From the book *Networks, Crowds, and Markets: Reasoning about a Highly Connected World.*
By David Easley and Jon Kleinberg. Cambridge University Press, 2010.
Complete preprint on-line at http://www.cs.cornell.edu/home/kleinber/networks-book/

Chapter 5

Positive and Negative Relationships

In our discussion of networks thus far, we have generally viewed the relationships contained in these networks as having positive connotations — links have typically indicated such things as friendship, collaboration, sharing of information, or membership in a group. The terminology of on-line social networks reflects a largely similar view, through its emphasis on the connections one forms with friends, fans, followers, and so forth. But in most network settings, there are also negative effects at work. Some relations are friendly, but others are antagonistic or hostile; interactions between people or groups are regularly beset by controversy, disagreement, and sometimes outright conflict. How should we reason about the mix of positive and negative relationships that take place within a network?

Here we describe a rich part of social network theory that involves taking a network and annotating its links (i.e., its edges) with positive and negative signs. Positive links represent friendship while negative links represent antagonism, and an important problem in the study of social networks is to understand the tension between these two forces. The notion of *structural balance* that we discuss in this chapter is one of the basic frameworks for doing this.

In addition to introducing some of the basics of structural balance, our discussion here serves a second, methodological purpose: it illustrates a nice connection between local and global network properties. A recurring issue in the analysis of networked systems is the way in which *local* effects — phenomena involving only a few nodes at a time — can have global consequences that are observable at the level of the network as a whole. Structural balance offers a way to capture one such relationship in a very clean way, and by purely mathematical analysis: we will consider a simple definition abstractly, and find that it inevitably leads to certain macroscopic properties of the network.

Draft version: June 10, 2010

5.1 Structural Balance

We focus here on perhaps the most basic model of positive and negative relationships, since it captures the essential idea. Suppose we have a social network on a set of people, in which everyone knows everyone else — so we have an edge joining each pair of nodes. Such a network is called a *clique*, or a *complete graph*. We then *label* each edge with either + or −; a + label indicates that its two endpoints are friends, while a − label indicates that its two endpoints are enemies.

Note that since there's an edge connecting each pair, we are assuming that each pair of people are either friends or enemies — no two people are indifferent to one another, or unaware of each other. Thus, the model we're considering makes the most sense for a group of people small enough to have this level of mutual awareness (e.g. a classroom, a small company, a sports team, a fraternity or sorority), or for a setting such as international relations, in which the nodes are countries and every country has an official diplomatic position toward every other.[1]

The principles underlying structural balance are based on theories in social psychology dating back to the work of Heider in the 1940s [216], and generalized and extended to the language of graphs beginning with the work of Cartwright and Harary in the 1950s [97, 126, 204]. The crucial idea is the following. If we look at any two people in the group in isolation, the edge between them can be labeled + or −; that is, they are either friends or enemies. But when we look at sets of *three* people at a time, certain configurations of +'s and −'s are socially and psychologically more plausible than others. In particular, there are four distinct ways (up to symmetry) to label the three edges among three people with +'s and −'s; see Figure 5.1. We can distinguish among these four possibilities as follows.

- Given a set of people A, B, and C, having three pluses among them (as in Figure 5.1(a)) is a very natural situation: it corresponds to three people who are mutual friends.

- Having a single plus and two minuses in the relations among the there people is also very natural: it means that two of the three are friends, and they have a mutual enemy in the third. (See Figure 5.1(c).)

- The other two possible labelings of the triangle on A, B, and C introduce some amount of psychological "stress" or "instability" into the relationships. A triangle with two pluses and one minus corresponds (as in Figure 5.1(b)) to a person A who is friends with each of B and C, but B and C don't get along with each other. In this type of situation, there would be implicit forces pushing A to try to get B and C to become

[1] Later, in Section 5.5, we will consider the more general setting in which not every pair of nodes is necessarily connected by an edge.

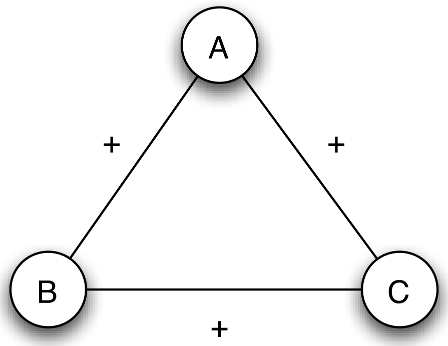

(a) *A, B, and C are mutual friends: balanced.*

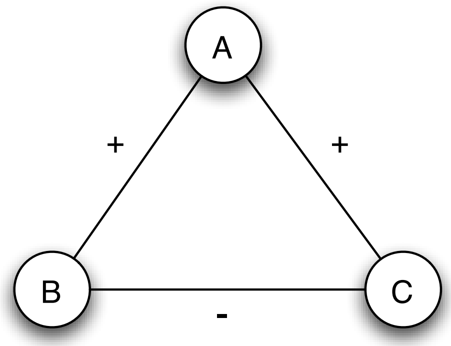

(b) *A is friends with B and C, but they don't get along with each other: not balanced.*

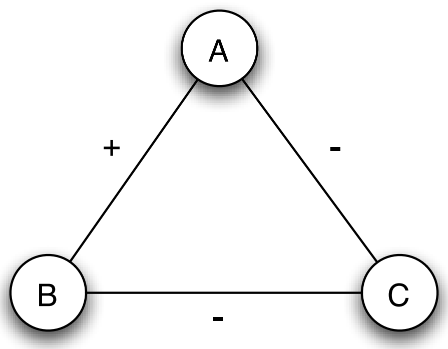

(c) *A and B are friends with C as a mutual enemy: balanced.*

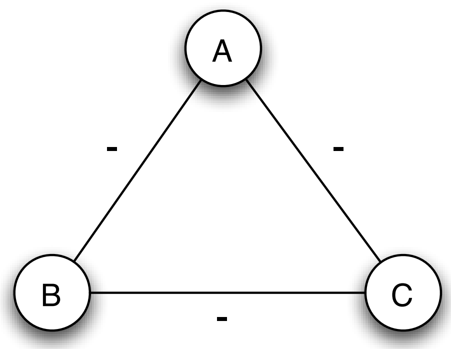

(d) *A, B, and C are mutual enemies: not balanced.*

Figure 5.1: Structural balance: Each labeled triangle must have 1 or 3 positive edges.

friends (thus turning the B-C edge label to $+$); or else for A to side with one of B or C against the other (turning one of the edge labels out of A to a $-$).

- Similarly, there are sources of instability in a configuration where each of A, B, and C are mutual enemies (as in Figure 5.1(d)). In this case, there would be forces motivating two of the three people to "team up" against the third (turning one of the three edge labels to a $+$).

Based on this reasoning, we will refer to triangles with one or three $+$'s as *balanced*, since they are free of these sources of instability, and we will refer to triangles with zero or two $+$'s as *unbalanced*. The argument of structural balance theorists is that because unbalanced triangles are sources of stress or psychological dissonance, people strive to minimize them in their personal relationships, and hence they will be less abundant in real social settings than

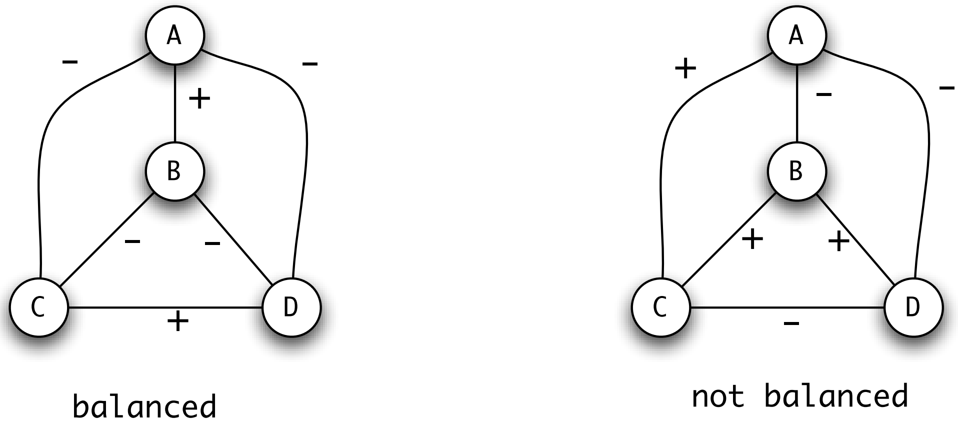

Figure 5.2: The labeled four-node complete graph on the left is balanced; the one on the right is not.

balanced triangles.

Defining Structural Balance for Networks. So far we have been talking about structural balance for groups of three nodes. But it is easy to create a definition that naturally generalizes this to complete graphs on an arbitrary number of nodes, with edges labeled by +'s and −'s.

Specifically, we say that a labeled complete graph is *balanced* if every one of its triangles is balanced — that is, if it obeys the following:

> *Structural Balance Property: For* every *set of three nodes, if we consider the three edges connecting them, either all three of these edges are labeled +, or else exactly one of them is labeled +.*

For example, consider the two labeled four-node networks in Figure 5.2. The one on the left is balanced, since we can check that each set of three nodes satisfies the Structural Balance Property above. On the other hand, the one on the right is not balanced, since among the three nodes A, B, C, there are exactly two edges labeled +, in violation of Structural Balance. (The triangle on B, C, D also violates the condition.)

Our definition of balanced networks here represents the limit of a social system that has eliminated all unbalanced triangles. As such, it is a fairly extreme definition — for example, one could instead propose a definition which only required that at least some large percentage of all triangles were balanced, allowing a few triangles to be unbalanced. But the version with all triangles balanced is a fundamental first step in thinking about this concept; and

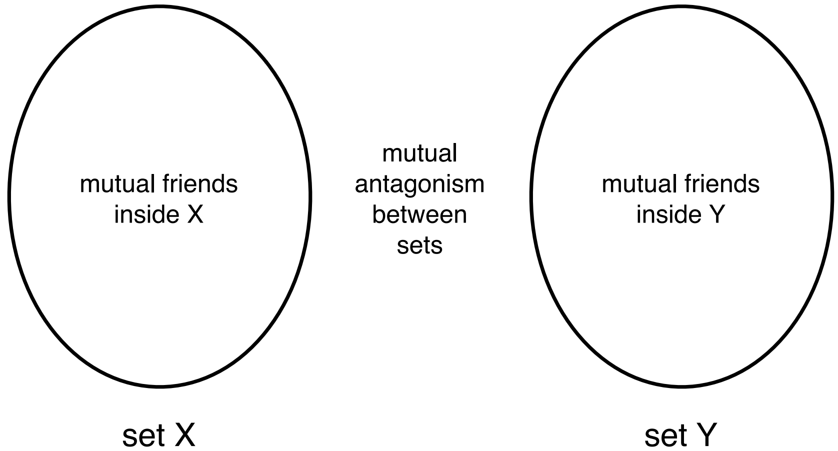

Figure 5.3: If a complete graph can be divided into two sets of mutual friends, with complete mutual antagonism between the two sets, then it is balanced. Furthermore, this is the only way for a complete graph to be balanced.

as we will see next, it turns out to have very interesting mathematical structure that in fact helps to inform the conclusions of more complicated models as well.

