

Chapter 3

Strong and Weak Ties

One of the powerful roles that networks play is to bridge the local and the global — to offer explanations for how simple processes at the level of individual nodes and links can have complex effects that ripple through a population as a whole. In this chapter, we consider some fundamental social network issues that illustrate this theme: how information flows through a social network, how different nodes can play structurally distinct roles in this process, and how these structural considerations shape the evolution of the network itself over time. These themes all play central roles throughout the book, adapting themselves to different contexts as they arise. Our context in this chapter will begin with the famous “strength of weak ties” hypothesis from sociology [190], exploring outward from this point to more general settings as well.

Let’s begin with some background and a motivating question. As part of his Ph.D. thesis research in the late 1960s, Mark Granovetter interviewed people who had recently changed employers to learn how they discovered their new jobs [191]. In keeping with earlier research, he found that many people learned information leading to their current jobs through personal contacts. But perhaps more strikingly, these personal contacts were often described by interview subjects as acquaintances rather than close friends. This is a bit surprising: your close friends presumably have the most motivation to help you when you’re between jobs, so why is it so often your more distant acquaintances who are actually to thank for crucial information leading to your new job?

The answer that Granovetter proposed to this question is striking in the way it links two different perspectives on distant friendships — one structural, focusing on the way these friendships span different portions of the full network; and the other interpersonal, considering the purely local consequences that follow from a friendship between two people being either strong or weak. In this way, the answer transcends the specific setting of job-

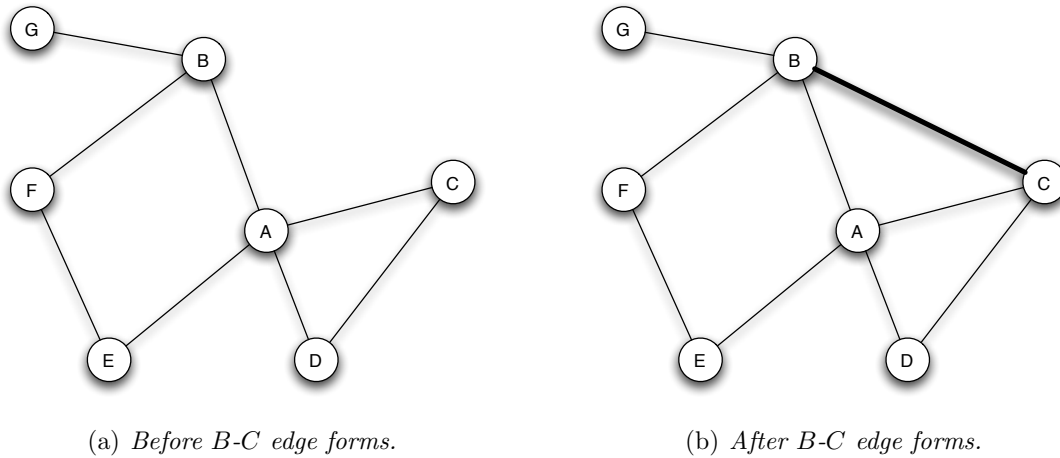


Figure 3.1: The formation of the edge between B and C illustrates the effects of triadic closure, since they have a common neighbor A .

seeking, and offers a way of thinking about the architecture of social networks more generally. To get at this broader view, we first develop some general principles about social networks and their evolution, and then return to Granovetter’s question.

3.1 Triadic Closure

In Chapter 2, our discussions of networks treated them largely as static structures — we take a snapshot of the nodes and edges at a particular moment in time, and then ask about paths, components, distances, and so forth. While this style of analysis forms the basic foundation for thinking about networks — and indeed, many datasets are inherently static, offering us only a single snapshot of a network — it is also useful to think about how a network evolves over time. In particular, what are the mechanisms by which nodes arrive and depart, and by which edges form and vanish?

The precise answer will of course vary depending on the type of network we’re considering, but one of the most basic principles is the following:

If two people in a social network have a friend in common, then there is an increased likelihood that they will become friends themselves at some point in the future [347].

We refer to this principle as *triadic closure*, and it is illustrated in Figure 3.1: if nodes B and C have a friend A in common, then the formation of an edge between B and C produces a situation in which all three nodes A , B , and C have edges connecting each other — a structure we refer to as a *triangle* in the network. The term “triadic closure” comes from




Figure 3.2: If we watch a network for a longer span of time, we can see multiple edges forming — some form through triadic closure while others (such as the $D-G$ edge) form even though the two endpoints have no neighbors in common.

the fact that the $B-C$ edge has the effect of “closing” the third side of this triangle. If we observe snapshots of a social network at two distinct points in time, then in the later snapshot, we generally find a significant number of new edges that have formed through this triangle-closing operation, between two people who had a common neighbor in the earlier snapshot. Figure 3.2, for example, shows the new edges we might see from watching the network in Figure 3.1 over a longer time span.

The Clustering Coefficient. The basic role of triadic closure in social networks has motivated the formulation of simple social network measures to capture its prevalence. One of these is the *clustering coefficient* [320, 411]. The clustering coefficient of a node A is defined as the probability that two randomly selected friends of A are friends with each other. In other words, it is the fraction of pairs of A ’s friends that are connected to each other by edges. For example, the clustering coefficient of node A in Figure 3.2(a) is $1/6$ (because there is only the single $C-D$ edge among the six pairs of friends $B-C$, $B-D$, $B-E$, $C-D$, $C-E$, and $D-E$), and it has increased to $1/2$ in the second snapshot of the network in Figure 3.2(b) (because there are now the three edges $B-C$, $C-D$, and $D-E$ among the same six pairs). In general, the clustering coefficient of a node ranges from 0 (when none of the node’s friends are friends with each other) to 1 (when all of the node’s friends are friends with each other), and the more strongly triadic closure is operating in the neighborhood of the node, the higher the clustering coefficient will tend to be.




Figure 3.3: The A - B edge is a *bridge*, meaning that its removal would place A and B in distinct connected components. Bridges provide nodes with access to parts of the network that are unreachable by other means.

Reasons for Triadic Closure. Triadic closure is intuitively very natural, and essentially everyone can find examples from their own experience. Moreover, experience suggests some of the basic reasons why it operates. One reason why B and C are more likely to become friends, when they have a common friend A , is simply based on the *opportunity* for B and C to meet: if A spends time with both B and C , then there is an increased chance that they will end up knowing each other and potentially becoming friends. A second, related reason is that in the process of forming a friendship, the fact that each of B and C is friends with A (provided they are mutually aware of this) gives them a basis for *trusting* each other that an arbitrary pair of unconnected people might lack.

A third reason is based on the *incentive* A may have to bring B and C together: if A is friends with B and C , then it becomes a source of latent stress in these relationships if B and C are not friends with each other. This premise is based in theories dating back to early work in social psychology [217]; it also has empirical reflections that show up in natural but troubling ways in public-health data. For example, Bearman and Moody have found that teenage girls who have a low clustering coefficient in their network of friends are significantly more likely to contemplate suicide than those whose clustering coefficient is high [48].

3.2 The Strength of Weak Ties

So how does all this relate to Mark Granovetter’s interview subjects, telling him with such regularity that their best job leads came from acquaintances rather than close friends? In fact, triadic closure turns out to be one of the crucial ideas needed to unravel what’s going on.

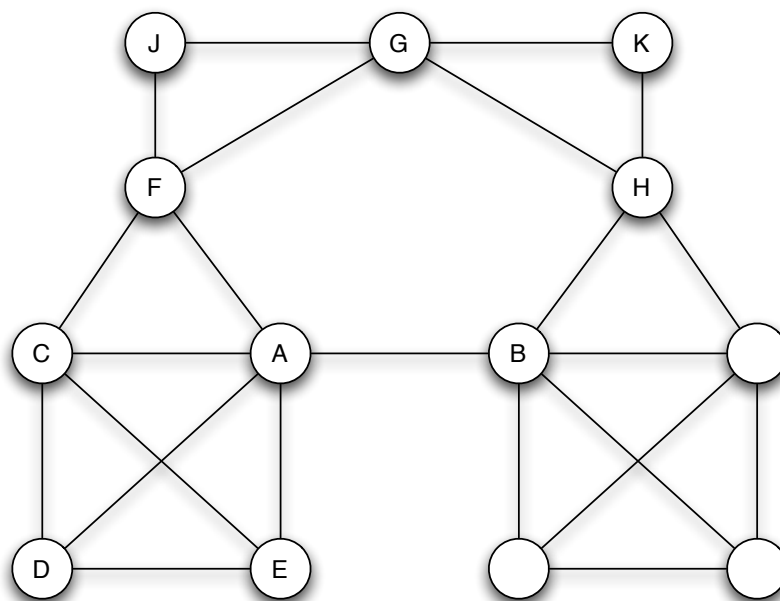


Figure 3.4: The A - B edge is a local bridge of span 4, since the removal of this edge would increase the distance between A and B to 4.

