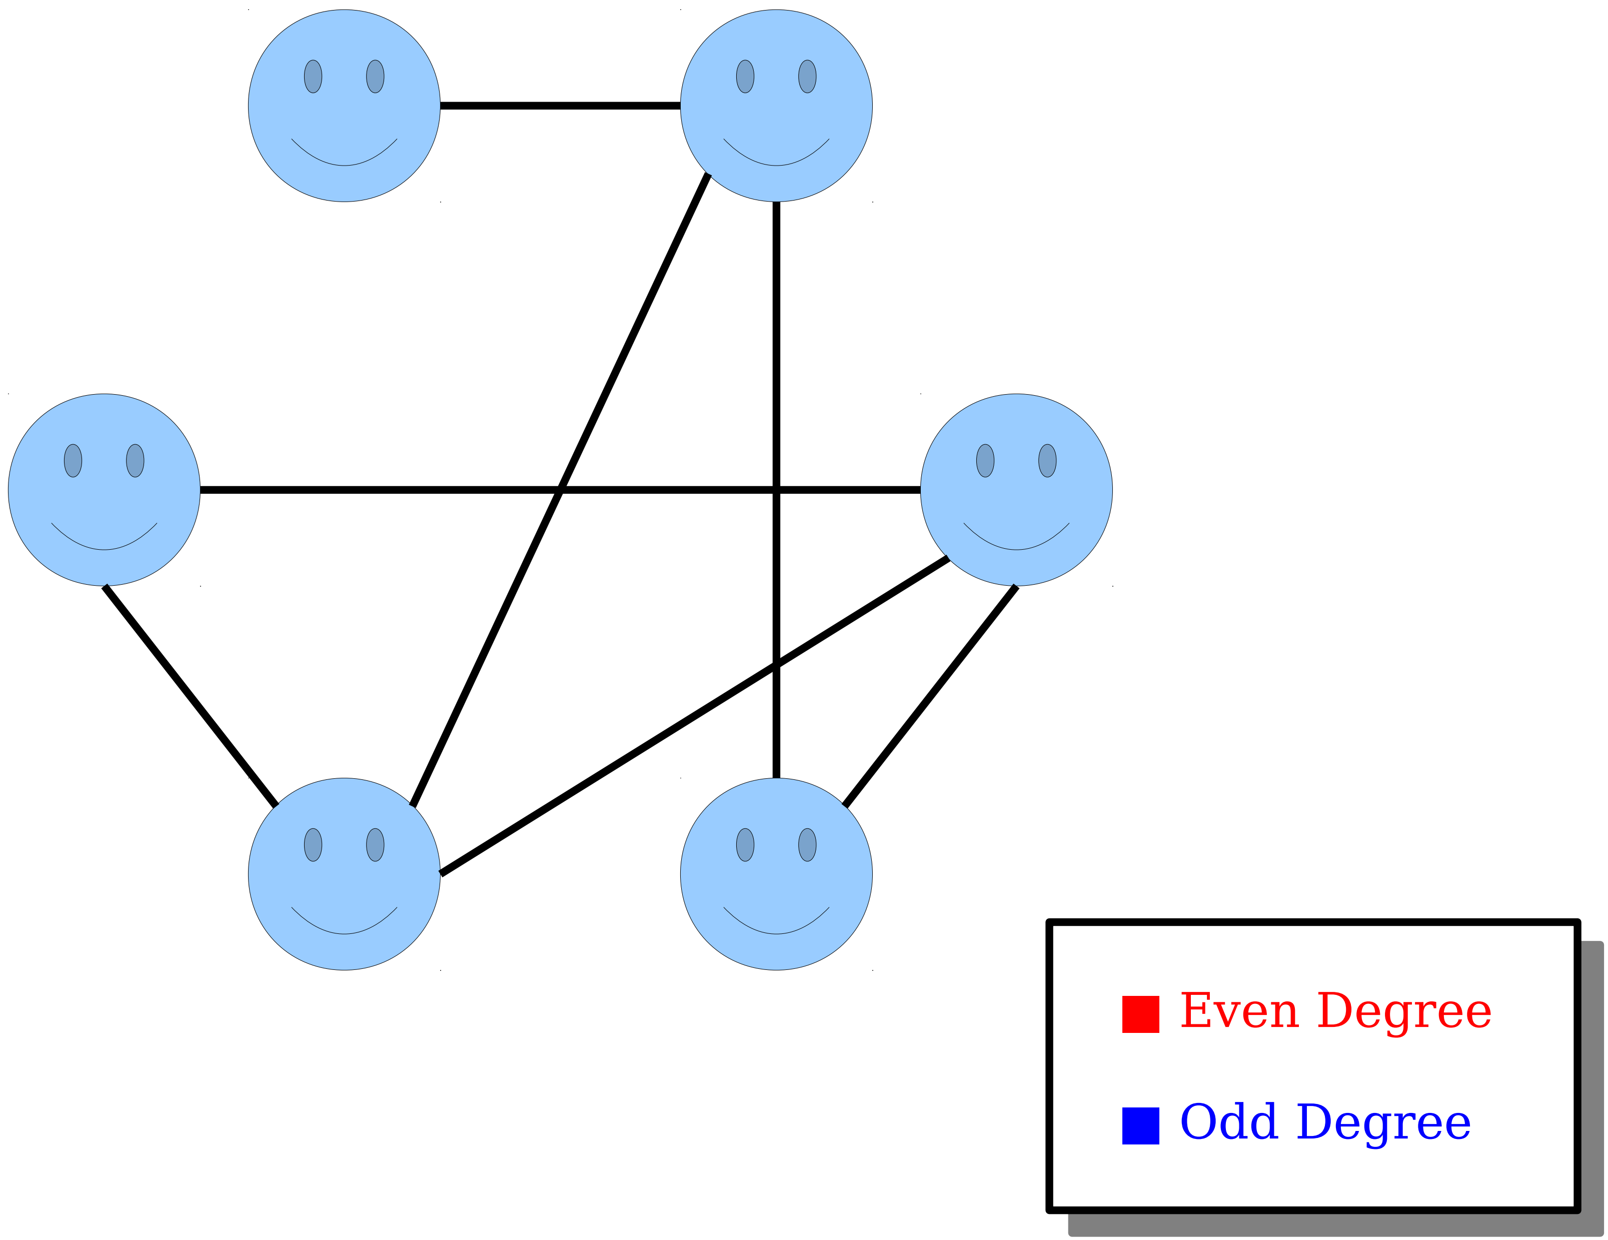
Even Degree
Odd Degree

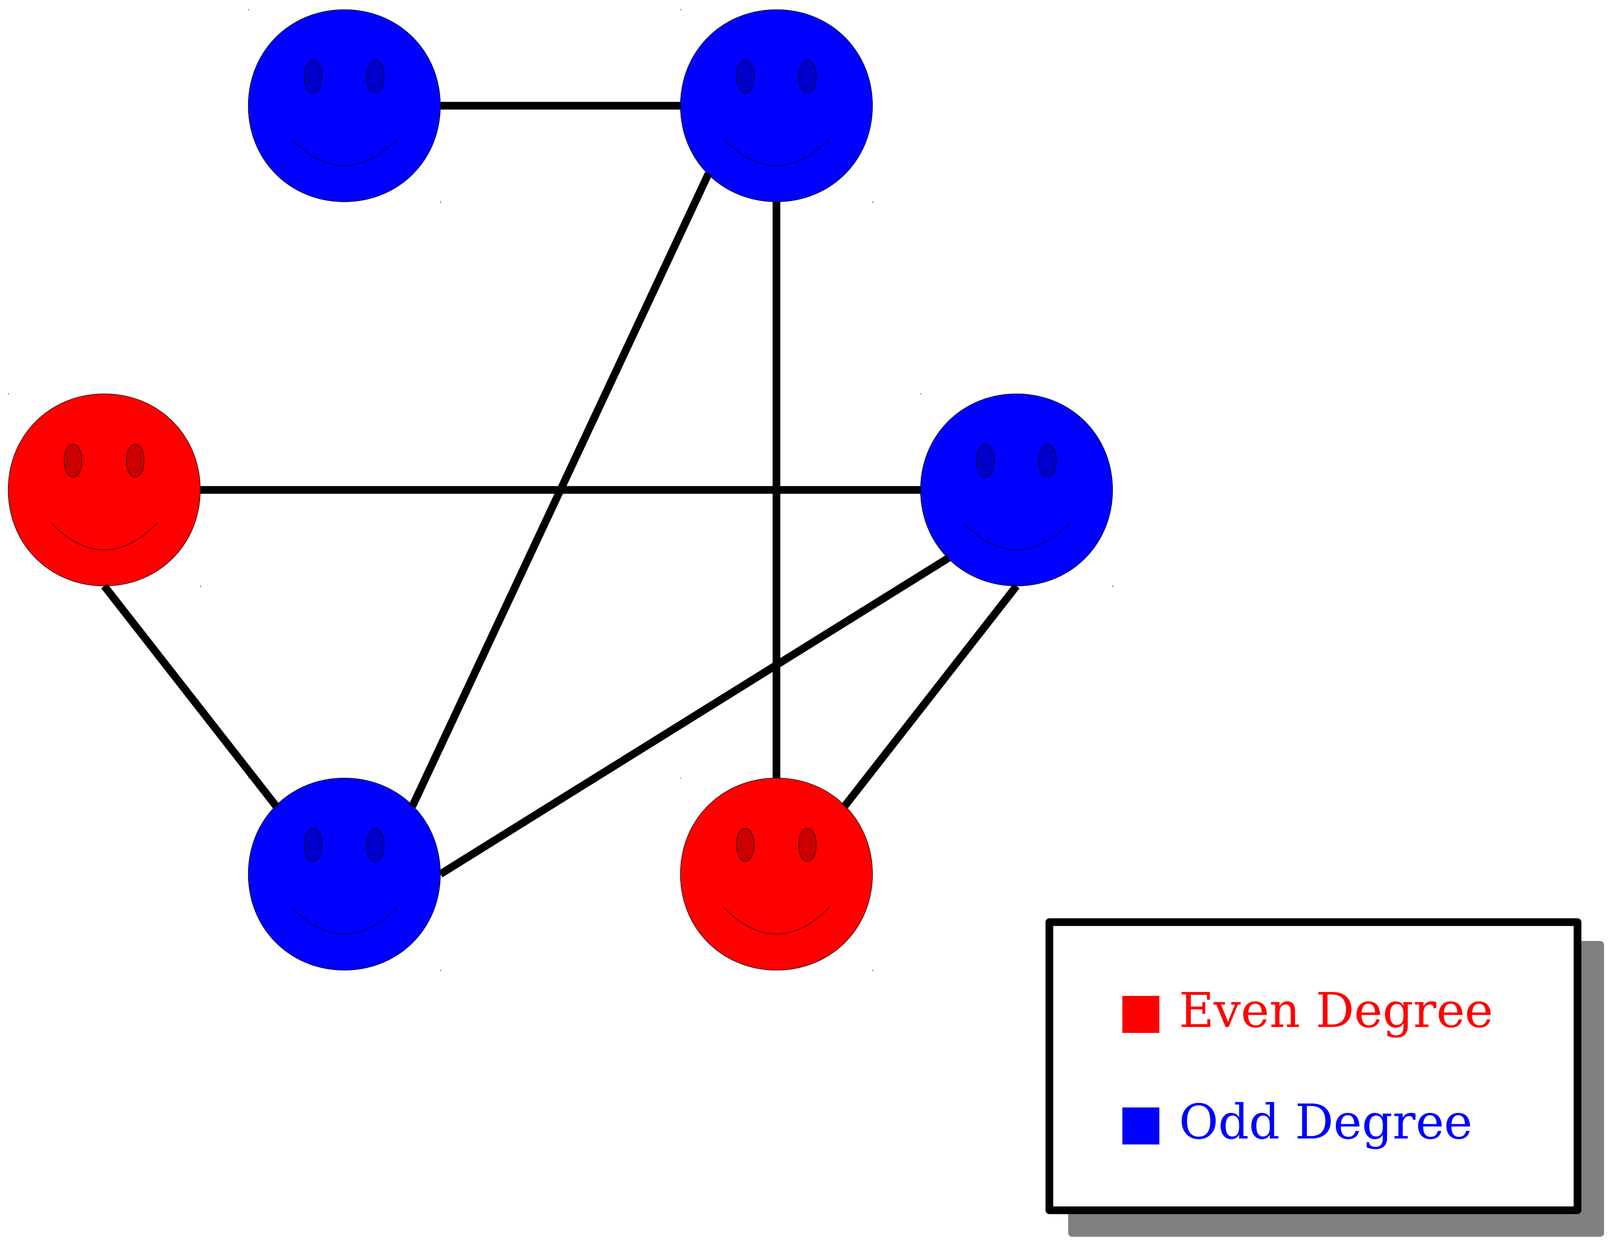