5.2 Characterizing the Structure of Balanced Networks

At a general level, what does a balanced network (i.e. a balanced labeled complete graph) look like? Given any specific example, we can check all triangles to make sure that they each obey the balance conditions; but it would be much better to have a simple conceptual description of what a balanced network looks like in general.

One way for a network to be balanced is if everyone likes each other; in this case, all triangles have three + labels. On the other hand, the left-hand side of Figure 5.2 suggests a slightly more complicated way for a network to be balanced: it consists of two groups of friends (A, B and C, D), with negative relations between people in different groups. This is actually true in general: suppose we have a labeled complete graph in which the nodes can be divided into two groups, X and Y, such that every pair of nodes in X like each other, every pair of nodes in Y like each other, and everyone in X is the enemy of everyone in Y. (See the schematic illustration in Figure 5.3.) You can check that such a network is balanced: a triangle contained entirely in one group or the other has three + labels, and a triangle with two people in one group and one in the other has exactly one + label.

So this describes two basic ways to achieve structural balance: either everyone likes each other; or the world consists of two groups of mutual friends with complete antagonism between the groups. The surprising fact is the following: these are the *only* ways to have a balanced network. We formulate this fact precisely as the following *Balance Theorem*, proved by Frank Harary in 1953 [97, 204]:

> *Balance Theorem: If a labeled complete graph is balanced, then either all pairs of nodes are friends, or else the nodes can be divided into two groups, X and Y, such that every pair of nodes in X like each other, every pair of nodes in Y like each other, and everyone in X is the enemy of everyone in Y.*

The Balance Theorem is not at all an obvious fact, nor should it be initially clear why it is true. Essentially, we're taking a purely *local* property, namely the Structural Balance Property, which applies to only three nodes at a time, and showing that it implies a strong *global* property: either everyone gets along, or the world is divided into two battling factions.

We're now going to show why this claim in fact is true.

Proving the Balance Theorem. Establishing the claim requires a proof: we're going to suppose we have an arbitrary labeled complete graph, assume only that it is balanced, and conclude that either everyone is friends, or that there are sets X and Y as described in the claim. Recall that we worked through a proof in Chapter 3 as well, when we used simple assumptions about triadic closure in a social network to conclude all local bridges in the network must be weak ties. Our proof here will be somewhat longer, but still very natural and straightforward — we use the definition of balance to directly derive the conclusion of the claim.

To start, suppose we have a labeled complete graph, and all we know is that it's balanced. We have to show that it has the structure in the claim. If it has no negative edges at all, then everyone is friends, and we're all set. Otherwise, there is at least one negative edge, and we need to somehow come up with a division of the nodes into sets of mutual friends X and Y, with complete antagonism between them. The difficulty is that, knowing so little about the graph itself other than that it is balanced, it's not clear how we're supposed to identify X and Y.

Let's pick any node in the network — we'll call it A — and consider things from A's perspective. Every other node is either a friend of A or an enemy of A. Thus, natural candidates to try for the sets X and Y would be to define X to be A and all its friends, and define Y to be all the enemies of A. This is indeed a division of all the nodes, since every node is either a friend or an enemy of A.

Recall what we need to show in order for these two sets X and Y to satisfy the conditions of the claim:

(i) Every two nodes in X are friends.

(ii) Every two nodes in Y are friends.

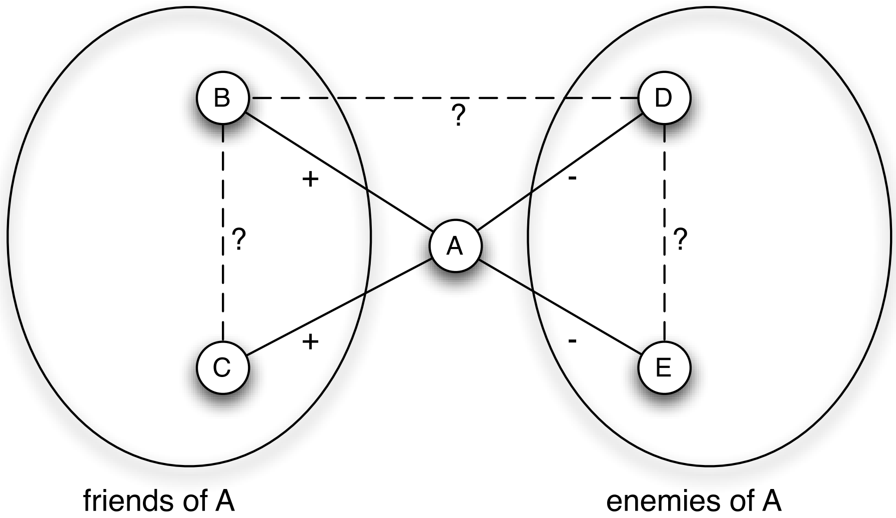

Figure 5.4: A schematic illustration of our analysis of balanced networks. (There may be other nodes not illustrated here.)

(iii) Every node in X is an enemy of every node in Y.

Let's argue that each of these conditions is in fact true for our choice of X and Y. This will mean that X and Y do satisfy the conditions of the claim, and will complete the proof. The rest of the argument, establishing (i), (ii), and (iii), is illustrated schematically in Figure 5.4.

For (i), we know that A is friends with every other node in X. How about two other nodes in X (let's call them B and C) — must they be friends? We know that A is friends with both B and C, so if B and C were enemies of each other, then A, B, and C would form a triangle with two + labels — a violation of the balance condition. Since we know the network is balanced, this can't happen, so it must be that B and C in fact are friends. Since B and C were the names of any two nodes in X, we have concluded that every two nodes in X are friends.

Let's try the same kind of argument for (ii). Consider any two nodes in Y (let's call them D and E) — must they be friends? We know that A is enemies with both D and E, so if D and E were enemies of each other, then A, D, and E would form a triangle with no + labels — a violation of the balance condition. Since we know the network is balanced, this can't happen, so it must be that D and E in fact are friends. Since D and E were the names of any two nodes in Y, we have concluded that every two nodes in Y are friends.

Finally, let's try condition (iii). Following the style of our arguments for (i) and (ii), consider a node in X (call if B) and a node in Y (call it D) — must they be enemies? We know A is friends with B and enemies with D, so if B and D were friends, then a, B, and

D would form a triangle with two + labels — a violation of the balance condition. Since we know the network is balanced, this can't happen, so it must be that B and D in fact are enemies. Since B and D were the names of any node in X and any node in Y, we have concluded that every such pair constitutes a pair of enemies.

So, in conclusion, assuming only that the network is balanced, we have described a division of the nodes into two sets X and Y, and we have checked conditions (i), (ii), and (iii) required by the claim. This completes the proof of the Balance Theorem.

5.3 Applications of Structural Balance

Structural balance has grown into a large area of study, and we've only described a simple but central example of the theory. In Section 5.5, we discuss two extensions to the basic theory: one to handle graphs that are not necessarily complete, and one to describe the structure of complete graphs that are "approximately balanced," in the sense that most but not all their triangles are balanced.

There has also been recent research looking at dynamic aspects of structural balance theory, modeling how the set of friendships and antagonisms in a complete graph — in other words, the labeling of the edges — might evolve over time, as the social network implicitly seeks out structural balance. Antal, Krapivsky, and Redner [20] study a model in which we start with a random labeling (choosing + or − randomly for each edge); we then repeatedly look for a triangle that is not balanced, and flip one of its labels to make it balanced. This captures a situation in which people continually reassess their likes and dislikes of others, as they strive for structural balance. The mathematics here becomes quite complicated, and turns out to resemble the mathematical models one uses for certain physical systems as they reconfigure to minimize their energy [20, 287].

In the remainder of this section, we consider two further areas in which the ideas of structural balance are relevant: international relations, where the nodes are different countries; and on-line social media sites where users can express positive or negative opinions about each other.

International Relations. International politics represents a setting in which it is natural to assume that a collection of nodes all have opinions (positive or negative) about one another — here the nodes are nations, and + and − labels indicate alliances or animosity. Research in political science has shown that structural balance can sometimes provide an effective explanation for the behavior of nations during various international crises. For example, Moore [306], describing the conflict over Bangladesh's separation from Pakistan in 1972, explicitly invokes structural balance theory when he writes, "[T]he United States's somewhat surprising support of Pakistan ... becomes less surprising when one considers that the USSR

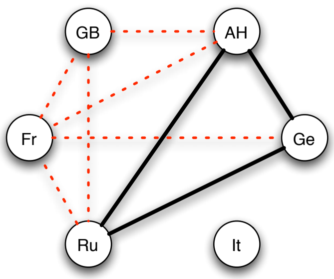

(a) *Three Emperors' League 1872–81*

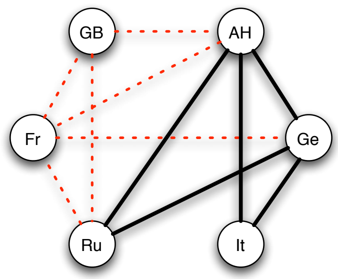

(b) *Triple Alliance 1882*

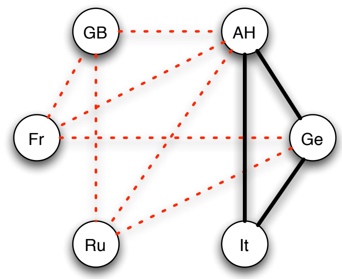

(c) *German-Russian Lapse 1890*

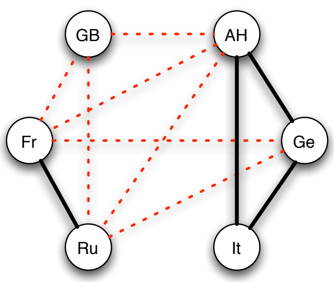

(d) *French-Russian Alliance 1891–94*

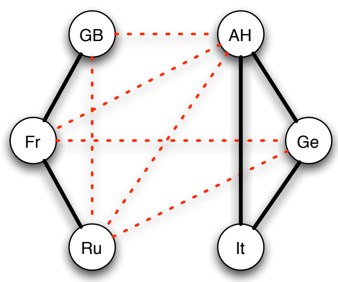

(e) *Entente Cordiale 1904*

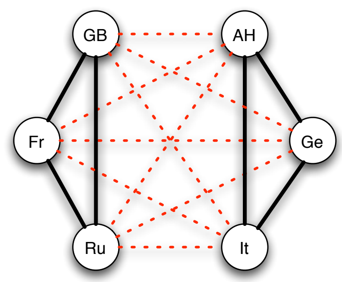

(f) *British Russian Alliance 1907*

Figure 5.5: The evolution of alliances in Europe, 1872-1907 (the nations GB, Fr, Ru, It, Ge, and AH are Great Britain, France, Russia, Italy, Germany, and Austria-Hungary respectively). Solid dark edges indicate friendship while dotted red edges indicate enmity. Note how the network slides into a balanced labeling — and into World War I. This figure and example are from Antal, Krapivsky, and Redner [20].

was China's enemy, China was India's foe, and India had traditionally bad relations with Pakistan. Since the U.S. was at that time improving its relations with China, it supported the enemies of China's enemies. Further reverberations of this strange political constellation became inevitable: North Vietnam made friendly gestures toward India, Pakistan severed diplomatic relations with those countries of the Eastern Bloc which recognized Bangladesh, and China vetoed the acceptance of Bangladesh into the U.N."