Bridges and Local Bridges. Let's start by positing that information about good jobs is something that is relatively scarce; hearing about a promising job opportunity from someone suggests that they have access to a source of useful information that you don't. Now consider this observation in the context of the simple social network drawn in Figure 3.3. The person labeled A has four friends in this picture, but one of her friendships is qualitatively different from the others: A 's links to C , D , and E connect her to a tightly-knit group of friends who all know each other, while the link to B seems to reach into a different part of the network. We could speculate, then, that the structural peculiarity of the link to B will translate into differences in the role it plays in A 's everyday life: while the tightly-knit group of nodes A , C , D , and E will all tend to be exposed to similar opinions and similar sources of information, A 's link to B offers her access to things she otherwise wouldn't necessarily hear about.

To make precise the sense in which the A - B link is unusual, we introduce the following definition. We say that an edge joining two nodes A and B in a graph is a *bridge* if deleting the edge would cause A and B to lie in two different components. In other words, this edge is literally the only route between its endpoints, the nodes A and B .

Now, if our discussion in Chapter 2 about giant components and small-world properties taught us anything, it's that bridges are presumably extremely rare in real social networks. You may have a friend from a very different background, and it may seem that your friendship is the only thing that bridges your world and his, but one expects in reality that there will

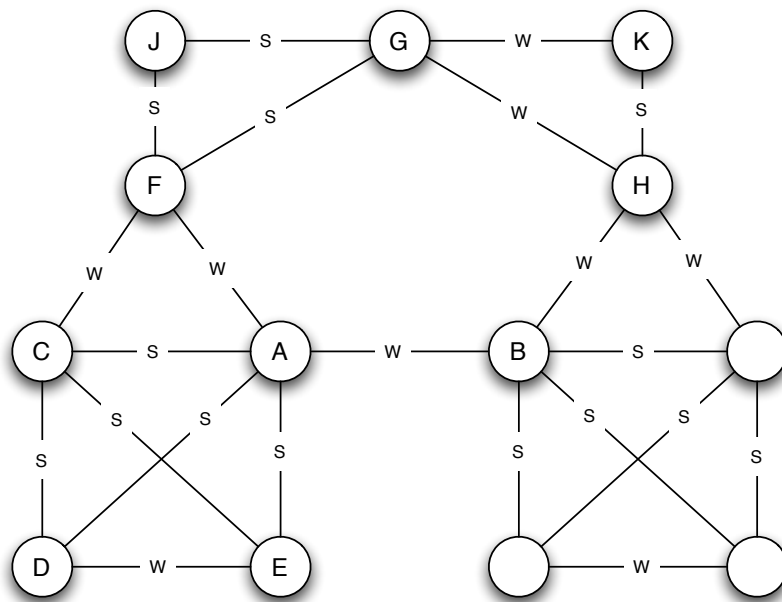


Figure 3.5: Each edge of the social network from Figure 3.4 is labeled here as either a *strong tie* (S) or a *weak tie* (W), to indicate the strength of the relationship. The labeling in the figure satisfies the Strong Triadic Closure Property at each node: if the node has strong ties to two neighbors, then these neighbors must have at least a weak tie between them.

be other, hard-to-discover, multi-step paths that also span these worlds. In other words, if we were to look at Figure 3.3 as it is embedded in a larger, ambient social network, we would likely see a picture that looks like Figure 3.4.

Here, the A - B edge isn't the only path that connects its two endpoints; though they may not realize it, A and B are also connected by a longer path through F , G , and H . This kind of structure is arguably much more common than a bridge in real social networks, and we use the following definition to capture it. We say that an edge joining two nodes A and B in a graph is a *local bridge* if its endpoints A and B have no friends in common — in other words, if deleting the edge would increase the distance between A and B to a value strictly more than two. We say that the *span* of a local bridge is the distance its endpoints would be from each other if the edge were deleted [190, 407]. Thus, in Figure 3.4, the A - B edge is a local bridge with span four; we can also check that no other edge in this graph is a local bridge, since for every other edge in the graph, the endpoints would still be at distance two if the edge were deleted. Notice that the definition of a local bridge already makes an implicit connection with triadic closure, in that the two notions form conceptual opposites: an edge is a local bridge precisely when it does not form a side of any triangle in the graph.

Local bridges, especially those with reasonably large span, still play roughly the same

role that bridges do, though in a less extreme way — they provide their endpoints with access to parts of the network, and hence sources of information, that they would otherwise be far away from. And so this is a first network context in which to interpret Granovetter’s observation about job-seeking: we might expect that if a node like *A* is going to get truly new information, the kind that leads to a new job, it might come unusually often (though certainly not always) from a friend connected by a local bridge. The closely-knit groups that you belong to, though they are filled with people eager to help, are also filled with people who know roughly the same things that you do.

The Strong Triadic Closure Property. Of course, Granovetter’s interview subjects didn’t say, “I learned about the job from a friend connected by a local bridge.” If we believe that local bridges were overrepresented in the set of people providing job leads, how does this relate to the observation that distant acquaintances were overrepresented as well?

To talk about this in any detail, we need to be able to distinguish between different levels of *strength* in the links of a social network. We deliberately refrain from trying to define “strength” precisely, but we mean it to align with the idea that stronger links represent closer friendship and greater frequency of interaction. In general, links can have a wide range of possible strengths, but for conceptual simplicity — and to match the friend/acquaintance dichotomy that we’re trying to explain — we’ll categorize all links in the social network as belonging to one of two types: *strong ties* (the stronger links, corresponding to friends), and *weak ties* (the weaker links, corresponding to acquaintances).¹

Once we have decided on a classification of links into strong and weak ties, we can take a social network and annotate each edge with a designation of it as either strong or weak. For example, assuming we asked the nodes in the social network of Figure 3.4 to report which of their network neighbors were close friends and which were acquaintances, we could get an annotated network as in Figure 3.5.

It is useful to go back and think about triadic closure in terms of this division of edges into strong and weak ties. If we recall the arguments supporting triadic closure, based on opportunity, trust, and incentive, they all act more powerfully when the edges involved are

¹In addition to the difficulty in reducing a range of possible link strengths to a two-category strong/weak distinction, there are many other subtleties in this type of classification. For example, in the discussion here, we will take this division of links into strong and weak ties as fixed in a single snapshot of the network. In reality, of course, the strength of a particular link can vary across different times and different situations. For example, an employee of a company who is temporarily assigned to work with a new division of the company for a few months may find that her full set of available social-network links remains roughly the same, but that her links to people within the new division have been temporarily strengthened (due to the sudden close proximity and increased contact), while her links to her old division have been temporarily weakened. Similarly, a high-school student may find that links to fellow members of a particular sports team constitute strong ties while that sport is in season, but that some of these links — to the teammates he knows less well outside of the team — become weak ties in other parts of the year. Again, for our purposes, we will consider a single distinction between strong and weak ties that holds throughout the analysis.

strong ties than when they are weak ties. This suggests the following qualitative assumption:

If a node A has edges to nodes B and C , then the B - C edge is especially likely to form if A 's edges to B and C are both strong ties.

To enable some more concrete analysis, Granovetter suggested a more formal (and somewhat more extreme version) of this, as follows.

We say that a node A violates the Strong Triadic Closure Property if it has strong ties to two other nodes B and C , and there is no edge at all (either a strong or weak tie) between B and C . We say that a node A satisfies the Strong Triadic Closure Property if it does not violate it.

You can check that no node in Figure 3.5 violates the Strong Triadic Closure Property, and hence all nodes satisfy the Property. On the other hand, if the A - F edge were to be a strong tie rather than a weak tie, then nodes A and F would both violate the Strong Triadic Closure Property: Node A would now have strong ties to nodes E and F without there being an E - F edge, and node F would have strong ties to both A and G without there being an A - G edge. As a further check on the definition, notice that with the labeling of edges as in Figure 3.5, node H satisfies the Strong Triadic Closure Property: H couldn't possibly violate the Property since it only has a strong tie to one other node.

Clearly the Strong Triadic Closure Property is too extreme for us to expect it hold across all nodes of a large social network. But it is a useful step as an abstraction to reality, making it possible to reason further about the structural consequences of strong and weak ties. In the same way that an introductory physics course might assume away the effects of air resistance in analyzing the flight of a ball, proposing a slightly too-powerful assumption in a network context can also lead to cleaner and conceptually more informative analysis. For now, then, let's continue figuring out where it leads us in this case; later, we'll return to the question of its role as a modeling assumption.

Local Bridges and Weak Ties. We now have a purely local, interpersonal distinction between kinds of links — whether they are weak ties or strong ties — as well as a global, structural notion — whether they are local bridges or not. On the surface, there is no direct connection between the two notions, but in fact using triadic closure we can establish a connection, in the following claim.

Claim: If a node A in a network satisfies the Strong Triadic Closure Property and is involved in at least two strong ties, then any local bridge it is involved in must be a weak tie.

In other words, assuming the Strong Triadic Closure Property and a sufficient number of strong ties, the local bridges in a network are necessarily weak ties.