Even Degree
Odd Degree

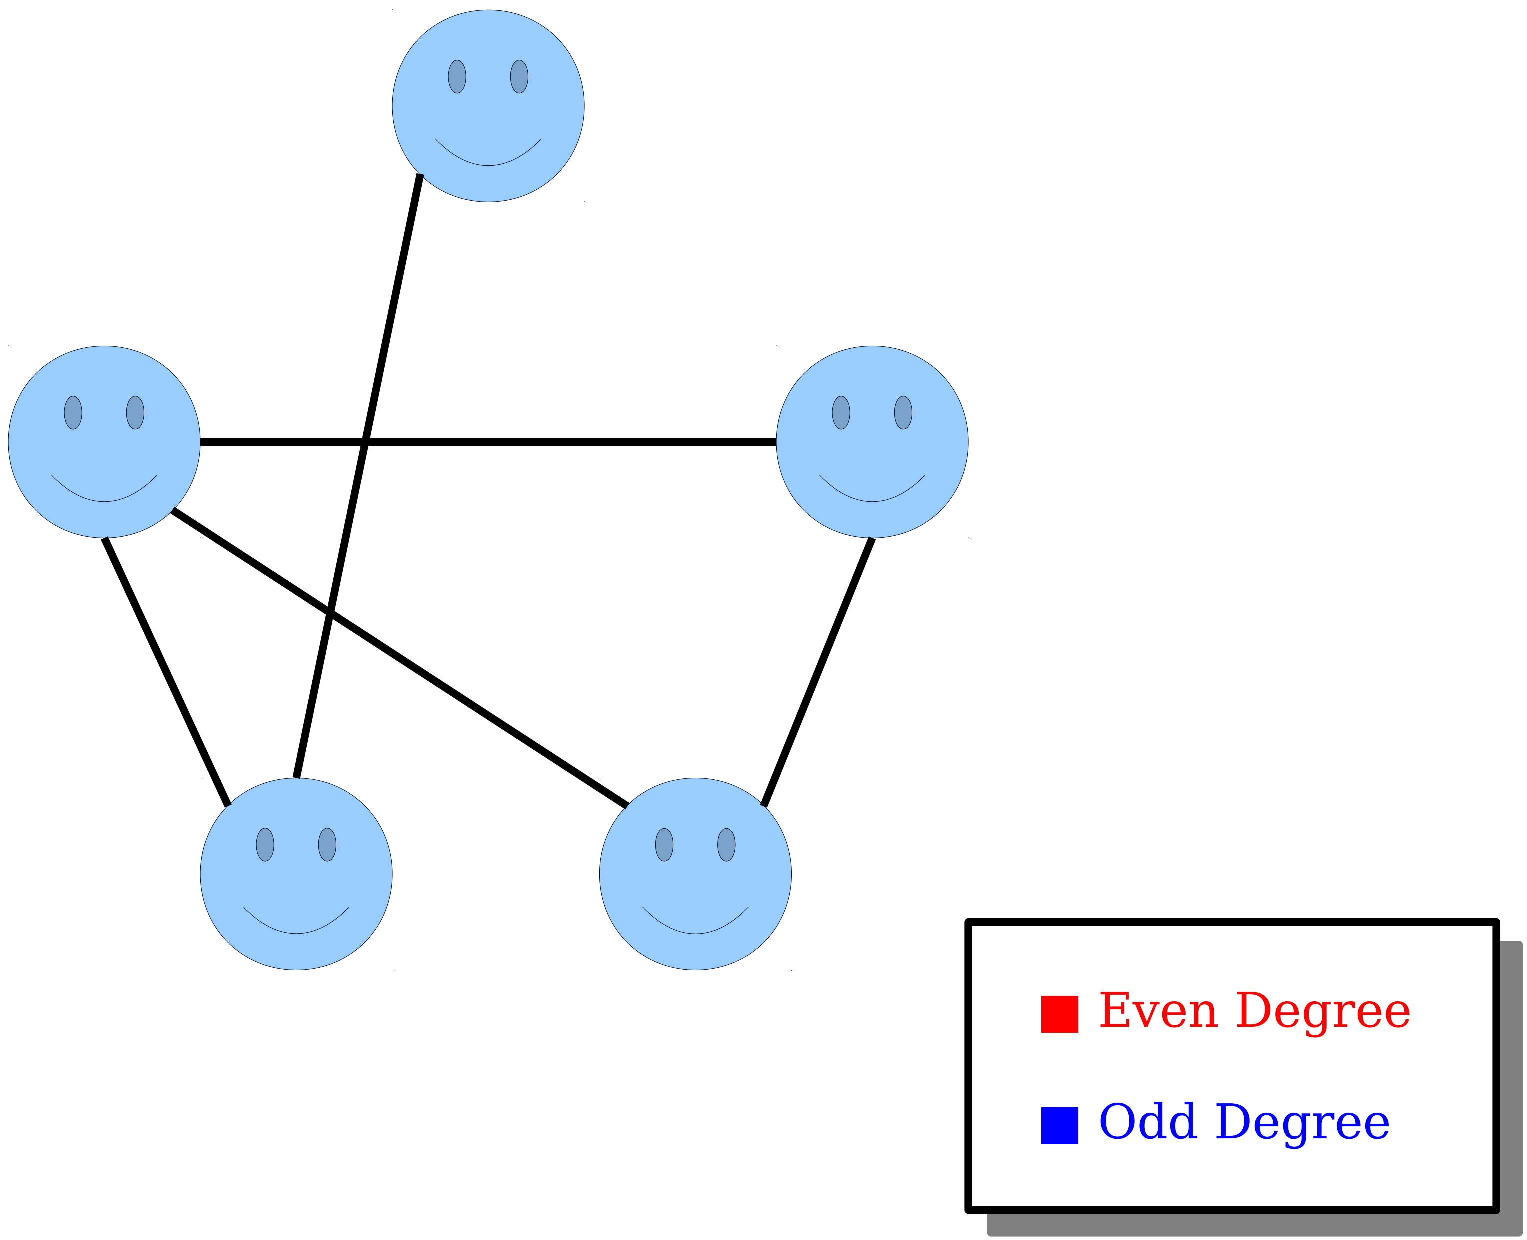
Even Degree