Antal, Krapivsky, and Redner use the shifting alliances preceding World War I as another example of structural balance in international relations — see Figure 5.5. This also reinforces the fact that structural balance is not necessarily a good thing: since its global outcome is often two implacably opposed alliances, the search for balance in a system can sometimes be seen as a slide into a hard-to-resolve opposition between two sides.

Trust, Distrust, and On-Line Ratings. A growing source for network data with both positive and negative edges comes from user communities on the Web where people can express positive or negative sentiments about each other. Examples include the technology news site Slashdot, where users can designate each other as a "friend" or a "foe" [266], and on-line product-rating sites such as Epinions, where a user can express evaluations of different products, and also express *trust* or *distrust* of other users.

Guha, Kumar, Raghavan, and Tomkins performed an analysis of the network of user evaluations on Epinions [201]; their work identified an interesting set of issues that show how the trust/distrust dichotomy in on-line ratings has both similarities and differences with the friend/enemy dichotomy in structural balance theory. One difference is based on a simple structural distinction: we have been considering structural balance in the context of undirected graphs, whereas user evaluations on a site like Epinions form a directed graph. That is, when a user A expresses trust or distrust of a user B, we don't necessarily know what B thinks of A, or whether B is even aware of A.

A more subtle difference between trust/distrust and friend/enemy relations becomes apparent when thinking about how we should expect triangles on three Epinions users to behave. Certain patterns are easy to reason about: for example, if user A trusts user B, and user B trusts user C, then it is natural to expect that A will trust C. Such triangles with three forward-pointing positive edges make sense here, by analogy with the all-positive (undirected) triangles of structural balance theory. But what if A distrusts B and B distrusts C: should we expect A to trust or to distrust C? There are appealing arguments in both directions. If we think of distrust as fundamentally a kind of enemy relationship, then the arguments from structural balance theory would suggest that A should trust C: otherwise we'd have a triangle with three negative edges. On the other hand, if A's distrust of B expresses A's belief that she is more knowledgeable and competent than B — and if B's distrust of C reflects a corresponding belief by B — then we might well expect that A will distrust C, and perhaps even more strongly than she distrusts B.

It is reasonable to expect that these two different interpretations of distrust may each apply, simply in different settings. And both might apply in the context of a single product-rating site like Epinions. For example, among users who are primarily rating best-selling books by political commentators, trust/distrust evaluations between users may become strongly aligned with agreement or disagreement in these users' own political orientations. In such a case, if A distrusts B and B distrusts C, this may suggest that A and C are close to each other on the underlying political spectrum, and so the prediction of structural balance theory that A should trust C may apply. On the other hand, among users who are primarily rating consumer electronics products, trust/distrust evaluations may largely reflect the relative expertise of users about the products (their respective features, reliability, and so forth). In such a case, if A distrusts B and B distrusts C, we might conclude that A is

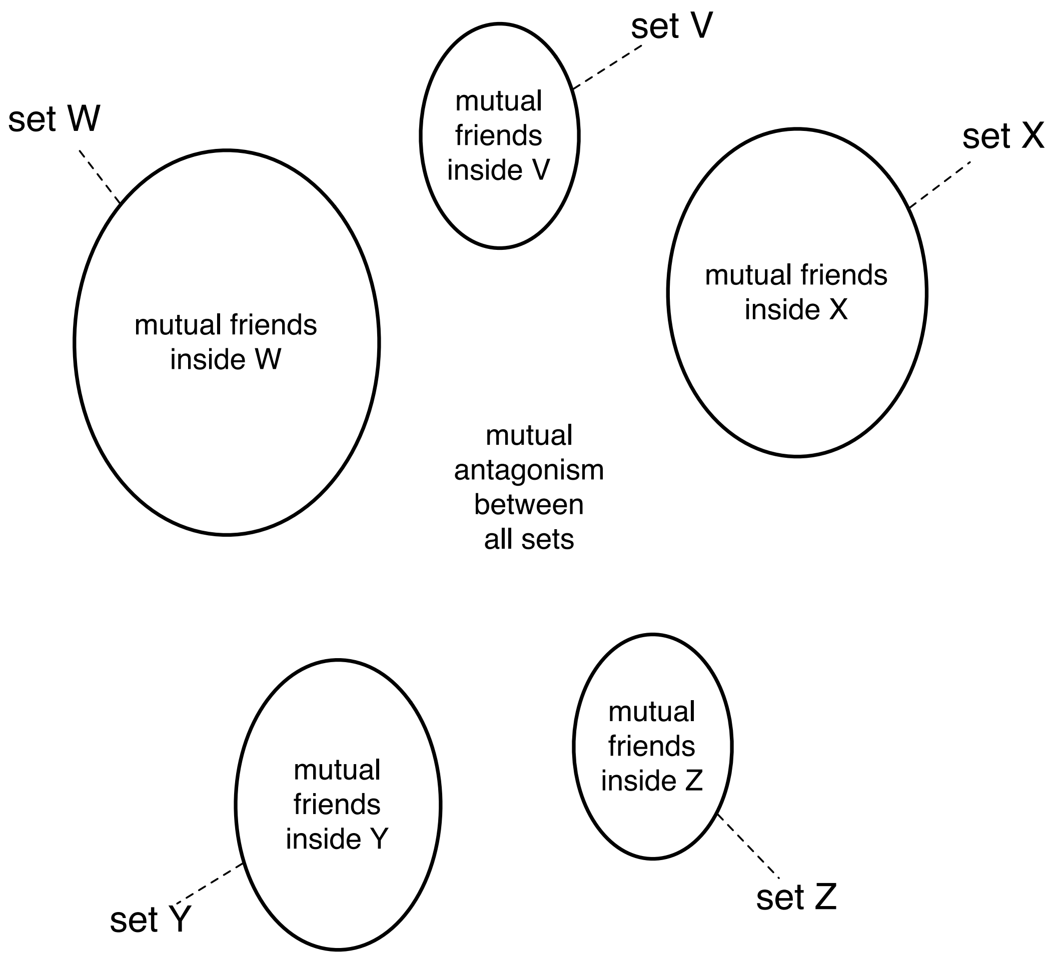

Figure 5.6: A complete graph is weakly balanced precisely when it can be divided into multiple sets of mutual friends, with complete mutual antagonism between each pair of sets.

far more expert than C, and so should distrust C as well.

Ultimately, understanding how these positive and negative relationships work is important for understanding the role they play on social Web sites where users register subjective evaluations of each other. Research is only beginning to explore these fundamental questions, including the ways in which theories of balance — as well as related theories — can be used to shed light on these issues in large-scale datasets [274].

5.4 A Weaker Form of Structural Balance

In studying models of positive and negative relationships on networks, researchers have also formulated alternate notions of structural balance, by revisiting the original assumptions we used to motivate the framework.

In particular, our analysis began from the claim that there are two kinds of structures on a group of three people that are inherently unbalanced: a triangle with two positive edges and one negative edge (as in Figure 5.1(b)); and a triangle with three negative edges (as in Figure 5.1(d)). In each of these cases, we argued that the relationships within the triangle contained a latent source of stress that the network might try to resolve. The underlying arguments in the two cases, however, were fundamentally different. In a triangle with two positive edges, we have the problem of a person whose two friends don't get along; in a triangle with three negative edges, there is the possibility that two of the nodes will ally themselves against the third.

James Davis and others have argued that in many settings, the first of these factors may be significantly stronger than the second [127]: we may see friends of friends trying to reconcile their differences (resolving the lack of balance in Figure 5.1(b)), while at the same time there could be less of a force leading any two of three mutual enemies (as in Figure 5.1(d)) to become friendly. It therefore becomes natural to ask what structural properties arise when we rule out only triangles with exactly two positive edges, while allowing triangles with three negative edges to be present in the network.

Characterizing Weakly Balanced Networks. More precisely, we will say that a complete graph, with each edge labeled by $+$ or $-$, is *weakly balanced* if the following property holds.

> *Weak Structural Balance Property: There is no set of three nodes such that the edges among them consist of exactly two positive edges and one negative edge.*

Since weak balance imposes less of a restriction on what the network can look like, we should expect to see a broader range of possible structures for weakly balanced networks — beyond what the Balance Theorem required for networks that were balanced under our original definition. And indeed, Figure 5.6 indicates a new kind of structure that can arise. Suppose that the nodes can be divided into an arbitrary number of groups (possibly more than two), so that two nodes are friends when they belong to the same group, and enemies when they belong to different groups. Then we can check that such a network is weakly balanced: in any triangle that contains at least two positive edges, all three nodes must belong to the same group. Therefore, the third edge of this triangle must be positive as well — in other words, the network contains no triangles with exactly two $+$ edges.

Just as the Balance Theorem established that all balanced networks must have a simple structure, an analogous result holds for weakly balanced networks: they must have the structure depicted in Figure 5.6, with any number of groups.