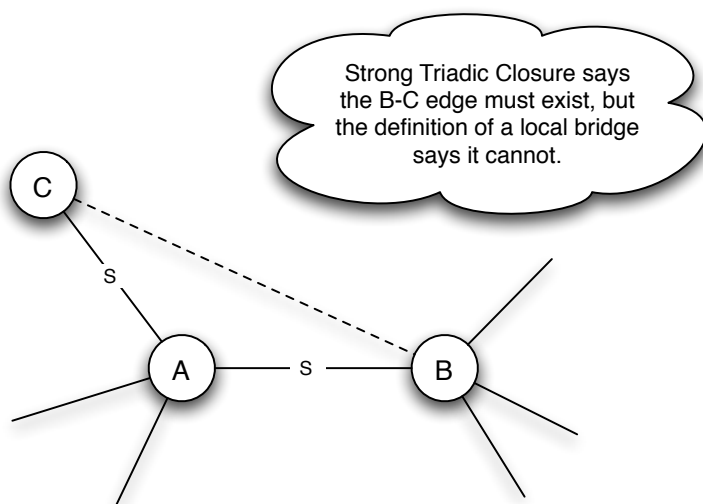


Figure 3.6: If a node satisfies Strong Triadic Closure and is involved in at least two strong ties, then any local bridge it is involved in must be a weak tie. The figure illustrates the reason why: if the A - B edge is a strong tie, then there must also be an edge between B and C , meaning that the A - B edge cannot be a local bridge.

We’re going to justify this claim as a mathematical statement – that is, it will follow logically from the definitions we have so far, without our having to invoke any as-yet-unformalized intuitions about what social networks ought to look like. In this way, it’s a different kind of claim from our argument in Chapter 2 that the global friendship network likely contains a giant component. That was a thought experiment (albeit a very convincing one), requiring us to believe various empirical statements about the network of human friendships — empirical statements that could later be confirmed or refuted by collecting data on large social networks. Here, on the other hand, we’ve constructed a small number of specific mathematical definitions — particularly, local bridges and the Strong Triadic Closure Property — and we can now justify the claim directly from these.

The argument is actually very short, and it proceeds by contradiction. Take some network, and consider a node A that satisfies the Strong Triadic Closure Property and is involved in at least two strong ties. Now suppose A is involved in a local bridge — say, to a node B — that is a strong tie. We want to argue that this is impossible, and the crux of the argument is depicted in Figure 3.6. First, since A is involved in at least two strong ties, and the edge to B is only one of them, it must have a strong tie to some other node, which we’ll call C . Now let’s ask: is there an edge connecting B and C ? Since the edge from A to B is a local bridge, A and B must have no friends in common, and so the B - C edge must not exist. But this contradicts Strong Triadic Closure, which says that since the A - B and

A - C edges are both strong ties, the B - C edge must exist. This contradiction shows that our initial premise, the existence of a local bridge that is a strong tie, cannot hold, finishing the argument.

This argument completes the connection we've been looking for between the local property of tie strength and the global property of serving as a local bridge. As such, it gives us a way to think about the way in which interpersonal properties of social-network links are related to broader considerations about the network's structure. But since the argument is based on some strong assumptions (mainly Strong Triadic Closure, since the other assumption is very mild), it is also worth reflecting on the role that simplifying assumptions play in a result like this.

First, simplifying assumptions are useful when they lead to statements that are robust in practice, making sense as qualitative conclusions that hold in approximate forms even when the assumptions are relaxed. This is the case here: the mathematical argument can be summarized more informally and approximately as saying that in real life, a local bridge between nodes A and B tends to be a weak tie because if it weren't, triadic closure would tend to produce short-cuts to A and B that would eliminate its role as a local bridge. Again, one is tempted to invoke the analogy to freshman physics: even if the assumptions used to derive the perfectly parabolic flight of a ball don't hold exactly in the real world, the conclusions about flight trajectories are a very useful, conceptually tractable approximation to reality.

Second, when the underlying assumptions are stated precisely, as they are here, it becomes possible to test them on real-world data. In the past few years researchers have studied the relationship of tie strength and network structure quantitatively across large populations, and have shown that the conclusions described here in fact hold in an approximate form. We describe some of this empirical research in the next section.

Finally, this analysis provides a concrete framework for thinking about the initially surprising fact that life transitions such as a new jobs are often rooted in contact with distant acquaintances. The argument is that these are the social ties that connect us to new sources of information and new opportunities, and their conceptual "span" in the social network (the local bridge property) is directly related to their weakness as social ties. This dual role as weak connections but also valuable conduits to hard-to-reach parts of the network — this is the surprising strength of weak ties.

3.3 Tie Strength and Network Structure in Large-Scale Data

The arguments connecting tie strength with structural properties of the underlying social network make intriguing theoretical predictions about the organization of social networks

in real life. For many years after Granovetter’s initial work, however, these predictions remained relatively untested on large social networks, due to the difficulty in finding data that reliably captured the strengths of edges in large-scale, realistic settings.

This state of affairs began to change rapidly once detailed traces of digital communication became available. Such “who-talks-to-whom” data exhibits the two ingredients we need for empirical evaluation of hypotheses about weak ties: it contains the network structure of communication among pairs of people, and we can use the total time that two people spend talking to each other as a proxy for the strength of the tie — the more time spent communicating during the course of an observation period, the stronger we declare the tie to be.

In one of the more comprehensive studies of this type, Onnela et al. studied the who-talks-to-whom network maintained by a cell-phone provider that covered roughly 20% of a national population [334]. The nodes correspond to cell-phone users, and there is an edge joining two nodes if they made phone calls to each other in both directions over an 18-week observation period. Because the cell phones in this population are generally used for personal communication rather than business purposes, and because the lack of a central directory means that cell-phone numbers are generally exchanged among people who already know each other, the underlying network can be viewed as a reasonable sampling of the conversations occurring within a social network representing a significant fraction of one country’s population. Moreover, the data exhibits many of the broad structural features of large social networks discussed in Chapter 2, including a *giant component* — a single connected component containing most (in this case 84%) of the individuals in the network.

Generalizing the Notions of Weak Ties and Local Bridges. The theoretical formulation in the preceding section is based on two definitions that impose sharp dichotomies on the network: an edge is either a strong tie or a weak tie, and it is either a local bridge or it isn’t. For both of these definitions, it is useful to have versions that exhibit smoother gradations when we go to examine real data at a large scale.

Above, we just indicated a way to do this for tie strength: we can make the strength of an edge a numerical quantity, defining it to be the total number of minutes spent on phone calls between the two ends of the edge. It is also useful to sort all the edges by tie strength, so that for a given edge we can ask what percentile it occupies this ordering of edges sorted by strength.

Since a very small fraction of the edges in the cell-phone data are local bridges, it makes sense to soften this definition as well, so that we can view certain edges as being “almost” local bridges. To do this, we define the *neighborhood overlap* of an edge connecting A and B to be the ratio

$$\frac{\text{number of nodes who are neighbors of both } A \text{ and } B}{\text{number of nodes who are neighbors of at least one of } A \text{ or } B}, \quad (3.1)$$




Figure 3.7: A plot of the neighborhood overlap of edges as a function of their percentile in the sorted order of all edges by tie strength. The fact that overlap increases with increasing tie strength is consistent with the theoretical predictions from Section 3.2. (Image from [334].)

where in the denominator we don't count A or B themselves (even though A is a neighbor of B and B is a neighbor of A). As an example of how this definition works, consider the edge $A-F$ in Figure 3.4. The denominator of the neighborhood overlap for $A-F$ is determined by the nodes $B, C, D, E, G,$ and J , since these are the ones that are a neighbor of at least one of A or F . Of these, only C is a neighbor of both A and F , so the neighborhood overlap is $1/6$.

The key feature of this definition is that this ratio in question is 0 precisely when the numerator is 0, and hence when the edge is a local bridge. So the notion of a local bridge is contained within this definition — local bridges are the edges of neighborhood overlap 0 — and hence we can think of edges with very small neighborhood overlap as being “almost” local bridges. (Since intuitively, edges with very small neighborhood overlap consist of nodes that travel in “social circles” having almost no one in common.) For example, this definition views the $A-F$ edge as much closer to being a local bridge than the $A-E$ edge is, which accords with intuition.

Empirical Results on Tie Strength and Neighborhood Overlap. Using these definitions, we can formulate some fundamental quantitative questions based on Granovetter’s theoretical predictions. First, we can ask how the neighborhood overlap of an edge depends on its strength; the strength of weak ties predicts that neighborhood overlap should grow as tie strength grows.