Odd Degree

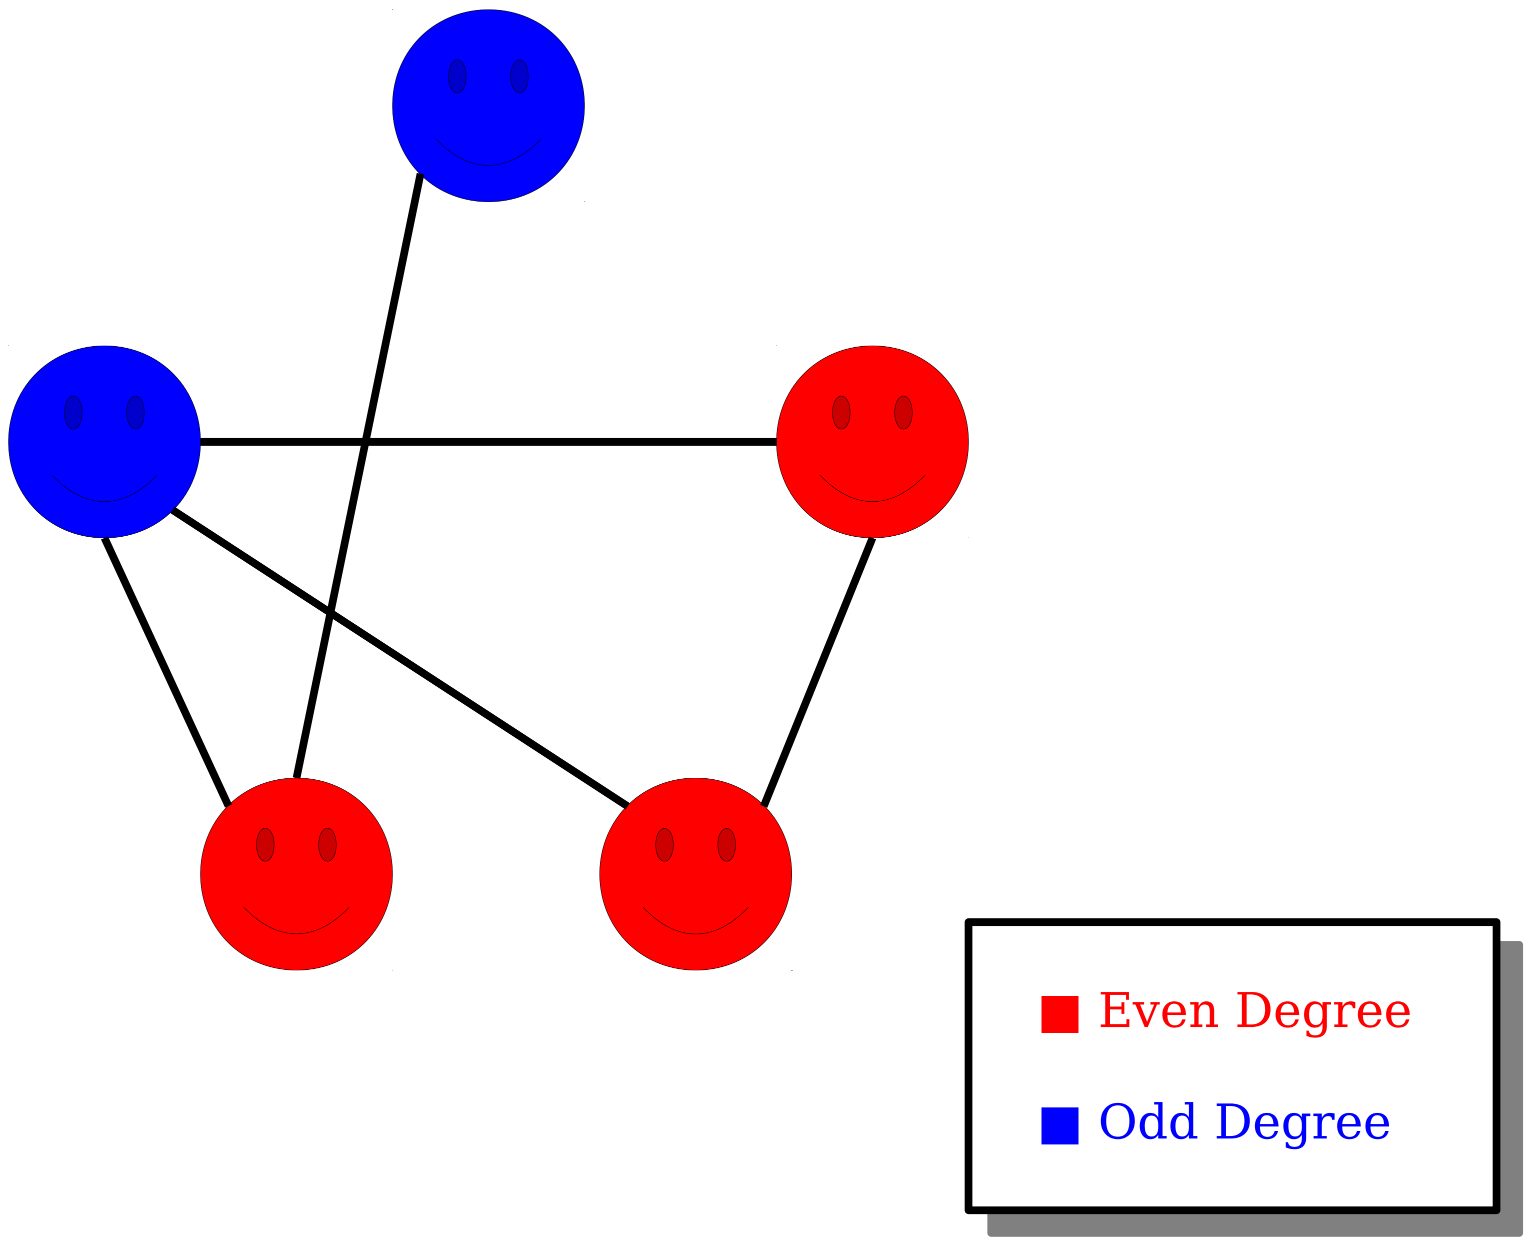
Even Degree
Odd Degree