> *Characterization of Weakly Balanced Networks: If a labeled complete graph is weakly balanced, then its nodes can be divided into groups in such a way that*

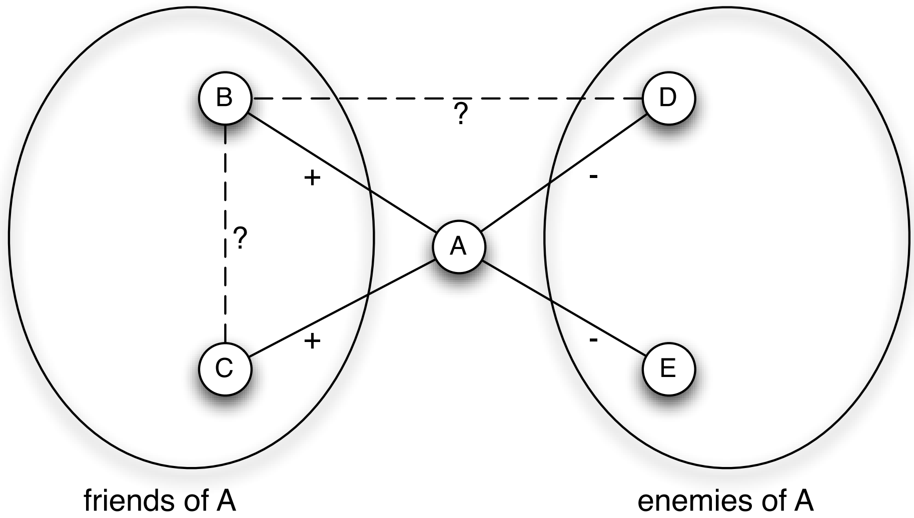

Figure 5.7: A schematic illustration of our analysis of weakly balanced networks. (There may be other nodes not illustrated here.)

> *every two nodes belonging to the same group are friends, and every two nodes belonging to different groups are enemies.*

The fact that this characterization is true in fact provided another early motivation for studying weak structural balance. The Cartwright-Harary notion of balance predicted only dichotomies (or mutual consensus) as its basic social structure, and thus did not provide a model for reasoning about situations in which a network is divided into more than two factions. Weak structural balance makes this possible, since weakly balanced complete graphs can contain any number of opposed groups of mutual friends [127].

Proving the Characterization. It is not hard to give a proof for this characterization, following the structure of our proof for the Balance Theorem, and making appropriate changes where necessary. Starting with a weakly balanced complete graph, the characterization requires that we produce a division of its nodes into groups of mutual friends, such that all relations between nodes in different groups are negative. Here is how we will construct this division.

First, we pick any node A, and we consider the set consisting of A and all its friends. Let's call this set of nodes X. We'd like to make X our first group, and for this to work, we need to establish two things:

(i) All of A's friends are friends with each other. (This way, we have indeed produced a group of mutual friends).

(ii) A and all his friends are enemies with everyone else in the graph. (This way, the people in this group will be enemies with everyone in other groups, however we divide up the rest of the graph.)

Fortunately, ideas that we already used inside the proof of the Balance Theorem can be adapted to our new setting here to establish (i) and (ii). The idea is shown in Figure 5.7. First, for (i), let's consider two nodes B and C who are both friends with A. If B and C were enemies of each other, then the triangle on nodes A, B, and C would have exactly two $+$ labels, which would violate weak structural balance. So B and C must indeed be friends with each other.

For (ii), we know that A is enemies with all nodes in the graph outside X, since the group X is defined to include all of A's friends. How about an edge between a node B in X and a node D outside X? If B and D were friends, then the triangle on nodes A, B, and D would have exactly two $+$ labels — again, a violation of weak structural balance. So B and D must be enemies.

Since properties (i) and (ii) hold, we can remove the set X — consisting of A and all his friends — from the graph and declare it to be the first group. We now have a smaller complete graph that is still weakly balanced; we find a second group in this graph, and proceed to remove groups in this way until all the nodes have been assigned to a group. Since each group consists of mutual friends (by property (i)), and each group has only negative relations with everyone outside the group (by property (ii)), this proves the characterization.

It is interesting to reflect on this proof in relation to the proof of the Balance Theorem — in particular, the contrast reflected by the small differences between Figures 5.4 and 5.7. In proving the Balance Theorem, we had to reason about the sign of the edge between D and E, to show that the enemies of the set X themselves formed a set Y of mutual friends. In characterizing weakly balanced complete graphs, on the other hand, we made no attempt to reason about the D-E edge, because weak balance imposes no condition on it: two enemies of A can be either friends or enemies. As a result, the set of enemies in Figure 5.7 might not be a set of mutual friends when only weak balance holds; it might consist of multiple groups of mutual friends, and as we extract these groups one by one over the course of the proof, we recover the multi-faction structure illustrated schematically in Figure 5.6.

5.5 Advanced Material: Generalizing the Definition of Structural Balance

In this section, we consider more general ways of formulating the idea of structural balance in a network. In particular, our definition of structural balance thus far is fairly demanding in two respects:

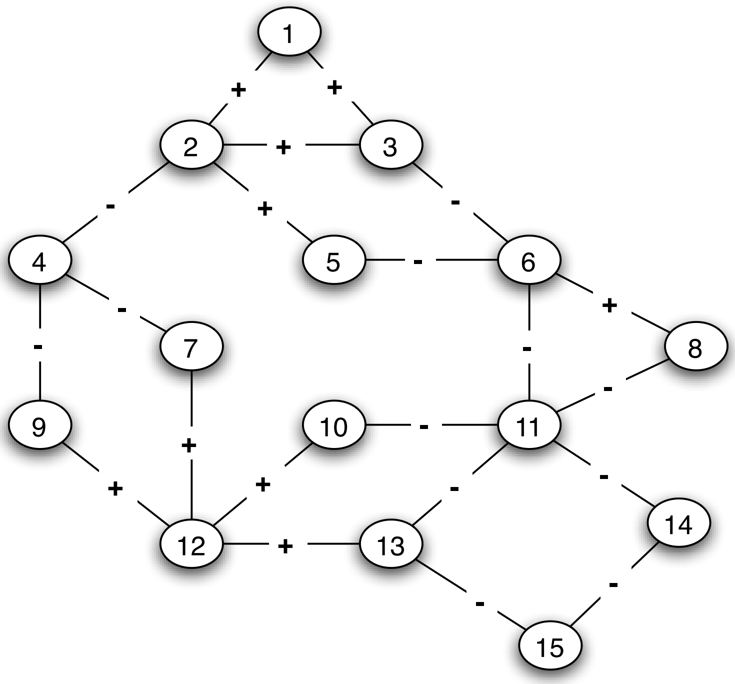

Figure 5.8: In graphs that are not complete, we can still define notions of structural balance when the edges that are present have positive or negative signs indicating friend or enemy relations.

1. It applies only to complete graphs: we require that each person know and have an opinion (positive or negative) on everyone else. What if only some pairs of people know each other?

2. The Balance Theorem, showing that structural balance implies a global division of the world into two factions [97, 204], only applies to the case in which *every* triangle is balanced. Can we relax this to say that if *most* triangles are balanced, then the world can be *approximately* divided into two factions?

In the two parts of this section, we discuss a pair of results that address these questions. The first is based on a graph-theoretic analysis involving the notion of breadth-first search from Chapter 2, while the second is typical of a style of proof known as a "counting argument." Throughout this section, we will focus on the original definition of structural balance from Sections 5.1 and 5.2, rather than the weaker version from Section 5.4.

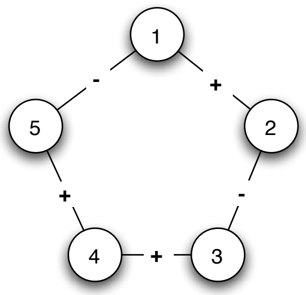

(a) *A graph with signed edges.*

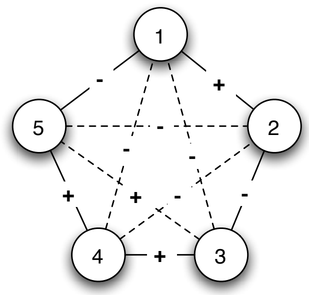

(b) *Filling in the missing edges to achieve balance.*

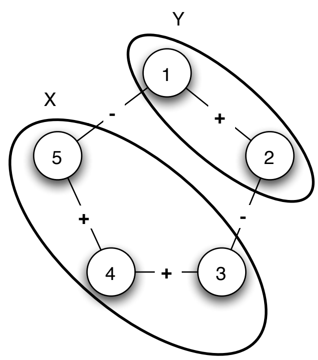

(c) *Dividing the graph into two sets.*

Figure 5.9: There are two equivalent ways to define structural balance for general (non-complete) graphs. One definition asks whether it is possible to fill in the remaining edges so as to produce a signed complete graph that is balanced. The other definition asks whether it is possible to divide the nodes into two sets X and Y so that all edges inside X and inside Y are positive, and all edges between X and Y are negative.

A. Structural Balance in Arbitrary (Non-Complete) Networks

First, let's consider the case of a social network that is not necessarily complete — that is, there are only edges between certain pairs of nodes, but each of these edges is still labeled with $+$ or $-$. So now there are three possible relations between each pair of nodes: a positive edge, indicating friendship; a negative edge, indicating enmity; or the absence of an edge, indicating that the two endpoints do not know each other. Figure 5.8 depicts an example of such a signed network.

Defining Balance for General Networks. Drawing on what we've learned from the special case of complete graphs, what would be a good definition of balance for this more general kind of structure? The Balance Theorem suggests that we can view structural balance in either of two equivalent ways: a *local* view, as a condition on each triangle of the network; or a *global* view, as a requirement that the world be divided into two mutually opposed sets of friends. Each of these suggests a way of defining structure balance for general signed graphs.

1. One option would be to treat balance for non-complete networks as a problem of filling in "missing values." Suppose we imagine, as a thought experiment, that all people in the group in fact do know and have an opinion on each other; the graph under consideration is not complete only because we have failed to observe the relations between some of the pairs. We could then say that the graph is balanced if it possible to fill in all the missing labeled edges in such a way that the resulting signed complete graph is balanced. In other words, a (non-complete) graph is balanced if it can be "completed" by adding edges to form a signed complete graph that is balanced.

   For example, Figure 5.9(a) shows a graph with signed edges, and Figure 5.9(b) shows how the remaining edges can be "filled in" to produce a balanced complete graph: we declare the missing edge between nodes 3 and 5 to be positive, and the remaining missing edges to be negative, and one can check that this causes all triangles to be balanced.

2. Alternately, we could take a more global view, viewing structural balance as implying a division of the network into two mutually opposed sets of friends. With this in mind, we could define a signed graph to be balanced if it is possible to divide the nodes into two sets X and Y, such that any edge with both ends inside X or both ends inside Y is positive, and any edge with one end in X and the other in Y is negative. That is, people in X are all mutual friends to the extent that they know each other; the same is true for people in Y; and people in X are all enemies of people in Y to the extent that they know each other.