In fact, this is borne out extremely cleanly by the data. Figure 3.7 shows the neighborhood overlap of edges as a function of their percentile in the sorted order of all edges by tie strength. Thus, as we go to the right on the x -axis, we get edges of greater and greater strength, and because the curve rises in a strikingly linear fashion, we also get edges of greater and greater neighborhood overlap. The relationship between these quantities thus aligns well with the theoretical prediction.²

The measurements underlying Figure 3.7 describe a connection between tie strength and network structure at a local level — in the neighborhoods of individual nodes. It is also interesting to consider how this type of data can be used to evaluate the more global picture suggested by the theoretical framework, that weak ties serve to link together different tightly-knit communities that each contain a large number of stronger ties. Here, Onnela et al. provided an indirect analysis to address this question, as follows. They first deleted edges from the network one at a time, starting with the strongest ties and working downward in order of tie strength. The giant component shrank steadily as they did this, its size going down gradually due to the elimination of connections among the nodes. They then tried the same thing, but starting from the weakest ties and working upward in order of tie strength. In this case, they found that the giant component shrank more rapidly, and moreover that its remnants broke apart abruptly once a critical number of weak ties had been removed. This is consistent with a picture in which the weak ties provide the more crucial connective structure for holding together disparate communities, and for keeping the global structure of the giant component intact.

Ultimately, this is just a first step toward evaluating theories of tie strength on network data of this scale, and it illustrates some of the inherent challenges: given the size and complexity of the network, we cannot simply look at the structure and “see what’s there.” Indirect measures must generally be used, and since one knows relatively little about the meaning or significance of any particular node or edge, it remains an ongoing research challenge to draw richer and more detailed conclusions in the way that one can on small datasets.

²It is of course interesting to note the deviation from this trend at the very right-hand edge of the plot in Figure 3.7, corresponding to the edges of greatest possible tie strength. It is not clear what causes this deviation, but it is certainly plausible that these extremely strong edges are associated with people who are using their cell-phones in some unusual fashion.

3.4 Tie Strength, Social Media, and Passive Engagement

As an increasing amount of social interaction moves on-line, the way in which we maintain and access our social networks begins to change as well. For example, as is well-known to users of social-networking tools, people maintain large explicit lists of friends in their profiles on these sites — in contrast to the ways in which such friendship circles were once much more implicit, and in fact relatively difficult for individuals even to enumerate or mentally access [244]. What effect does this have on social network structure more broadly? Understanding the changes arising from these forms of technological mediation is a challenge that was already being articulated in the early 1990s by researchers including Barry Wellman [414, 413], as the Internet began making remote interaction possible for a broad public; these issues have of course grown steadily more pervasive between then and now.

Tie strength can provide an important perspective on such questions, providing a language for asking how on-line social activity is distributed across different kinds of links — and in particular, how it is distributed across links of different strengths. When we see people maintaining hundreds of friendship links on a social-networking site, we can ask how many of these correspond to strong ties that involve frequent contact, and how many of these correspond to weak ties that are activated relatively rarely.

Tie Strength on Facebook. Researchers have begun to address such questions of tie strength using data from some of the most active social media sites. At Facebook, Cameron Marlow and his colleagues analyzed the friendship links reported in each user’s profile, asking to what extent each link was actually *used* for social interaction, beyond simply being reported in the profile [286]. In other words, where are the strong ties among a user’s friends? To make this precise using the data they had available, they defined three categories of links based on usage over a one-month observation period.

- A link represents *reciprocal (mutual) communication*, if the user both sent messages to the friend at the other end of the link, and also received messages from them during the observation period.
- A link represents *one-way communication* if the user sent one or more messages to the friend at the other end of the link (whether or not these messages were reciprocated).
- A link represents a *maintained relationship* if the user followed information about the friend at the other end of the link, whether or not actual communication took place; “following information” here means either clicking on content via Facebook’s News Feed service (providing information about the friend) or visiting the friend’s profile more than once.

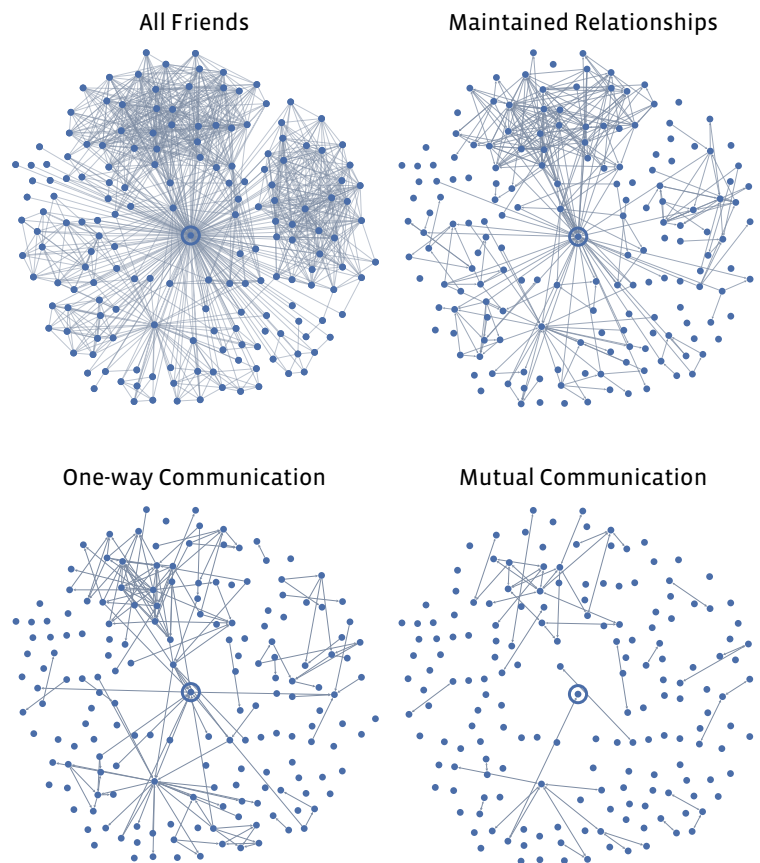


Figure 3.8: Four different views of a Facebook user’s network neighborhood, showing the structure of links corresponding respectively to all declared friendships, maintained relationships, one-way communication, and reciprocal (i.e. mutual) communication. (Image from [286].)

Notice that these three categories are not mutually exclusive — indeed, the links classified as reciprocal communication always belong to the set of links classified as one-way communication.

This stratification of links by their use lets us understand how a large set of declared friendships on a site like Facebook translates into an actual pattern of more active social interaction, corresponding approximately to the use of stronger ties. To get a sense of the relative volumes of these different kinds of interaction through an example, Figure 3.8 shows the network neighborhood of a sample Facebook user — consisting of all his friends, and all links among his friends. The picture in the upper-left shows the set of all declared friendships in this user’s profile; the other three pictures show how the set of links becomes sparser once we consider only maintained relationships, one-way communication, or reciprocal communi-




Figure 3.9: The number of links corresponding to maintained relationships, one-way communication, and reciprocal communication as a function of the total neighborhood size for users on Facebook. (Image from [286].)

cation. Moreover, as we restrict to stronger ties, certain parts of the network neighborhood thin out much faster than others. For example, in the neighborhood of the sample user in Figure 3.8, we see two distinct regions where there has been a particularly large amount of triadic closure: one in the upper part of the drawing, and one on the right-hand side of the drawing. However, when we restrict to links representing communication or a maintained relationship, we see that a lot of the links in the upper region survive, while many fewer of the links in the right-hand region do. One could conjecture that the right-hand region represents a set of friends from some earlier phase of the user’s life (perhaps from high school) who declare each other as friends, but do not actively remain in contact; the upper region, on the other hand, consists of more recent friends (perhaps co-workers) for whom there is more frequent contact.

We can make the relative abundance of these different types of links quantitative through the plot in Figure 3.9. On the x -axis is the total number of friends a user declares, and the curves then show the (smaller) numbers of other link types as a function of this total. There are several interesting conclusions to be drawn from this. First, it confirms that even for users who report very large numbers of friends on their profile pages (on the order of 500),




Figure 3.10: The total number of a user’s strong ties (defined by multiple directed messages) as a function of the number of followees he or she has on Twitter. (Image from [222].)

the number with whom they actually communicate is generally between 10 and 20, and the number they follow even passively (e.g. by reading about them) is under 50. But beyond this observation, Marlow and his colleagues draw a further conclusion about the power of media like Facebook to enable this kind of *passive engagement*, in which one keeps up with friends by reading news about them even in the absence of communication. They argue that this passive network occupies an interesting middle ground between the strongest ties maintained by regular communication and the weakest ties from one’s distant past, preserved only in lists on social-networking profile pages. They write, “The stark contrast between reciprocal and passive networks shows the effect of technologies such as News Feed. If these people were required to talk on the phone to each other, we might see something like the reciprocal network, where everyone is connected to a small number of individuals. Moving to an environment where everyone is passively engaged with each other, some event, such as a new baby or engagement can propagate very quickly through this highly connected network.”