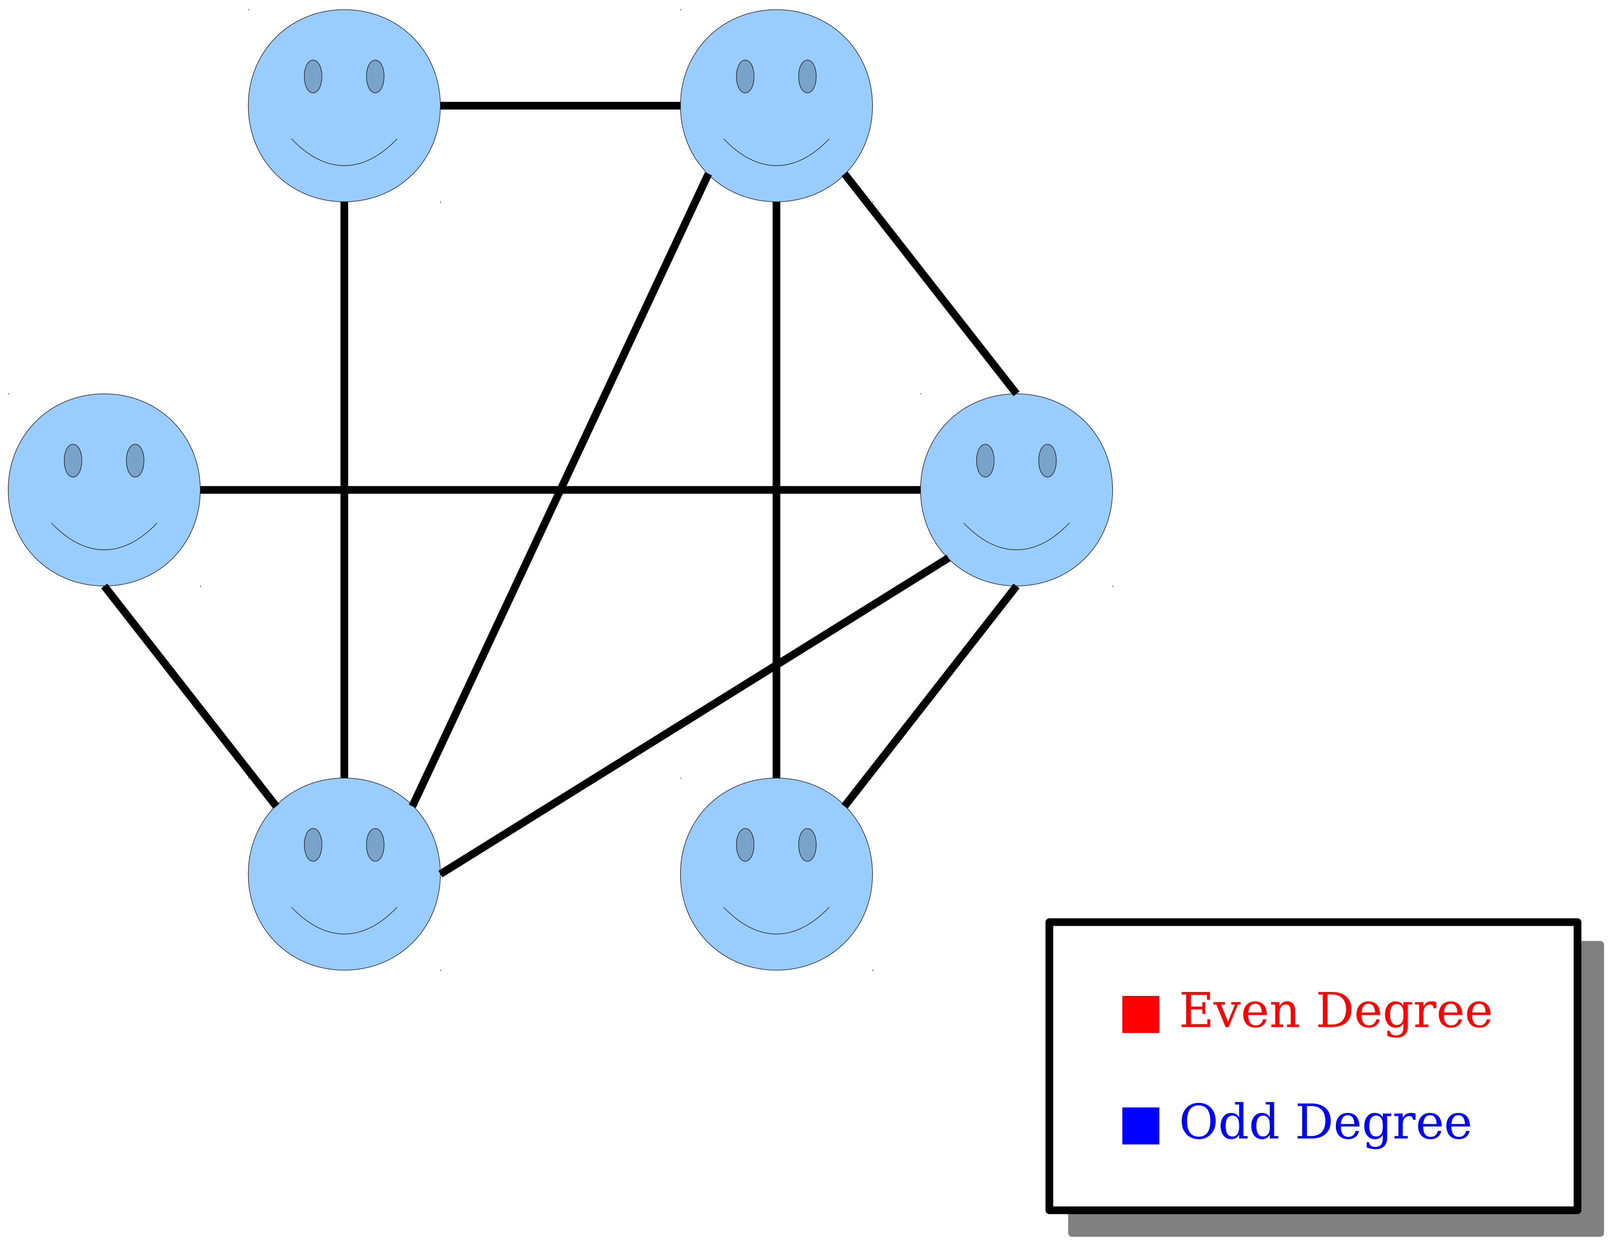

Even Degree
Odd Degree

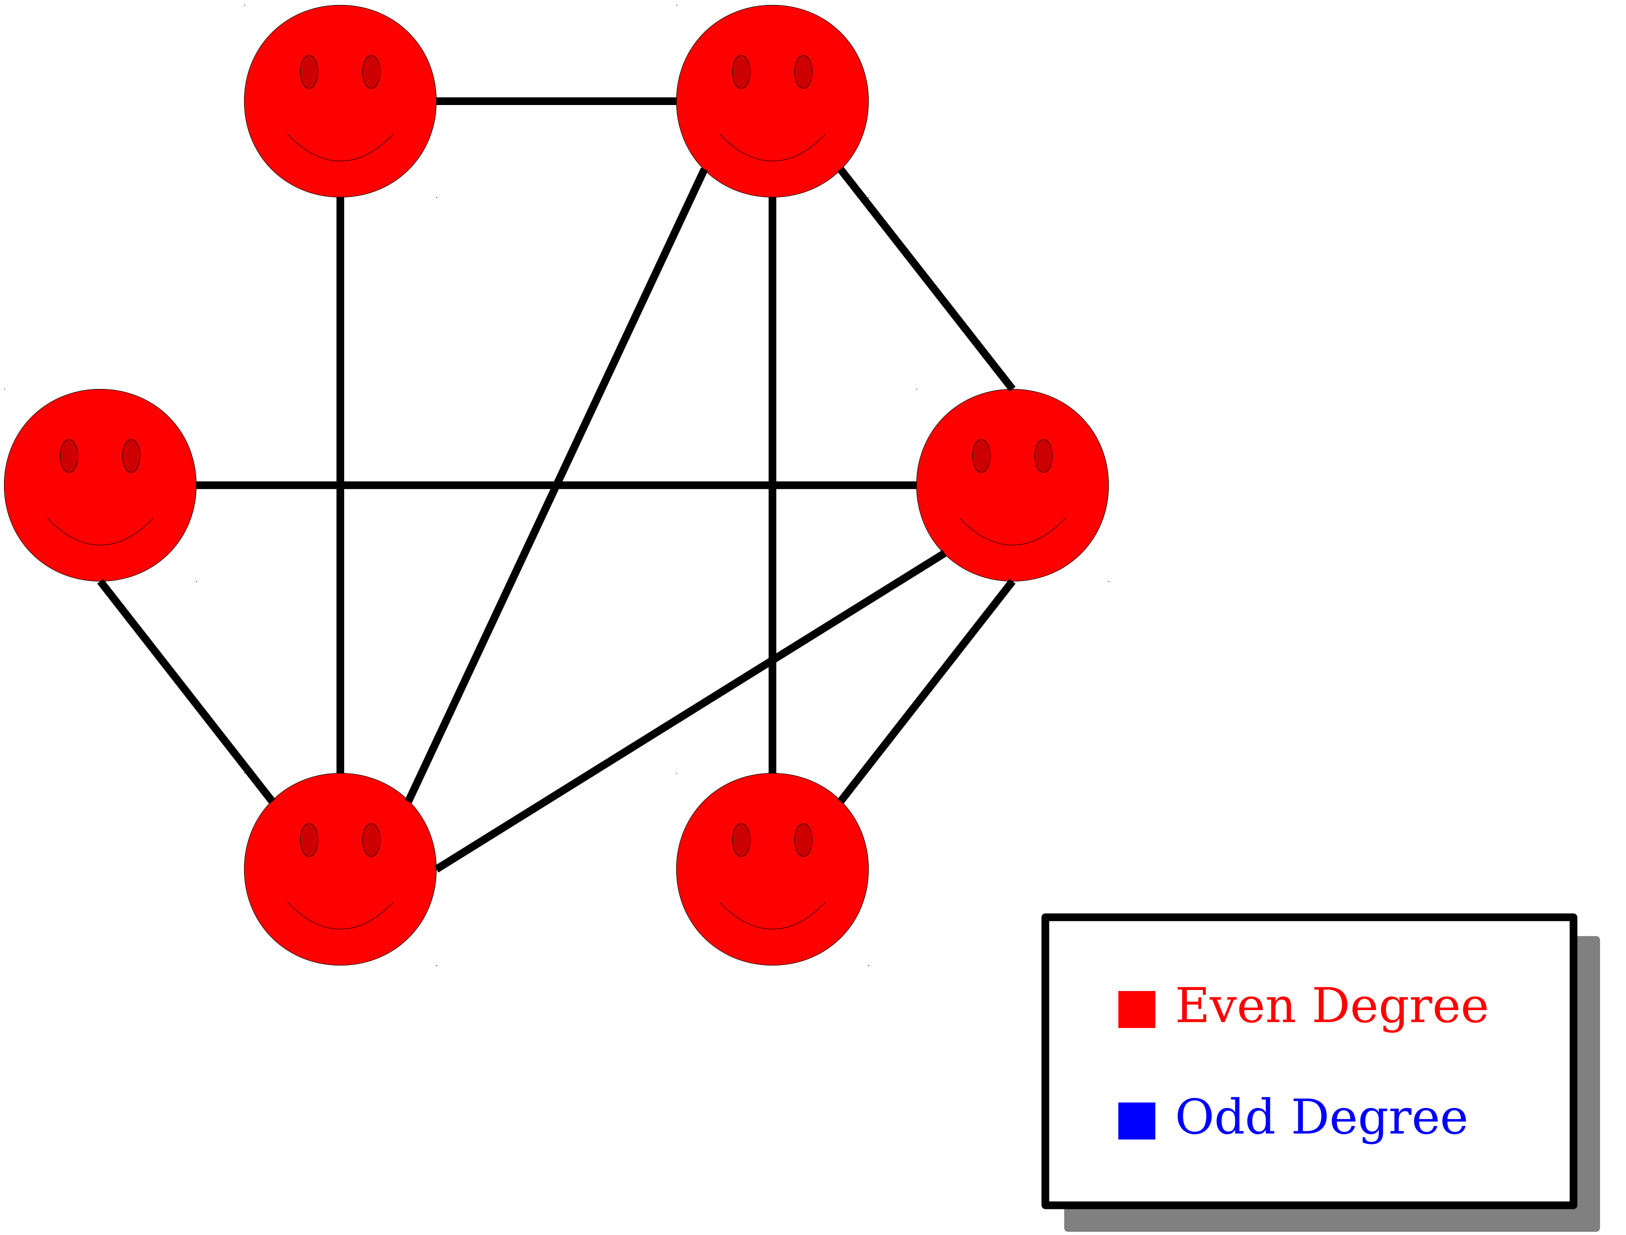