   Continuing the example from Figure 5.9(a), in Figure 5.9(c) we show how to divide this graph into two sets with the desired properties.

This example hints at a principle that is true in general: these two ways of defining balance are equivalent. An arbitrary signed graph is balanced under the first definition if and only if it is balanced under the second definition.

This is actually not hard to see. If a signed graph is balanced under the first definition, then after filling in all the missing edges appropriately, we have a signed complete graph to which we can apply the Balance Theorem. This gives us a division of the network into two sets X and Y that satisfies the properties of the second definition. On the other hand, if a signed graph is balanced under the second definition, then after finding a division of the nodes into sets X and Y, we can fill in positive edges inside X and inside Y, and fill in

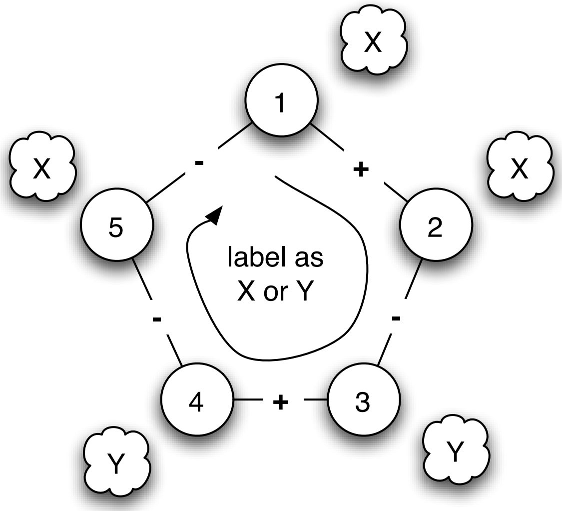

Figure 5.10: If a signed graph contains a cycle with an odd number of negative edges, then it is not balanced. Indeed, if we pick one of the nodes and try to place it in X, then following the set of friend/enemy relations around the cycle will produce a conflict by the time we get to the starting node.

negative edges between X and Y, and then we can check that all triangles will be balanced. So this gives a "filling-in" that satisfies the first definition.

The fact that the two definitions are equivalent suggests a certain "naturalness" to the definition, since there are fundamentally different ways to arrive at it. It also lets us use either definition, depending on which is more convenient in a given situation. As the example in Figure 5.9 suggests, the second definition is generally more useful to work with — it tends to be much easier to think about dividing the nodes into two sets than to reason about filling in edges and checking triangles.

Characterizing Balance for General Networks. Conceptually, however, there is something not fully satisfying about either definition: the definitions themselves do not provide much insight into how to easily check that a graph is balanced. There are, after all, lots of ways to choose signs for the missing edges, or to choose ways of splitting the nodes into sets X and Y. And if a graph is not balanced, so that there is no way to do these things successfully, what could you show someone to convince them of this fact? To take just a small example to suggest some of the difficulties, it may not be obvious from a quick inspection of Figure 5.8 that this is not a balanced graph — or that if we change the edge connecting nodes 2 and 4 to be positive instead of negative, it becomes a balanced graph.

In fact, however, all these problems can be remedied if we explore the consequences of the definitions a little further. What we will show is a simple characterization of balance in general signed graphs, also due to Harary [97, 204]; and the proof of this characterization also provides an easy method for checking whether a graph is balanced.

The characterization is based on considering the following question: what prevents a graph from being balanced? Figure 5.10 shows a graph that is not balanced (obtained from Figure 5.9(a) and changing the sign of the edge from node 4 to node 5). It also illustrates a reason why it's not balanced, as follows. If we start at node 1 and try to divide the nodes into sets X and Y, then our choices are forced at every step. Suppose we initially decide that node 1 should belong to X. (For the first node, it doesn't matter, by symmetry.) Then since node 2 is friends with node 1, it too must belong to X. Node 3, an enemy of 2, must therefore belong to Y; hence node 4, a friend of 3, must belong to Y as well; and node 5, an enemy of 4, must belong to X. The problem is that if we continue this reasoning one step further, then node 1, an enemy of 5, should belong to Y — but we had already decided at the outset to put it into X. We had no freedom of choice during this process — so this shows that there is no way to divide the nodes in sets X and Y so as to satisfy the mutual-friend/mutual-enemy conditions of structural balance, and hence the signed graph in Figure 5.10 is not balanced.

The reasoning in the previous paragraph sounds elaborate, but in fact it followed a simple principle: we were walking around a cycle, and every time we crossed a negative edge, we had to change the set into which we were putting nodes. The difficulty was that getting back around to node 1 required crossing an *odd number* of negative edges, and so our original decision to put node 1 into X clashed with the eventual conclusion that node 1 ought to be in Y.

This principle applies in general: if the graph contains a cycle with an odd number of negative edges, then this implies the graph is not balanced. Indeed, if we start at any node A in the cycle and place it in one of the two sets, and then we walk around the cycle placing the other nodes where they must go, the identity of the set where we're placing nodes switches an odd number of times as we go around the cycle. Thus we end up with the "wrong set" by the time we make it back to A.

A cycle with an odd number of negative edges is thus a very simple-to-understand reason why a graph is not balanced: you can show someone such a cycle and immediately convince them that the graph is not balanced. For example, the cycle back in Figure 5.8 consisting of nodes, 2, 3, 6, 11, 13, 12, 9, 4 contains five negative edges, thus supplying a succinct reason why this graph is not balanced. But are there other, more complex reasons why a graph is not balanced?

In fact, though it may seem initially surprising, cycles with an odd number of negative edges are the only obstacles to balance. This is the crux of the following claim [97, 204].

Claim: A signed graph is balanced if and only if it contains no cycle with an odd

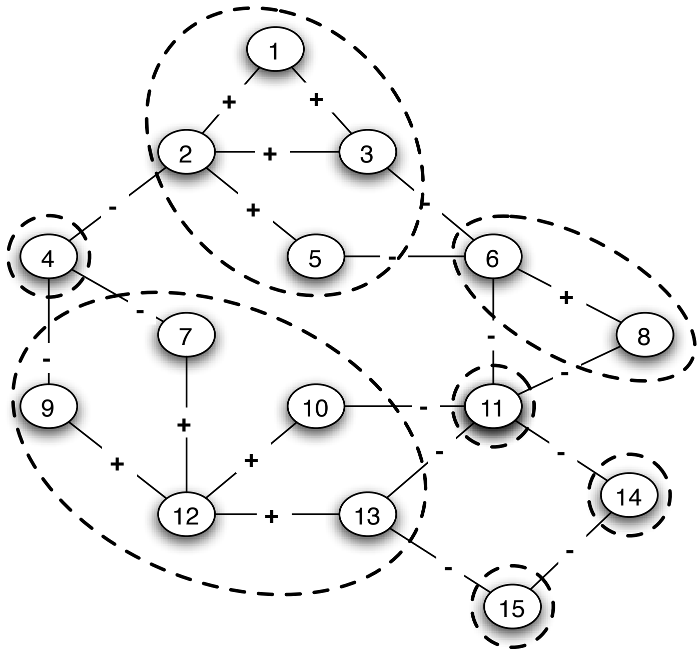

Figure 5.11: To determine if a signed graph is balanced, the first step is to consider only the positive edges, find the connected components using just these edges, and declare each of these components to be a *supernode*. In any balanced division of the graph into X and Y, all nodes in the same supernode will have to go into the same set.

number of negative edges.

We now show how to prove this claim; this is done by designing a method that analyzes the graph and either finds a division into the desired sets X and Y, or else finds a cycle with an odd number of negative edges.

Proving the Characterization: Identifying Supernodes. Let's recall what we're trying to do: find a division of the nodes into sets X and Y so that all edges inside X and Y are positive, and all edges crossing between X and Y are negative. We will call a partitioning into sets X and Y with these properties a *balanced division*. We now describe a procedure that searches for a balanced division of the nodes into sets X and Y; either it succeeds, or it stops with a cycle containing an odd number of negative edges. Since these are the only two possible outcomes for the procedure, this will give a proof of the claim.

The procedure works in two main steps: the first step is to convert the graph to a reduced one in which there are only negative edges, and the second step is to solve the problem on this reduced graph. The first step works as follows. Notice that whenever two nodes are

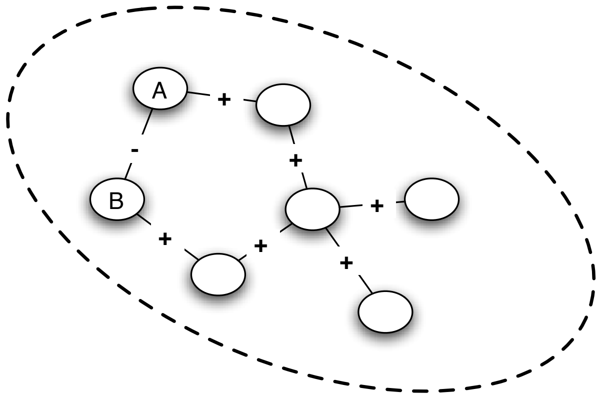

Figure 5.12: Suppose a negative edge connects two nodes A and B that belong to the same supernode. Since there is also a path consisting entirely of positive edges that connects A and B through the inside of the supernode, putting this negative edge together with the all-positive path produces a cycle with an odd number of negative edges.

connected by a positive edge, they must belong to the same one of the sets X or Y in a balanced division. So we begin by considering what the connected components of the graph would be if we were to only consider positive edges. These components can be viewed as a set of contiguous "blobs" in the overall graph, as shown in Figure 5.11. We will refer to each of these blobs as a *supernode*: each supernode is connected internally via positive edges, and the only edges going between two different supernodes are negative. (If there were a positive edge linking two different supernodes, we should have combined them together into a single supernode.)

Now, if any supernode contains a negative edge between some pair of nodes A and B, then we already have a cycle with an odd number of negative edges, as illustrated in the example of Figure 5.12. Consider the path of positive edges that connects A and B inside the supernode, and then close off a cycle by including the negative edge joining A and B. This cycle has only a single negative edge, linking A and B, and so it shows that the graph is not balanced.