Tie Strength on Twitter. Similar lines of investigation have been carried out recently on the social media site Twitter, where individual users engage in a form of *micro-blogging* by posting very short, 140-character public messages known as “tweets.” Twitter also includes social-network features, and these enable one to distinguish between stronger and weaker ties: each user can specify a set of other users whose messages he or she will follow, and each user can also direct messages specifically to another user. (In the latter case, the message

remains public for everyone to read, but it is marked with a notation indicating that it is intended for a particular user.) Thus, the former kind of interaction defines a social network based on more passive, weak ties — it is very easy for a user to follow many people’s messages without ever directly communicating with any of them. The latter kind of interaction — especially when we look at users directing multiple messages to others — corresponds to a stronger kind of direct interaction.

In a style analogous to the work of Marlow et al., Huberman, Romero, and Wu analyzed the relative abundance of these two kinds of links on Twitter [222]. Specifically, for each user they considered the number of users whose messages she followed (her “followees”), and then defined her strong ties to consist of the users to whom she had directed at least two messages over the course of an observation period. Figure 3.10 shows how the number of strong ties varies as a function of the number of followees. As we saw for Facebook, even for users who maintain very large numbers of weak ties on-line, the number of strong ties remains relatively modest, in this case stabilizing at a value below 50 even for users with over 1000 followees.

There is another useful way to think about the contrast between the ease of forming links and the relative scarcity of strong ties in environments like Facebook and Twitter. By definition, each strong tie requires the continuous investment of time and effort to maintain, and so even people who devote a lot of their energy to building strong ties will eventually reach a limit — imposed simply by the hours available in a day — on the number of ties that they can maintain in this way. The formation of weak ties is governed by much milder constraints — they need to be established at their outset but not necessarily maintained continuously — and so it is easier for someone to accumulate them in large numbers. We will encounter this distinction again in Chapter 13, when we consider how social networks differ at a structural level from information networks such as the World Wide Web.

Understanding the effect that on-line media have on the maintenance and use of social networks is a complex problem for which the underlying research is only in its early stages. But some of these preliminary studies already highlight the ways in which networks of strong ties can still be relatively sparse even in on-line settings where weak ties abound, and how the nature of the underlying on-line medium can affect the ways in which different links are used for conveying information.

3.5 Closure, Structural Holes, and Social Capital

Our discussion thus far suggests a general view of social networks in terms of tightly-knit groups and the weak ties that link them. The analysis has focused primarily on the roles that different kinds of edges of a network play in this structure — with a few edges spanning different groups while most are surrounded by dense patterns of connections.




Figure 3.11: The contrast between densely-knit groups and boundary-spanning links is reflected in the different positions of nodes *A* and *B* in the underlying social network.

There is a lot of further insight to be gained by asking about the roles that different *nodes* play in this structure as well. In social networks, access to edges that span different groups is not equally distributed across all nodes: some nodes are positioned at the interface between multiple groups, with access to boundary-spanning edges, while others are positioned in the middle of a single group. What is the effect of this heterogeneity? Following the expositional lead of social-network researchers including Ron Burt [87], we can formulate an answer to this question as a story about the different experiences that nodes have in a network like the one in Figure 3.11 — particularly in the contrast between the experience of a node such as *A*, who sits at the center of a single tightly-knit group, and node *B*, who sits at the interface between several groups.

Embeddedness. Let's start with node *A*. Node *A*'s set of network neighbors has been subject to considerable triadic closure; *A* has a high clustering coefficient. (Recall that the clustering coefficient is the fraction of pairs of neighbors who are themselves neighbors).

To talk about the structure around *A* it is useful to introduce an additional definition. We define the *embeddedness* of an edge in a network to be the number of common neighbors the two endpoints have. Thus, for example, the *A-B* edge has an embeddedness of two, since *A* and *B* have the two common neighbors *E* and *F*. This definition relates to two notions from earlier in the chapter. First, the embeddedness of an edge is equal to the numerator in

the ratio that defines the *neighborhood overlap* in Equation (3.1) from Section 3.3. Second, we observe that local bridges are precisely the edges that have an embeddedness of zero — since they were defined as those edges whose endpoints have no neighbors in common.

In the example shown in Figure 3.11, what stands out about A is the way in which all of his edges have significant embeddedness. A long line of research in sociology has argued that if two individuals are connected by an embedded edge, then this makes it easier for them to trust one another, and to have confidence in the integrity of the transactions (social, economic, or otherwise) that take place between them [117, 118, 193, 194, 395]. Indeed, the presence of mutual friends puts the interactions between two people “on display” in a social sense, even when they are carried out in private; in the event of misbehavior by one of the two parties to the interaction, there is the potential for social sanctions and reputational consequences from their mutual friends. As Granovetter writes, “My mortification at cheating a friend of long standing may be substantial even when undiscovered. It may increase when a friend becomes aware of it. But it may become even more unbearable when our mutual friends uncover the deceit and tell one another” [194].

No similar kind of deterring threat exists for edges with zero embeddedness, since there is no one who knows both people involved in the interaction. In this respect, the interactions that B has with C and D are much riskier than the embedded interactions that A experiences. Moreover, the constraints on B ’s behavior are made complicated by the fact that she is subject to potentially contradictory norms and expectations from the different groups she associates with [116].

Structural holes. Thus far we have been discussing the advantages that accrue to node A in Figure 3.11 from the closure in his network neighborhood, and the embedded edges that result from this. But a related line of research in sociology, catalyzed by influential work of Burt [86], has argued that network positions such as node B ’s, at the ends of multiple local bridges, confer a distinct set of equally fundamental advantages.

The canonical setting for this argument is the social network within an organization or company, consisting of people who are in some ways collaborating on common objectives and in other ways implicitly competing for career advancement. Note that although we may be thinking about settings in which there is a formal organizational hierarchy — encoding who reports to whom — we’re interested in the more informal network of who knows whom, and who talks to whom on a regular basis. Empirical studies of managers in large corporations has correlated an individual’s success within a company to their access to local bridges [86, 87]. At a more abstract level, the central arguments behind these studies are also supported by the network principles we have been discussing, as we now explore further.

Let’s go back to the network in Figure 3.11, imagining the network to represent the interaction and collaboration among managers in a large company. In Burt’s language,

node B , with her multiple local bridges, spans a *structural hole* in the organization — the “empty space” in the network between two sets of nodes that do not otherwise interact closely. (Unlike the term “local bridge,” which has a precise mathematical definition in terms of the underlying graph, we will keep the term “structural hole” somewhat informal in this discussion.) The argument is that B ’s position offers advantages in several dimensions relative to A ’s. The first kind of advantage, following the observations in the previous section, is an informational one: B has early access to information originating in multiple, non-interacting parts of the network. Any one person has a limited amount of energy they can invest in maintaining contacts across the organization, and B is investing her energy efficiently by reaching out to different groups rather than basing all her contacts in the same group.

A second, related kind of advantage is based on the way in which standing at one end of a local bridge can be an amplifier for creativity [88]. Experience from many domains suggests that innovations often arise from the unexpected synthesis of multiple ideas, each of them on their own perhaps well-known, but well-known in distinct and unrelated bodies of expertise. Thus, B ’s position at the interface between three non-interacting groups gives her not only access to the combined information from these groups, but also the opportunity for novel ideas by combining these disparate sources of information in new ways.

Finally, B ’s position in the network provides an opportunity for a kind of social “gate-keeping” — she regulates the access of both C and D to the tightly-knit group she belongs to, and she controls the ways in which her own group learns about information coming from C ’s and D ’s groups. This provides B with a source of power in the organization, and one could imagine that certain people in this situation might try to prevent triangles from forming around the local bridges they’re part of — for example, another edge from C or D into B ’s group would reduce B ’s gatekeeping role.

This last point highlights a sense in which the interests of node B and of the organization as a whole may not be aligned. For the functioning of the organization, accelerating the flow of information between groups could be beneficial, but this building of bridges would come at the expense of B ’s latent power at the boundaries of these groups. It also emphasizes that our analysis of structural holes is primarily a static one: we look at the network at a single point in time, and consider the effects of the local bridges. How long these local bridges last before triadic closure produces short-cuts around them, and the extent to which people in an organization are consciously, strategically seeking out local bridges and trying to maintain them, is less well understood; it is a topic of ongoing research [90, 188, 252, 259].

Ultimately, then, there are trade-offs in the relative positions of A and B . B ’s position at the interface between groups means that her interactions are less embedded within a single group, and less protected by the presence of mutual network neighbors. On the other hand, this riskier position provides her with access to information residing in multiple groups, and

the opportunity to both regulate the flow of this information and to synthesize it in new ways.

Closure and Bridging as Forms of Social Capital. All of these arguments are framed in terms of individuals and groups deriving benefits from an underlying social structure or social network; as such, they are naturally related to the notion of *social capital* [117, 118, 279, 342, 344]. Social capital is a term in increasingly widespread use, but it is a famously difficult one to define [138]. In Alejandro Portes’s review of the topic, he writes, “Consensus is growing in the literature that social capital stands for the ability of actors to secure benefits by virtue of membership in social networks or other social structures” [342].