Even Degree
Odd Degree

Theorem (The Handshaking Lemma): Let $G = (V, E)$ be a graph. Then each connected component of G has an even number of nodes of odd degree.

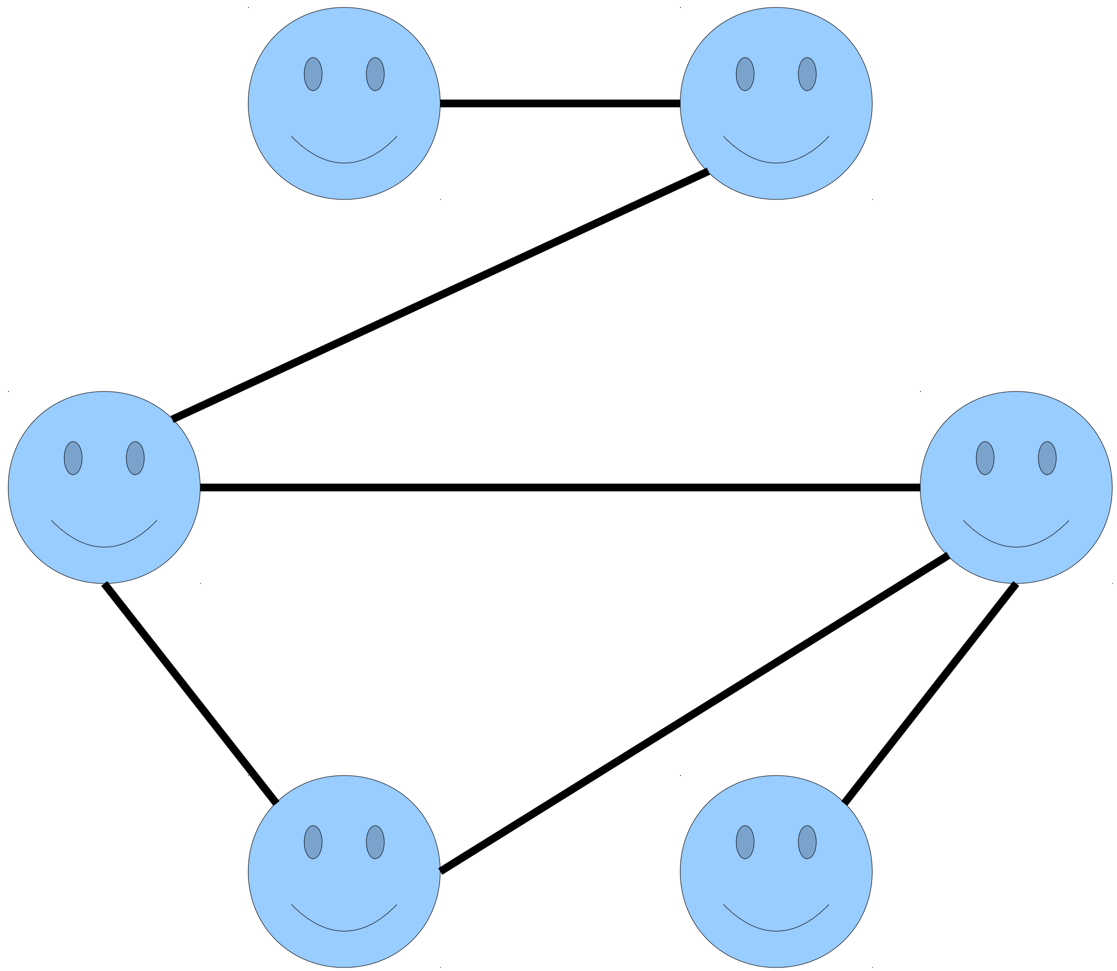

There are $2m$ total coins here, where m is the number of edges.