If there are no negative edges inside any of the supernodes, then there is no "internal" problem with declaring each supernode to belong entirely to one of X or Y. So the problem is now how to assign a single label "X" or "Y" to each supernode, in such a way that these choices are all consistent with each other Since the decision-making is now at the level of supernodes, we create a new version of the problem in which there is a node for each

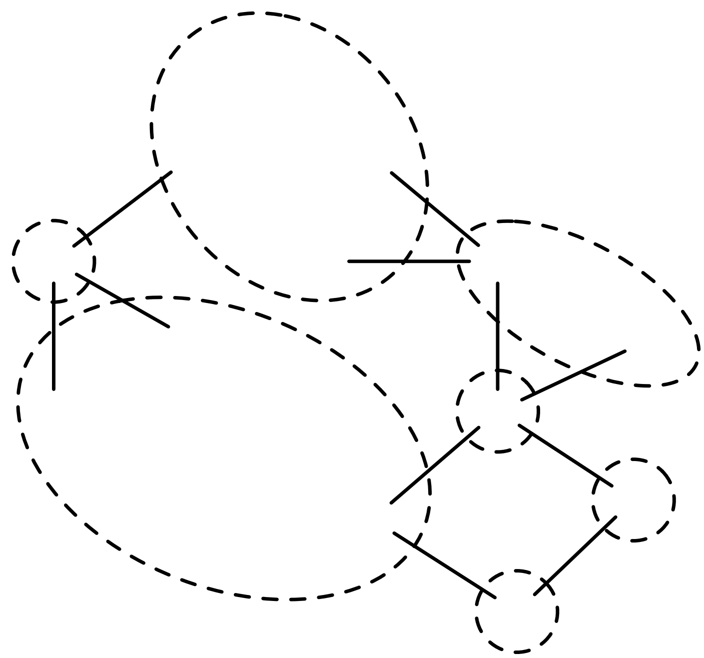

Figure 5.13: The second step in determining whether a signed graph is balanced is to look for a labeling of the supernodes so that adjacent supernodes (which necessarily contain mutual enemies) get opposite labels. For this purpose, we can ignore the original nodes of the graph and consider a *reduced graph* whose nodes are the supernodes of the original graph.

supernode, and an edge joining two supernodes if there is an edge in the original that connects the two supernodes. Figure 5.13 shows how this works for the example of Figure 5.11: we essentially forget about the individual nodes inside the supernodes, and build a new graph at the level of the large "blobs." Of course, having done so, we can draw the graph in a less blob-like way, as in Figure 5.14.

We now enter the second step of the procedure, using this *reduced graph* whose nodes are the supernodes of the original graph.

Proving the Characterization: Breadth-First Search of the Reduced Graph. Recall that only negative edges go between supernodes (since a positive edge between two supernodes would have merged them together into a single one). As a result, our reduced graph has only negative edges. The remainder of the procedure will produce one of two possible outcomes.

1. The first possible outcome is to label each node in the reduced graph as either X or Y, in such a way that every edge has endpoints with opposite labels. From this we

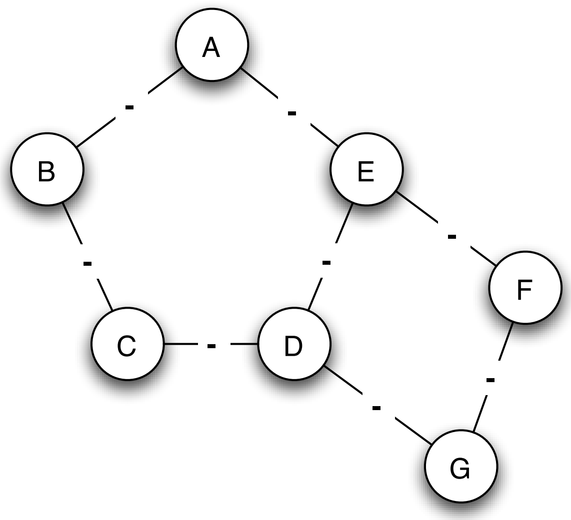

Figure 5.14: A more standard drawing of the reduced graph from the previous figure. A negative cycle is visually apparent in this drawing.

can create a balanced division of the original graph, by labeling each node the way its supernode is labeled in the reduced graph.

2. The second possible outcome will be to find a cycle in the reduced graph that has an odd number of edges. We can then convert this to a (potentially longer) cycle in the original graph with an odd number of negative edges: the cycle in the reduced graph connects supernodes, and corresponds to a set of negative edges in the original graph. We can simply "stitch together" these negative edges using paths consisting entirely of positive edges that go through the insides of the supernodes. This will be a path containing an odd number of negative edges in the original graph.

   For example, the odd-length cycle in Figure 5.14 through nodes A through E can be realized in the original graph as the darkened negative edges shown in Figure 5.15. This can then be turned into a cycle in the original graph by including paths through the supernodes – in this example using the additional nods 3 and 12.

In fact, this version of the problem when there are only negative edges is known in graph theory as the problem of determining whether a graph is *bipartite*: whether its nodes can be divided into two groups (in this case X and Y) so that each edge goes from one group to the other. We saw bipartite graphs when we considered affiliation networks in Chapter 4, but there the fact that the graphs were bipartite was apparent from the ready-made division of the nodes into people and social foci. Here, on the other hand, we are handed a graph "in

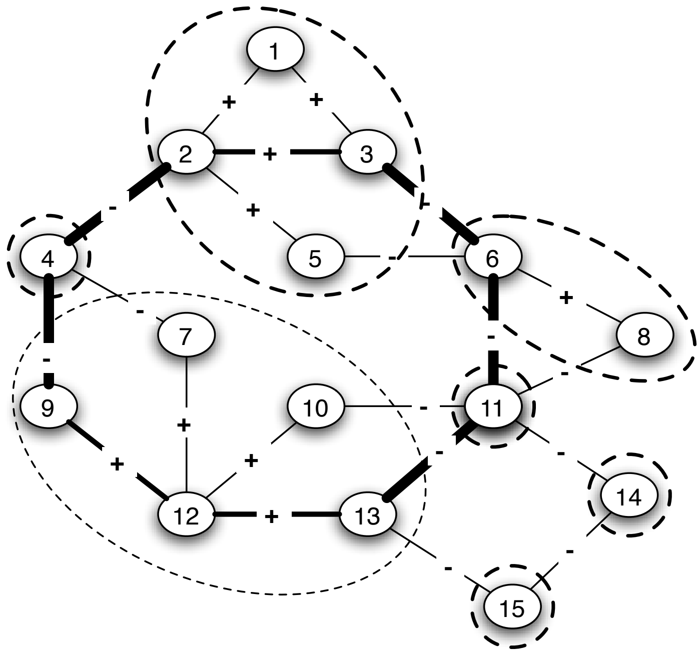

Figure 5.15: Having found a negative cycle through the supernodes, we can then turn this into a cycle in the original graph by filling in paths of positive edges through the inside of the supernodes. The resulting cycle has an odd number of negative edges.

the wild," with no pre-specified division into two sets, and we want to know if it is possible to identify such a division. We now show a way to do this using the idea of breadth-first search from Chapter 2, resulting either in the division we seek, or in a cycle of odd length.

We simply perform breadth-first search starting from any "root" node in the graph, producing a set of layers at increasing distances from this root. Figure 5.16 shows how this is done for the reduced graph in Figure 5.14, with node G as the starting root node. Now, because edges cannot jump over successive layers in breadth-first search, each edge either connects two nodes in adjacent layers or it connects two nodes in the same layer. If all edges are of the first type, then we can find the desired division of nodes into sets X and Y: we simply declare all nodes in even-numbered layers to belong to X, and all nodes in odd-numbered layers to belong to Y. Since edges only go between adjacent layers, all edges have one end in X and the other end in Y, as desired.

Otherwise, there is an edge connecting two nodes that belong to the same layer. Let's call them A and B (as they are in Figure 5.16). For each of these two nodes, there is a path that descends layer-by-layer from the root to it. Consider the last node that is common to these two paths — let's call this node D (as it is in Figure 5.16). The D-A path and the

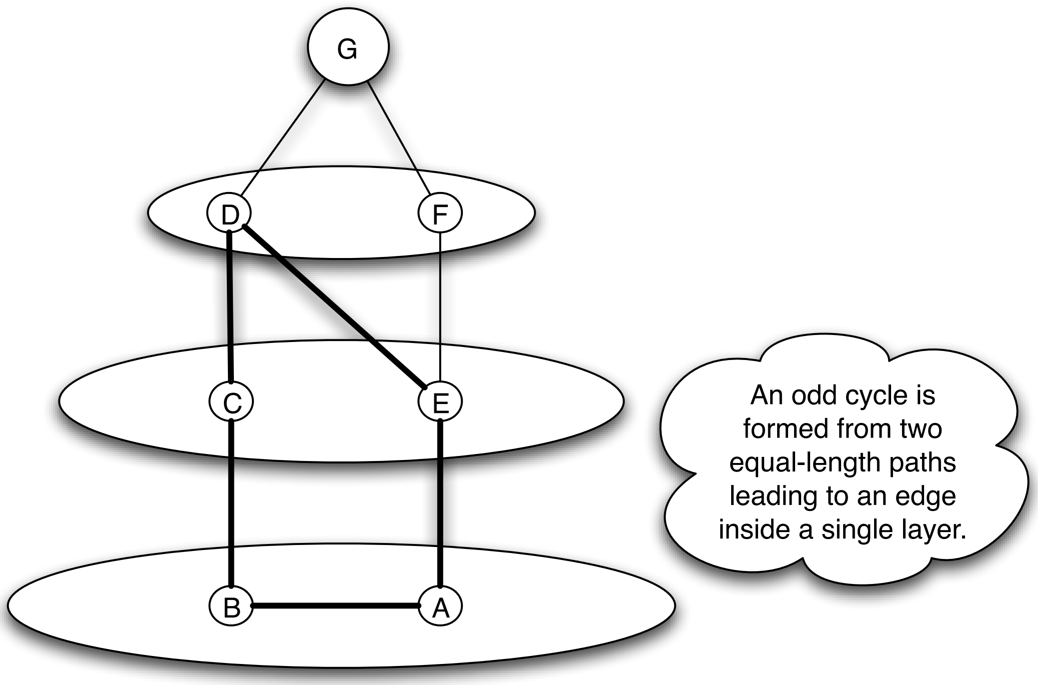

Figure 5.16: When we perform a breadth-first search of the reduced graph, there is either an edge connecting two nodes in the same layer or there isn't. If there isn't, then we can produce the desired division into X and Y by putting alternate layers in different sets. If there is such an edge (such as the edge joining A and B in the figure), then we can take two paths of the same length leading to the two ends of the edge, which together with the edge itself forms an odd cycle.

D-B path have the same length k, so a cycle created from the two of these plus the A-B edge must have length $2k + 1$: an odd number. This is the odd cycle we seek.