The term “social capital” is designed to suggest its role as part of an array of different forms of capital, all of which serve as tangible or intangible resources that can be mobilized to accomplish tasks. James Coleman and others speak of social capital alongside *physical capital* — the implements and technologies that help perform work — and *human capital* — the skills and talents that individual people bring to a job or goal [118]. Pierre Bourdieu offers a related but distinct taxonomy, considering social capital in relation to *economic capital* — consisting of monetary and physical resources — and *cultural capital* — the accumulated resources of a culture that exist at a level beyond any one individual’s social circle, conveyed through education and other broad social institutions [17, 75].

Borgatti, Jones, and Everett [74], summarizing discussions within the sociology community, observe two important sources of variation in the use of the term “social capital.” First, social capital is sometimes viewed as a property of a group, with some groups functioning more effectively than others because of favorable properties of their social structures or networks. Alternately, it has also been considered as a property of an individual; used in this sense, a person can have more or less social capital depending on his or her position in the underlying social structure or network. A second, related, source of terminological variation is based on whether social capital is a property that is purely intrinsic to a group — based only on the social interactions among the group’s members — or whether it is also based on the interactions of the group with the outside world.

A view at this level of generality does not yet specify what kinds of network structures are the most effective for creating social capital, and our discussion earlier in this section highlights several different perspectives on the question. The writings of Coleman and others on social capital emphasize the benefits of triadic closure and embedded edges for the reasons discussed above: they enable the enforcement of norms and reputational effects, and hence can help protect the integrity of social and economic transactions. Burt, on the other hand, discusses social capital as a tension between *closure* and *brokerage* — with the former referring to Coleman’s conception and the latter referring to benefits arising from the ability to “broker” interactions at the interface between different groups, across structural holes.

In addition to the structural distinctions between these perspectives, they also illustrate different focuses on groups versus individuals, and on the activity within a group versus its contacts with a larger population. The contrasts are also related to Robert Putnam’s dichotomy between *bonding capital* and *bridging capital* [344]; these terms, while intended informally, correspond roughly to the kinds of social capital arising respectively from connections within a tightly-knit group and from connections between such groups.

The notion of social capital thus provides a framework for thinking about social structures as facilitators of effective action by individuals and groups, and a way of focusing discussions of the different kinds of benefits conferred by different structures. Networks are at the heart of such discussions — both in the way they produce closed groups where transactions can be trusted, and in the way they link different groups and thereby enable the fusion of different sources of information residing in these groups.

3.6 Advanced Material: Betweenness Measures and Graph Partitioning

This is the first in a series of sections throughout the book labeled “Advanced Material.” Each of these sections comes at the end of a chapter, and it explores mathematically more sophisticated aspects of some of the models developed earlier in the chapter. They are strictly optional, in that nothing later in the book builds on them. Also, while these sections are technically more involved, they are written to be completely self-contained, except where specific pieces of mathematical background are needed; this necessary background is spelled out at the beginnings of the sections where it is required.

In this section, we will try formulating more concrete mathematical definitions for some of the basic concepts from earlier in the chapter. The discussion in this chapter has articulated a way of thinking about networks in terms of their tightly-knit regions and the weaker ties that link them together. We have formulated precise definitions for some of the underlying concepts, such as the clustering coefficient and the definition of a local bridge. In the process, however, we have refrained from trying to precisely delineate what we mean by a “tightly-knit region,” and how to formally characterize such regions.

For our purposes so far, it has been useful to be able to speak in this more general, informal way about tightly-knit regions; it helps to be flexible since the exact characterization of the notion may differ depending on the different domains in which we encounter it. But there are also settings in which having a more precise, formal definition is valuable. In particular, a formal definition can be crucial if we are faced with a network dataset and actually want to identify densely connected groups of nodes within it.

This will be our focus here: describing a method that can take a network and break it down into a set of tightly-knit regions, with sparser interconnections between the regions.




Figure 3.12: A co-authorship network of physicists and applied mathematicians working on networks [322]. Within this professional community, more tightly-knit subgroups are evident from the network structure.

We will refer to this as the problem of *graph partitioning*, and the constituent parts the network is broken into as the *regions* arising from the partitioning method. Formulating a method for graph partitioning will implicitly require working out a set of definitions for all these notions that are both mathematically tractable and also useful on real datasets.

To give a sense for what we might hope to achieve from such a method, let's consider two examples. The first, shown in Figure 3.12, depicts the co-authorships among a set of physicists and applied mathematicians working on networks [322]. Recall that we discussed co-authorship networks in Chapter 2 as a way of encoding the collaborations within a professional community. It's clear from the picture that there are tightly-knit groups within this community, and some people who sit on the boundaries of their respective groups. Indeed it resembles, at a somewhat larger scale, some of the pictures of tightly-knit groups and weak ties that we drew in schematic form earlier, in examples such as Figure 3.11. Is there a general way to pull these groups out of the data, beyond using just our visual intuition?




Figure 3.13: A karate club studied by Wayne Zachary [421] — a dispute during the course of the study caused it to split into two clubs. Could the boundaries of the two clubs be predicted from the network structure?

A second example, in Figure 3.13, is a picture of the social network of a karate club studied by Wayne Zachary [421] and discussed in Chapter 1: a dispute between the club president (node 34) and the instructor (node 1) led the club to split into two. Figure 3.13 shows the network structure, with the membership in the two clubs after the division indicated by the shaded and unshaded nodes. Now, a natural question is whether the structure itself contains enough information to predict the fault line. In other words, did the split occur along a weak interface between two densely connected regions? Unlike the network in Figure 3.12, or in some of the earlier examples in the chapter, the two conflicting groups here are still heavily interconnected. So to identify the division in this case, we need to look for more subtle signals in the way in which edges between the groups effectively occur at lower “density” than edges within the groups. We will see that this is in fact possible, both for the definitions we consider here as well as other definitions.

A. A Method for Graph Partitioning

Many different approaches have been developed for the problem of graph partitioning, and for networks with clear divisions into tightly-knit regions, there is often a wide range of methods that will prove to be effective. While these methods can differ considerably in their specifics, it is useful to identify the different general styles that motivate their designs.


General Approaches to Graph Partitioning. One class of methods focuses on identifying and removing the “spanning links” between densely-connected regions. Once these links are removed, the network begins to fall apart into large pieces; within these pieces, further spanning links can be identified, and the process continues. We will refer to these as *divisive* methods of graph partitioning, since they divide the network up as they go.

An alternate class of methods starts from the opposite end of the problem, focusing on the most tightly-knit parts of the network, rather than the connections at their boundaries. Such methods find nodes that are likely to belong to the same region and merge them together. Once this is done, the network consists of a large number of merged chunks, each containing the seeds of a densely-connected region; the process then looks for chunks that should be further merged together, and in this way the regions are assembled “bottom-up.” We refer to these as *agglomerative* methods of graph partitioning, since they glue nodes together into regions as they go.


To illustrate the conceptual differences between these two kinds of approaches, let’s consider the simple graph in Figure 3.14(a). Intuitively, as indicated in Figure 3.14(b), there appears to be a broad separation between one region consisting of nodes 1-7, and another consisting of nodes 8-14. Within each of these regions, there is a further split: on the left into nodes 1-3 and nodes 4-6; on the right into nodes 9-11 and nodes 12-14. Note how this simple example already illustrates that the process of graph partitioning can usefully be viewed as producing regions in the network that are naturally *nested*: larger regions potentially containing several smaller, even more tightly-knit regions “nested” within them. This is of course a familiar picture from everyday life, where — for example — a separation of the global population into national groups can be further subdivided into sub-populations within particular local areas within countries.

In fact, a number of graph partitioning methods will find the nested set of regions indicated in Figure 3.14(b). Divisive methods will generally proceed by breaking apart the graph first at the 7-8 edge, and subsequently at the remaining edges into nodes 7 and 8. Agglomerative methods will arrive at the same result from the opposite direction, first merging the four triangles into clumps, and then finding that the triangles themselves can be naturally paired off.

This is a good point at which to make the discussion more concrete, and to do so we focus on a particular divisive method proposed by Girvan and Newman [184, 322]. The Girvan-Newman method has been applied very widely in recent years, and to social network data in particular. Again, however, we emphasize that graph partitioning is an area in which there is an especially wide range of different approaches in use. The approach we discuss is an elegant and particular widely-used one; however, understanding which types of methods work best in different situations remains a subject of active research.



(a) *A sample network*



(b) *Tightly-knit regions and their nested structure*

Figure 3.14: In many networks, there are tightly-knit regions that are intuitively apparent, and they can even display a *nested* structure, with smaller regions nesting inside larger ones.

The Notion of Betweenness. To motivate the design of a divisive method for graph partitioning, let’s think about some general principles that might lead us to remove the 7-8 edge first in Figure 3.14(a).