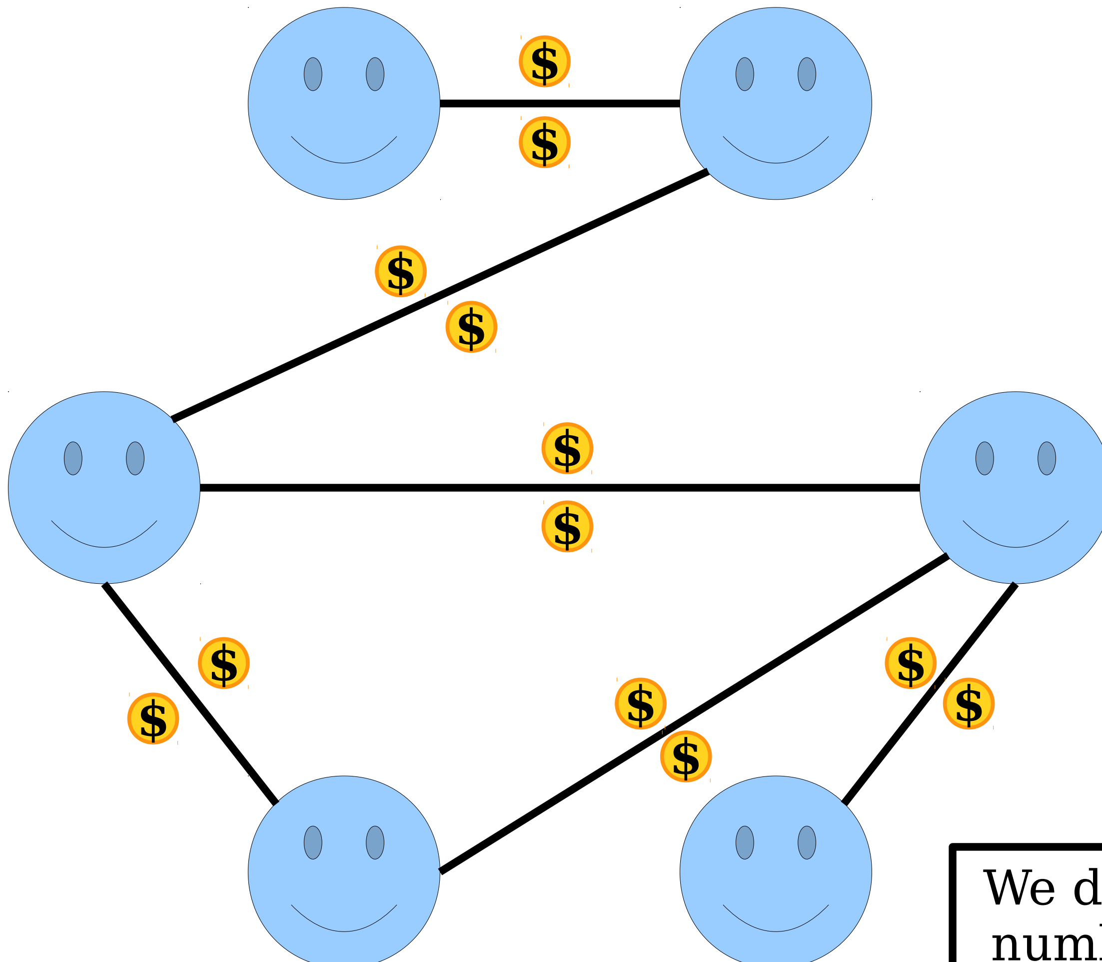

We distributed an even number of coins into a collection of nodes. Therefore, the number of nodes of odd degree is even!

Theorem (Handshaking Lemma): If G is a graph, then each connected component of G has an even number of nodes of odd degree.

Proof: Let $G = (V, E)$ be a graph and let C be a connnected component of G. Place one coin on each node in C for each edge in E incident to it. Notice that the number of coins on any node v is equal to $\deg(v)$.

We claim that there are an even total coins distributed across all the nodes of G. Notice that each edge contributes two coins to the total, one for each of its endpoints. This means that there are $2m$ total coins distributed across the nodes of V, where m is the number of edges adjacent to nodes in C, and $2m$ is even.

Since there are an even number of coins distributed across the nodes, our earlier theorem tells us that the number of nodes in G with an odd number of coins on them must be even. The number of coins on each node is the degree of that node, and therefore there must be an even number of nodes of odd degree. ■

A Fun Corollary

- A ***corollary*** of a theorem is a statement that follows nicely from the theorem.
- The previous theorem has this lovely follow-up:
- ***Corollary:*** If G is a graph with exactly two nodes of odd degree, those nodes are connected.

Corollary: If G is a graph with exactly two nodes of odd degree, then those two nodes are connected in G.

Proof: Let G be a graph with exactly two nodes u and v of odd degree. Consider the connected component C containing the node u. By the Handshaking Lemma, we know that C must contain an even number of nodes of odd degree. Therefore, C must contain at least one node of odd degree other than u, since otherwise C would have exactly one node of odd degree. Since v is the only node in G aside from u that has odd degree, we see that v must belong to C. Overall, this means that u and v are in the same connected component, so u and v are connected in G, as required. ■

Some Applications

- The corollary we just presented has some pretty unexpected applications:
  - The ***mountain-climbing theorem***. Suppose that two people start climbing the same mountain, beginning at any two spots they'd like. Assume the mountain has no regions that are perfectly flat, the two climbers can each choose a path to the summit such that they arrive at the same time, and have the same altitude throughout the entire journey.
  - ***Sperner's lemma***. A powerful mathematical primitive that lets you find equitable ways to split the rent in an apartment or show that no matter how you stir your coffee, there's always some particle that remains in the same place.
- And, as you saw on Problem Set Two, looking at parity is a powerful way to prove that certain objects must exist!

Next Time

- ***No class on Monday – you have a midterm!***
- ***Then, on Wednesday...***
  - ***Mathematical Induction***
    - Proofs on stepwise processes
  - ***Applications of Induction***
    - … to numbers!
    - … to data compression!
    - … to puzzles!
    - … to algorithms!