And this completes the proof. To recap: if all edges in the reduced graph connect nodes in adjacent layers of the breadth-first search, then we have a way to label the nodes in the reduced graph as into X and Y, which in turn provides a balanced division of the nodes in the original graph into X and Y. In this case, we've established that the graph is balanced. Otherwise, there is an edge connecting two nodes in the same layer of the breadth-first search, in which case we produce an odd cycle in the reduced graph as in Figure 5.16. In this case, we can convert into this to a cycle in the original graph containing an odd number of negative edges, as in Figure 5.15. Since these are the only two possibilities, this proves the claim.

B. Approximately Balanced Networks

We now return to the case in which the graph is complete, so that every node has a positive or negative relation with every other node, and we think about a different way of generalizing the characterization of structural balance.

First let's write down the original Balance Theorem again, with some additional formatting to make its logical structure clear.

> *Claim: If all triangles in a labeled complete graph are balanced, then either*
>
> *(a) all pairs of nodes are friends, or else*
>
> *(b) the nodes can be divided into two groups, X and Y, such that*
>
> *(i) every pair of nodes in X like each other,*
>
> *(ii) every pair of nodes in Y like each other, and*
>
> *(iii) everyone in X is the enemy of everyone in Y.*

The conditions of this theorem are fairly extreme, in that we require every single triangle to be balanced. What if we only know that most triangles are balanced? It turns out that the conditions of the theorem can be relaxed in a very natural way, allowing us to prove statements like the following one. We phrase it so that the wording remains completely parallel to that of the Balance Theorem.

> *Claim: If at least 99.9% of all triangles in a labeled complete graph are balanced, then either*
>
> *(a) there is a set consisting of at least 90% of the nodes in which at least 90% of all pairs are friends, or else*
>
> *(b) the nodes can be divided into two groups, X and Y, such that*
>
> *(i) at least 90% of the pairs in X like each other,*
>
> *(ii) at least 90% of the pairs in Y like each other, and*
>
> *(iii) at least 90% of the pairs with one end in X and the other end in Y are enemies.*

This is a true statement, though the choice of numbers is very specific. Here is a more general statement that includes both the Balance Theorem and the preceding claim as special cases.

> *Claim: Let ε be any number such that $0 \leq \varepsilon < \frac{1}{8}$, and define $\delta = \sqrt[3]{\varepsilon}$. If at least $1 - \varepsilon$ of all triangles in a labeled complete graph are balanced, then either*
>
> *(a) there is a set consisting of at least $1 - \delta$ of the nodes in which at least $1 - \delta$ of all pairs are friends, or else*

(b) the nodes can be divided into two groups, X and Y, such that

(i) at least $1 - \delta$ of the pairs in X like each other,

(ii) at least $1 - \delta$ of the pairs in Y like each other, and

(iii) at least $1 - \delta$ of the pairs with one end in X and the other end in Y are enemies.

Notice that the Balance Theorem is the case in which $\varepsilon = 0$, and the other claim above is the case in which $\varepsilon = .001$ (since in this latter case, $\delta = \sqrt[3]{\varepsilon} = .1$).

We now prove this last claim. The proof is self-contained, but it is most easily read with some prior experience in what is sometimes called the analysis of "permutations and combinations" — counting the number of ways to choose particular subsets of larger sets.

The proof loosely follows the style of the proof we used for the Balance Theorem: we will define the two sets X and Y to be the friends and enemies, respectively, of a designated node A. Things are trickier here, however, because not all choices of A will give us the structure we need — in particular, if a node is personally involved in too many unbalanced triangles, then splitting the graph into its friends and enemies may give a very disordered structure. Consequently, the proof consists of two steps. We first find a "good" node that is not involved in too many unbalanced triangles. We then show that if we divide the graph into the friends and enemies of this good node, we have the desired properties.

Warm-Up: Counting Edges and Triangles. Before launching into the proof itself, let's consider some basic counting questions that will show up as ingredients in the proof. Recall that we have a complete graph, with an (undirected) edge joining each pair of nodes. If N is the number of nodes in the graph, how many edges are there? We can count this quantity as follows. There are N possible ways to choose one of the two endpoints, and then $N-1$ possible ways to choose a different node as the other endpoint, for a total of $N(N-1)$ possible ways to choose the two endpoints in succession. If we write down a list of all these possible pairs of endpoints, then an edge with endpoints A and B will appear twice on the list: once as AB and once as BA. In general, each edge will appear twice on the list, and so the total number of edges is $N(N-1)/2$.

A very similar argument lets us count the total number of triangles in the graph. Specifically, there are N ways to pick the first corner, then $N-1$ ways to pick a different node as the second corner, and then $N-2$ ways to pick a third corner different from the first two. This yields a total of $N(N-1)(N-2)$ sequences of three corners. If we write down this list of $N(N-1)(N-2)$ sequences, then a triangle with corners A, B, and C will appear six times: as ABC, ACB, BAC, BCA, CAB, and CBA. In general, each triangle will appear six times in this list, and so the total number of triangles is

$$\frac{N(N-1)(N-2)}{6}.$$

The First Step: Finding a "Good" Node. Now let's move on to the first step of the proof, which is to find a node that isn't involved in too many unbalanced triangles.

Since we are assuming that at most an ε fraction of triangles are unbalanced, and the total number of triangles in the graph is $N(N-1)(N-2)/6$, it follows that the total number of unbalanced triangles is at most $\varepsilon N(N-1)(N-2)/6$. Suppose we define the *weight* of a node to be the number of unbalanced triangles that it is a part of; thus, a node of low weight will be precisely what we're seeking — a node that is in relatively few unbalanced triangles.

One way to count the total weight of all nodes would be to list — for each node — the unbalanced triangles that it belongs to, and then look at the length of all these lists combined. In these combined lists, each triangle will appear three times — once in the list for each of its corners — and so the total weight of all nodes is exactly three times the number of unbalanced triangles. As a result, the total weight of all nodes is at most $3\varepsilon N(N-1)(N-2)/6 = \varepsilon N(N-1)(N-2)/2$.

There are N nodes, so the *average* weight of a node is at most $\varepsilon(N-1)(N-2)/2$. It's not possible for all nodes to have weights that are strictly above the average, so there is at least one node whose weight is equal to the average or below it. Let's pick one such node and call it A. This will be our "good" node: a node whose weight is at most $\varepsilon(N-1)(N-2)/2$.[2] Since $(N-1)(N-2) < N^2$, this good node is in at most $\varepsilon N^2/2$ triangles, and because the algebra is a bit simpler with this slightly larger quantity, we will use it in the rest of the analysis.

The Second Step: Splitting the Graph According to the Good Node. By analogy with the proof of the Balance Theorem, we divide the graph into two sets: a set X consisting of A and all its friends, and a set Y consisting of all the enemies of A, as illustrated in Figure 5.17. Now, using the definition of unbalanced triangles, and the fact that node A is not involved in too many of them, we can argue that there are relatively few negative edges inside each of X and Y, and relatively few positive edges between them. Specifically, this works as follows.

- Each negative edge connecting two nodes in X creates a distinct unbalanced triangle involving node A. Since there are at most $\varepsilon N^2/2$ unbalanced triangles involving A, there are at most $\varepsilon N^2/2$ negative edges inside X.

- A closely analogous argument applies to Y: Each negative edge connecting two nodes in Y creates a distinct unbalanced triangle involving node A, and so there are at most $\varepsilon N^2/2$ negative edges inside Y.

[2] This is a very common trick in counting arguments, referred to as the *pigeonhole principle*: to compute the average value of a set of objects, and then argue that there must be at least one node that is equal to the average or below. (Also, of course, there must be at least at least one object that is equal to the average or above, although this observation isn't useful for our purposes here.)

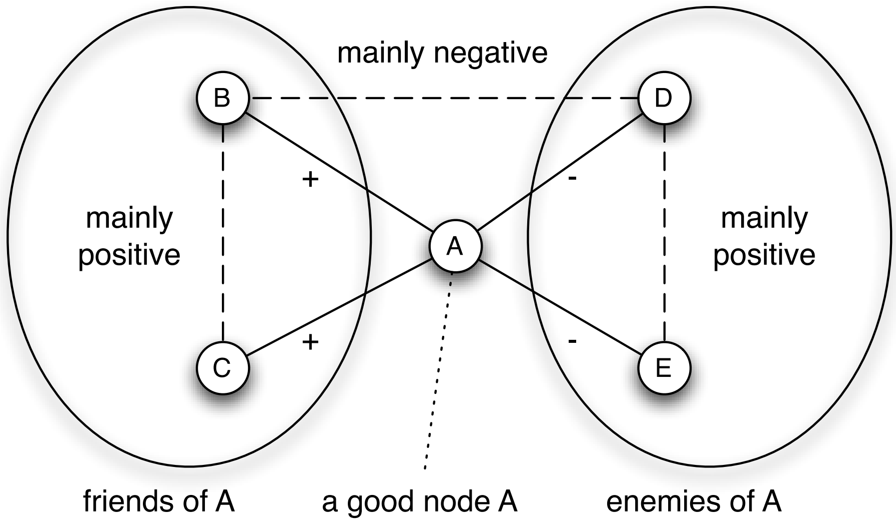

Figure 5.17: The characterization of approximately balanced complete graphs follows from an analysis similar to the proof of the original Balance Theorem. However, we have to be more careful in dividing the graph by first finding a "good" node that isn't involved in too many unbalanced triangles.

- And finally, an analogous argument applies to edges with one end in X and the other end in Y. Each such edge that is positive creates a distinct unbalanced triangle involving A, and so there are at most $\varepsilon N^2/2$ positive edges with one end in X and the other end in Y.

We now consider several possible cases, depending on the sizes of the sets X and Y. Essentially, if either of X or Y consists of almost the entire graph, then we show that alternative (a) in the claim holds. Otherwise, if each of X and Y contain a non-negligible number of nodes, then we show that alternative (b) in the claim holds. We're also going to assume, to make the calculations simpler, that N is even and that the quantity δN is a whole number, although this is not in fact necessary for the proof.

To start, let x be the number of nodes in X and y be the number of nodes in Y. Suppose first that $x \geq (1-\delta)N$. Since $\varepsilon < \frac{1}{8}$ and $\delta = \sqrt[3]{\varepsilon}$, it follows that $\delta < \frac{1}{2}$, and so $x > \frac{1}{2}N$. Now, recall our earlier counting argument that gave a formula for the number of edges in a complete graph, in terms of its number of nodes. In this case, X has x nodes, so it has $x(x-1)/2$ edges. Since $x > \frac{1}{2}N$, this number of edges is at least $(\frac{1}{2}N+1)(\frac{1}{2}N)/2 \geq (\frac{1}{2}N)^2/2 = N^2/8$. There are at most $\varepsilon N^2/2$ negative edges inside X, and so the fraction of negative edges inside X is at most

$$\frac{\varepsilon N^2/2}{N^2/8} = 4\varepsilon = 4\delta^3 < \delta,$$

where we use the facts that $\varepsilon = \delta^3$ and $\delta < \frac{1}{2}$. We thus conclude that if X contains at least $(1-\delta)N$ nodes, then it is a set containing at least a $1-\delta$ fraction of the nodes in which at least $1-\delta$ of all pairs are friends, satisfying part (a) in the conclusion of the claim.