A first idea, motivated by the discussion earlier in this chapter, is that since bridges and local bridges often connect weakly interacting parts of the network, we should try removing these bridges and local bridges first. This is in fact an idea along the right lines; the problem is simply that it’s not strong enough, for two reasons. First, when there are several bridges, it doesn’t tell us which to remove first. As we see in Figure 3.14(a), where there are five bridges, certain bridges can produce more reasonable splits than others. Second, there can be graphs where no edge is even a local bridge, because every edge belongs to a triangle — and yet there is still a natural division into regions. Figure 3.15 shows a simple example, where we might want to identify nodes 1-5 and nodes 7-11 as tightly-knit regions, despite




Figure 3.15: A network can display tightly-knit regions even when there are no bridges or local bridges along which to separate it.

the fact that there are no local bridges to remove.

However, if we think more generally about what bridges and local bridges are doing, then we can arrive at a notion that forms the central ingredient of the Girvan-Newman method. Local bridges are important because they form part of the shortest path between pairs of nodes in different parts of the network — without a particular local bridge, paths between many pairs of nodes may have to be “re-routed” a longer way. We therefore define an abstract notion of “traffic” on the network, and look for the edges that carry the most of this traffic. Like crucial bridges and highway arteries, we might expect these edges to link different densely-connected regions, and hence be good candidates for removal in a divisive method.

We define our notion of traffic as follows. For each pair of nodes A and B in the graph that are connected by a path, we imagine having one unit of fluid “flow” along the edges from A to B . (If A and B belong to different connected components, then no fluid flows between them.) The flow between A and B divides itself evenly along all the possible *shortest* paths from A to B : so if there are k shortest paths from A and B , then $1/k$ units of flow pass along each one.

We define the *betweenness* of an edge to be the total amount of flow it carries, counting flow between all pairs of nodes using this edge. For example, we can determine the betweenness of each edge in Figure 3.14(a) as follows.

- Let’s first consider the 7-8 edge. For each node A in the left half of the graph, and each node B in the right half of the graph, their full unit of flow passes through the 7-8 edge. On the other hand, no flow passing between pairs of nodes that both lie in the same half uses this edge. As a result, the betweenness of the 7-8 edge is $7 \cdot 7 = 49$.
- The 3-7 edge carries the full unit of flow from each node among 1, 2, and 3 to each

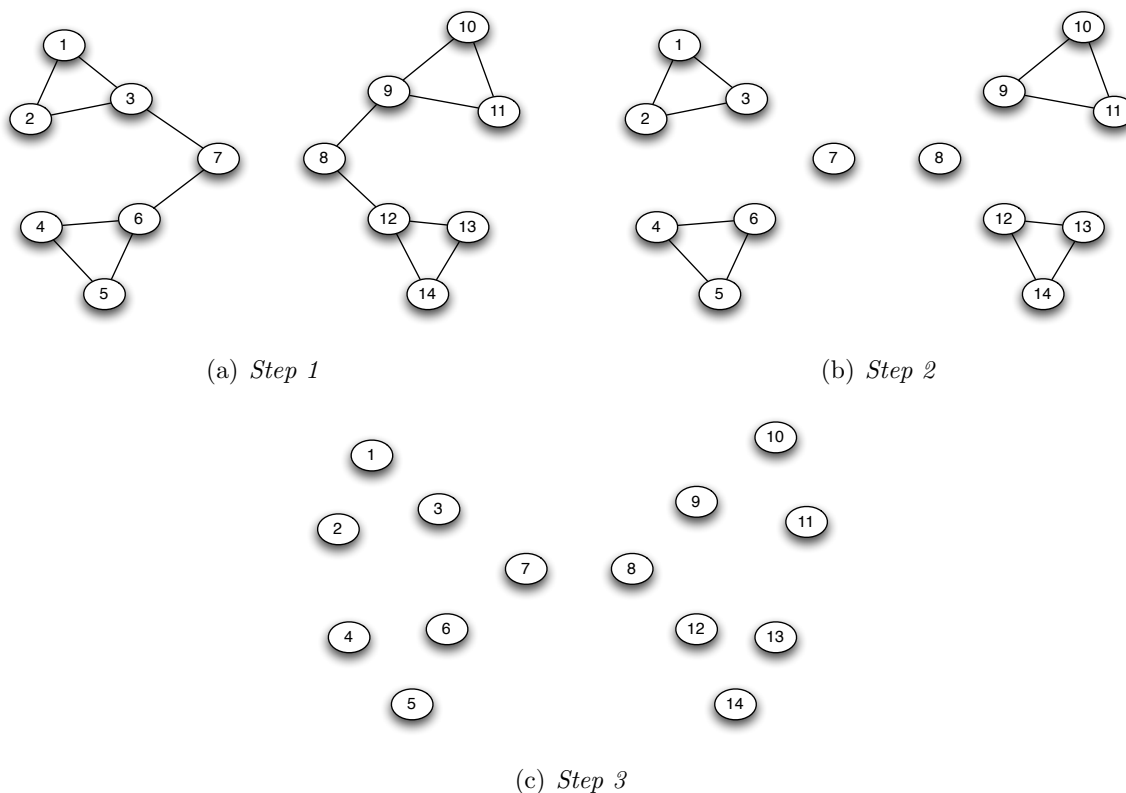


Figure 3.16: The steps of the Girvan-Newman method on the network from Figure 3.14(a).

node among 4-14. Thus, the betweenness of this edge is $3 \cdot 11 = 33$. The same goes for the edges 6-7, 8-9, and 8-12.

- The 1-3 edge carries all the flow from 1 to every other node except 2. As a result, its betweenness is 12. By strictly symmetric reasoning, the other edges linked from 3, 6, 9, and 12 into their respective triangles have betweenness 12 as well.
- Finally, the 1-2 edge only carries flow between its endpoints, so its betweenness is 1. This also holds for the edges 4-5, 10-11, and 13-14.

Thus, betweenness has picked out the 7-8 edge as the one carrying the most traffic.

In fact, the idea of using betweenness to identify important edges draws on a long history in sociology, where most attribute its first explicit articulation to Linton Freeman [73, 168, 169]. Its use by sociologists has traditionally focused more on nodes than on edges, where the definition the same: the betweenness of a node is the total amount of flow that it carries, when a unit of flow between each pair of nodes is divided up evenly over shortest paths. Like edges of high betweenness, nodes of high betweenness occupy critical roles in the network




Figure 3.17: The steps of the Girvan-Newman method on the network from Figure 3.15.

structure — indeed, because carrying a large amount of flow suggests a position at the interface between tightly-knit groups, there are clear relationships of betweenness with our earlier discussions of nodes that span structural holes in a social network [86].

The Girvan-Newman Method: Successively Deleting Edges of High Betweenness.

Edges of high betweenness are the ones that, over all pairs of nodes, carry the highest volume of traffic along shortest paths. Based on the premise that these are the most “vital” edges for connecting different regions of the network, it is natural to try removing these first. This is the crux of the Girvan-Newman method, which can now be summarized as follows.

- (1) Find the edge of highest betweenness — or multiple edges of highest betweenness, if there is a tie — and remove these edges from the graph. This may cause the graph to separate into multiple components. If so, this is the first level of regions in the partitioning of the graph.
- (2) Now recalculate all betweennesses, and again remove the edge or edges of highest betweenness. This may break some of the existing components into smaller components; if so, these are regions nested within the larger regions.
- (...) Proceed in this way as long as edges remain in graph, in each step recalculating all betweennesses and removing the edge or edges of highest betweenness.

Thus, as the graph falls apart first into large pieces and then into smaller ones, the method naturally exposes a nested structure in the tightly-knit regions. In Figures 3.16 and 3.17

we show how the method operates on the graphs from Figures 3.14(a) and 3.15 respectively. Note how smaller regions emerge from larger ones as edges are successively removed.

The sequence of steps in Figure 3.17 in fact exposes some interesting points about how the method works.

- When we calculate the betweennesses in the first step, the 5-7 edge carries all the flow from nodes 1-5 to nodes 7-11, for a betweenness of 25. The 5-6 edge, on the other hand, only carries flow from node 6 to each of nodes 1-5, for a betweenness of 5. (Similarly for the 6-7 edge.)
- Once the 5-7 edge is deleted, however, we recalculate all the betweennesses for the second step. At this point, all 25 units of flow that used to be on this deleted edge have shifted onto the path through nodes 5, 6, and 7, and so the betweenness of the 5-6 edge (and also the 6-7 edge) has increased to $5 + 25 = 30$. This is why these two edges are deleted next.

In their original presentation of the method, Girvan and Newman showed its effectiveness at partitioning a number of real network datasets into intuitively reasonable sets of regions. For example, on Zachary's karate club network in Figure 3.13, when the method is used to remove edges until the graph first separates into two pieces, the resulting partition agrees with the actual split that occurred in the club except for a single person — node 9 in the figure. In real life, node 9 went with the instructor's club, even though the graph partitioning analysis here would predict that he would join the president's club.