The same argument can be applied if Y contains at least $(1-\delta)N$ nodes. Thus we are left with the case in which both X and Y contain strictly fewer than $(1-\delta)N$, and in this case we will show that part (b) in the conclusion of the claim holds. First, of all the edges with one end in X and the other in Y, what fraction are positive? The total number of edges with one end in X and the other end in Y can be counted as follows: there are x ways to choose the end in X, and then y ways to choose the end in Y, for a total of xy such edges. Now, since each of x and y are less than $(1-\delta)N$, and they add up to N, this product xy is at least $(\delta N)(1-\delta)N = \delta(1-\delta)N^2 \geq \delta N^2/2$, where the last inequality follows from the fact that $\delta < \frac{1}{2}$. There are at most $\varepsilon N^2/2$ positive edges with one end in X and the other in Y, so as a fraction of the total this is at most

$$\frac{\varepsilon N^2/2}{\delta N^2/2} = \frac{\varepsilon}{\delta} = \delta^2 < \delta.$$

Finally, what fraction of edges inside each of X and Y are negative? Let's calculate this for X; the argument for Y is exactly the same. There are $x(x-1)/2$ edges inside X in total, and since we're in the case where $x > \delta N$, this total number of edges is at least $(\delta N + 1)(\delta N)/2 \geq (\delta N)^2/2 = \delta^2 N^2/2$. There are at most $\varepsilon N^2/2$ negative edges inside X, so as a fraction of the total this is at most

$$\frac{\varepsilon N^2/2}{\delta^2 N^2/2} = \frac{\varepsilon}{\delta^2} = \delta.$$

Thus, the division of nodes into sets X and Y satisfies all the requirements in conclusion (b) of the claim, and so the proof is complete.

As a final comment on the claim and its proof, one might feel that the difference between $1-\varepsilon$ in the assumption of the claim and $1-\sqrt[3]{\varepsilon}$ is a bit excessive: as we saw above, when $\varepsilon = .001$, this means we need to assume that 99.9% of all triangles are balanced in order to get sets with a 90% density of edges having the correct sign. But in fact, it is possible to construct examples showing that this relationship between ε and δ is in fact essentially the best one can do. In short, the claim provides the kind of approximate version of the Balance Theorem that we wanted at a qualitative level, but we need to assume a fairly small fraction of unbalanced triangles in order to be able to start drawing strong conclusions.

5.6 Exercises

1. Suppose that a team of anthropologists is studying a set of three small villages that neighbor one another. Each village has 30 people, consisting of 2-3 extended families.

Everyone in each village knows all the people in their own village, as well as the people in the other villages.

When the anthropologists build the social network on the people in all three villages taken together, they find that each person is friends with all the other people in their own village, and enemies with everyone in the two other villages. This gives them a network on 90 people (i.e., 30 in each village), with positive and negative signs on its edges.

According to the definitions in this chapter, is this network on 90 people balanced? Give a brief explanation for your answer.

2. Consider the network shown in Figure 5.18: there is an edge between each pair of nodes, with five of the edges corresponding to positive relationships, and the other five of the edges corresponding to negative relationships.

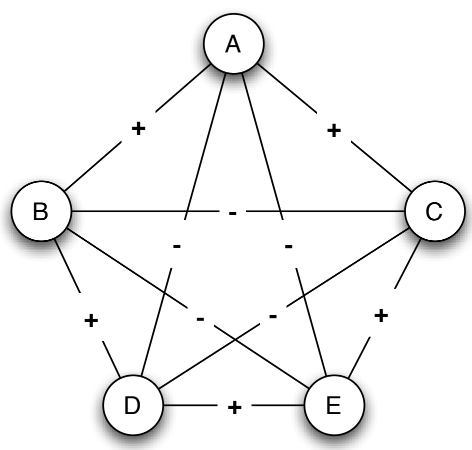

Figure 5.18: A network with five positive edges and five negative edges.

Each edge in this network participates in three triangles: one formed by each of the additional nodes who is not already an endpoint of the edge. (For example, the A-B edge participates in a triangle on A, B, and C, a triangle on A, B, and D, and a triangle on A, B, and E. We can list triangles for the other edges in a similar way.)

For each edge, how many of the triangles it participates in are balanced, and how many are unbalanced. (Notice that because of the symmetry of the network, the answer will be the same for each positive edge, and also for each negative edge; so it is enough to consider this for one of the positive edges and one of the negative edges.)

3. When we think about structural balance, we can ask what happens when a new node

tries to join a network in which there is existing friendship and hostility. In Figures 5.19–5.22, each pair of nodes is either friendly or hostile, as indicated by the + or − label on each edge.

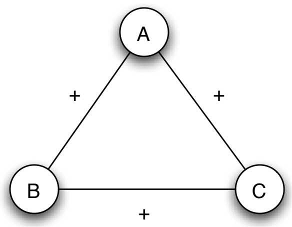

Figure 5.19: A 3-node social network in which all pairs of nodes know each other, and all pairs of nodes are friendly toward each other.

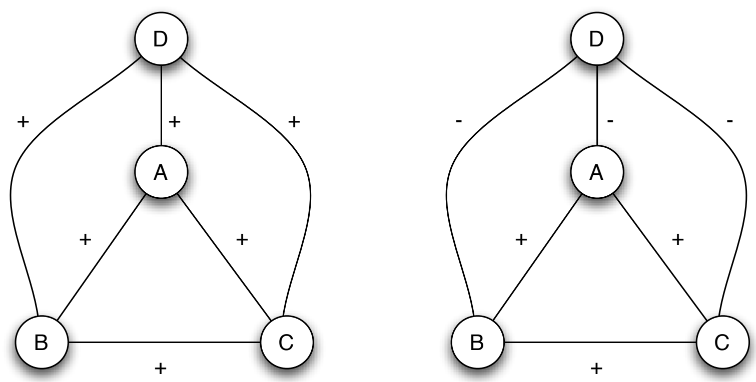

(a) *D joins the network by becoming friends with all nodes.*

(b) *D joins the network by becoming enemies with all nodes.*

Figure 5.20: There are two distinct ways in which node D can join the social network from Figure 5.19 without becoming involved in any unbalanced triangles.

First, consider the 3-node social network in Figure 5.19, in which all pairs of nodes know each other, and all pairs of nodes are friendly toward each other. Now, a fourth node D wants to join this network, and establish either positive or negative relations with each existing node A, B, and C. It wants to do this in such a way that it doesn't become involved in any unbalanced triangles. (I.e. so that after adding D and the labeled edges from D, there are no unbalanced triangles that contain D.) Is this possible?

In fact, in this example, there are two ways for D to accomplish this, as indicated in Figure 5.20. First, D can become friends with all existing nodes; in this way, all the triangles containing it have three positive edges, and so are balanced. Alternately, it can become enemies with all existing nodes; in this way, each triangle containing it has exactly one positive edge, and again these triangles would be balanced.

So for this network, it was possible for D to join without becoming involved in any unbalanced triangles. However, the same is not necessarily possible for other networks.

We now consider this kind of question for some other networks.

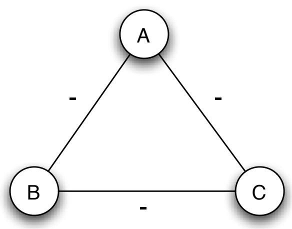

Figure 5.21: All three nodes are mutual enemies.

(a) Consider the 3-node social network in Figure 5.21, in which all pairs of nodes know each other, and each pair is either friendly or hostile as indicated by the $+$ or $-$ label on each edge. A fourth node D wants to join this network, and establish either positive or negative relations with each existing node A, B, and C. Can node D do this in such a way that it doesn't become involved in any unbalanced triangles?

- If there is a way for D to do this, say how many different such ways there are, and give an explanation. (That is, how many different possible labelings of the edges out of D have the property that all triangles containing D are balanced?)
- If there is no such way for D to do this, give an explanation why not.

(In this and the subsequent questions, it possible to work out an answer by reasoning about the new node's options without having to check all possibilities.)

(b) Same question, but for a different network. Consider the 3-node social network in Figure 5.22, in which all pairs of nodes know each other, and each pair is either friendly or hostile as indicated by the $+$ or $-$ label on each edge. A fourth node D wants to join this network, and establish either positive or negative relations with each existing node A, B, and C. Can node D do this in such a way that it doesn't become involved in any unbalanced triangles?

- If there is a way for D to do this, say how many different such ways there are, and give an explanation. (That is, how many different possible labelings

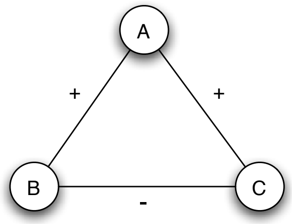

Figure 5.22: Node A is friends with nodes B and C, who are enemies with each other.

of the edges out of D have the property that all triangles containing D are balanced?)

- If there is no such way for D to do this, give an explanation why not.

(c) Using what you've worked out in Questions 2 and 3, consider the following question. Take *any* labeled complete graph — on any number of nodes — that is not balanced; i.e. it contains at least one unbalanced triangle. (Recall that a labeled complete graph is a graph in which there is an edge between each pair of nodes, and each edge is labeled with either $+$ or $-$.) A new node X wants to join this network, by attaching to each node using a positive or negative edge. When, if ever, is it possible for X to do this in such a way that it does not become involved in any unbalanced triangles? Give an explanation for your answer. *(Hint: Think about any unbalanced triangle in the network, and how X must attach to the nodes in it.)*

4. Together with some anthropologists, you're studying a sparsely populated region of a rain forest, where 50 farmers live along a 50-mile-long stretch of river. Each farmer lives on a tract of land that occupies a 1-mile stretch of the river bank, so their tracts exactly divide up the 50 miles of river bank that they collectively cover. (The numbers are chosen to be simple and to make the story easy to describe.)

   The farmers all know each other, and after interviewing them, you've discovered that each farmer is friends with all the other farmers that live at most 20 miles from him or her, and is enemies with all the farmers that live more than 20 miles from him or her.

   You build the signed complete graph corresponding to this social network, and you wonder whether it satisfies the Structural Balance property. This is the question: is the network structurally balanced or not? Provide an explanation for your answer.