Zachary's original analysis of the karate club employed a different approach that also used the network structure. He first supplemented the network with numerical estimates of tie strength for the edges, based on his empirical study of the relationships within the karate club. He then identified a set of edges of minimum total strength whose removal would place node 1 and node 34 (the rival leaders) in different connected components, and he predicted this as the split. The approach Zachary used, deleting edges of minimum total strength so as to separate two specified nodes, is known as the problem of finding a *minimum cut* in a graph, and it has been the subject of extensive research and applications [8, 164, 253]. On the karate-club network, this minimum-cut approach produced the same split as the Girvan-Newman method: it agreed with the split that actually occurred except for the outcome of node 9, an alignment of predictions that emphasizes how different approaches to graph partitioning can produce corresponding results. It is also interesting to note that Zachary traced the anomalous nature of node 9 to a fact that the network structure could not capture: at the time of the actual split, the person corresponding to node 9 was three weeks away from completing a four-year quest to obtain a black belt, which he could only do with the instructor (node 1).

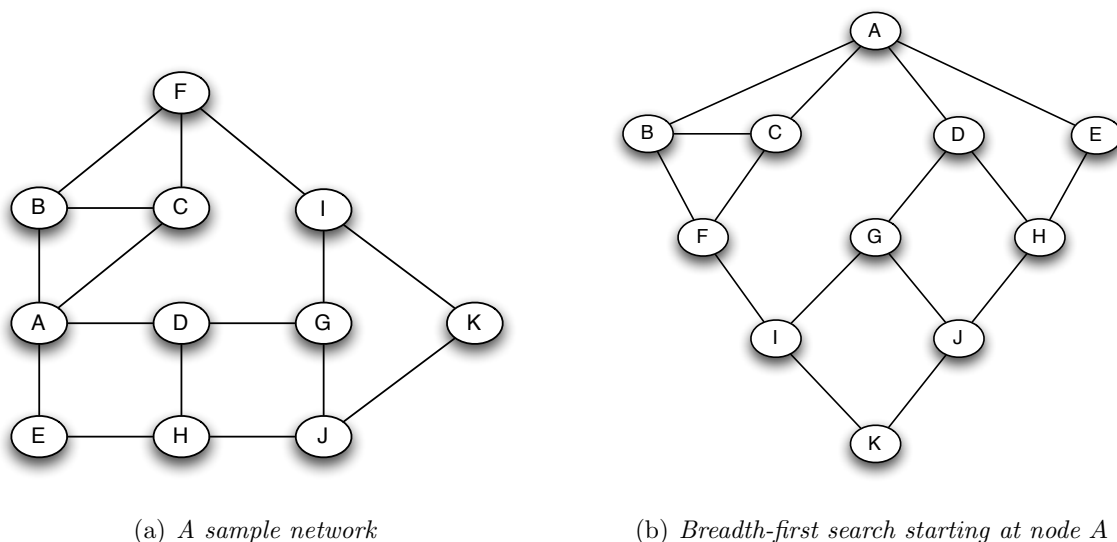


Figure 3.18: The first step in the efficient method for computing betweenness values is to perform a breadth-first search of the network. Here the results of breadth-first from node A are shown; over the course of the method, breadth-first search is performed from each node in turn.

Among the other examples discussed by Girvan and Newman, they provide a partition of the co-authorship network from Figure 3.12, with the top level of regions suggested by the different shadings of the nodes in that figure.

Ultimately, it is a challenge to rigorously evaluate graph partitioning methods and to formulate ways of asserting that one is better than another — both because the goal is hard to formalize, and because different methods may be more or less effective on different kinds of networks. Moreover, a line of recent work by Leskovec et al. has argued that in real social-network data, it is much easier to separate a tightly-knit region from the rest of the network when it is relatively small, on the order of at most a few hundred nodes [275]. Studies on a range of different social and information networks suggest that beyond this size, sets of nodes become much more “inextricable” from the rest of the network, suggesting that graph partitioning approaches on this type of data may produce qualitatively different kinds of results for small networks and small regions than for large ones. This is an area of ongoing investigation.

In the remainder of this section, we address a final important issue: how to actually compute the betweenness quantities that are needed in order to make the Girvan-Newman method work.

B. Computing Betweenness Values

In order to perform the Girvan-Newman method, we need a way to find the edges of highest betweenness in each step. This is done by computing all the betweennesses of all edges and then looking for the ones with the highest values. The tricky part is that the definition of betweenness involves reasoning about the set of *all* the shortest paths between pairs of nodes. Since there could be a very large number of such shortest paths, how can we efficiently compute betweenness without the overhead of actually listing out all such paths? This is crucial for implementing the method on a computer to work with datasets of any reasonable size.

In fact, there is a clever way to compute betweennesses efficiently [77, 317], and it is based on the notion of breadth-first search from Section 2.3. We will consider the graph from the perspective of one node at a time; for each given node, we will compute how the total flow from that node to all others is distributed over the edges. If we do this for every node, then we can simply add up the flows from all of them to get the betweennesses on every edge.

So let's consider how we would determine the flow from one node to all other nodes in the graph. As an example, we'll look at the graph in Figure 3.18(a), focusing on how the flow from node A reaches all other nodes. We do this in three high-level steps; below we explain the details of how each of these steps works.

- (1) Perform a breadth-first search of the graph, starting at A .
- (2) Determine the number of shortest paths from A to each other node.
- (3) Based on these numbers, determine the amount of flow from A to all other nodes that uses each edge.

For the first step, recall that breadth-first search divides a graph into *layers* starting at a given node (A in our case), with all the nodes in layer d having distance d from A . Moreover, the shortest paths from A to a node X in layer d are precisely the paths that move downward from A to X one layer at a time, thereby taking exactly d steps. Figure 3.18(b) shows the result of breadth-first search from A in our graph, with the layers placed horizontally going downward from A . Thus, for example, some inspection of the figure shows that there are two shortest paths (each of length two) from A to F : one using nodes A , B , and F , and the other using nodes A , C , and F .

Counting Shortest Paths. Now, let's consider the second step: determining the number of shortest paths from A to each other node. There is a remarkably clean way to do this, by working down through the layers of the breadth-first search.

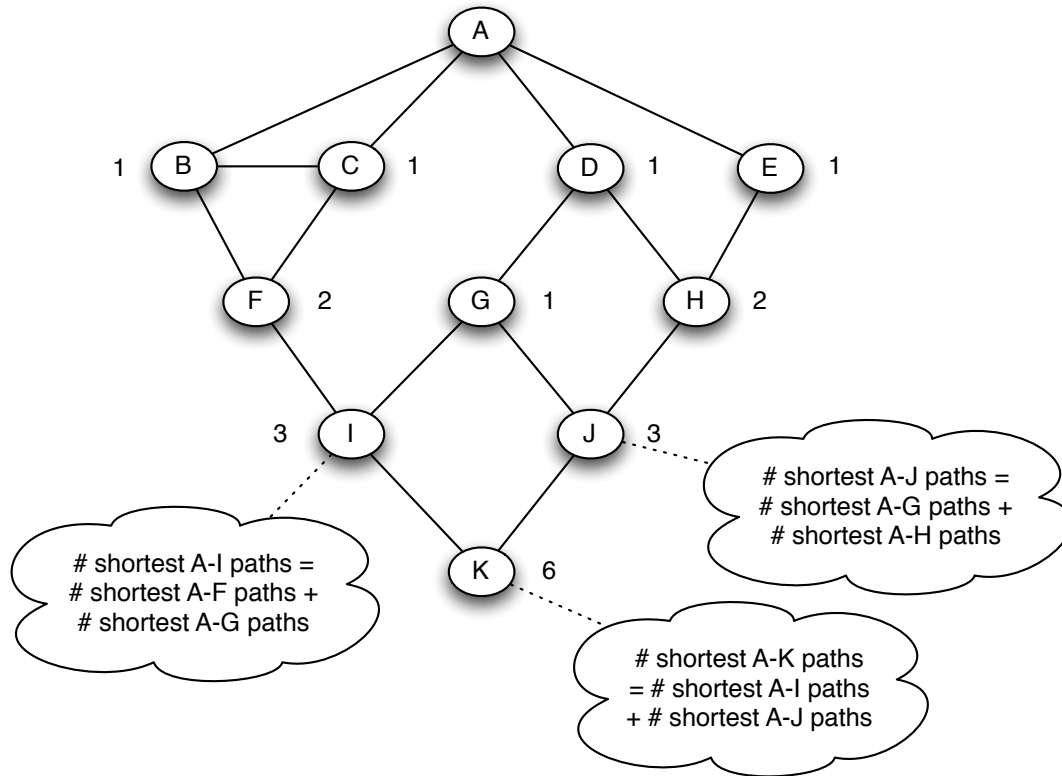


Figure 3.19: The second step in computing betweenness values is to count the number of shortest paths from a starting node A to all other nodes in the network. This can be done by adding up counts of shortest paths, moving downward through the breadth-first search structure.

To motivate this, consider a node like I in Figure 3.18(b). All shortest-paths from A to I must take their last step through either F or G , since these are the two nodes above it in the breadth-first search. (For terminological convenience, we will say that a node X is *above* a node Y in the breadth-first search if X is in the layer immediately preceding Y , and X has an edge to Y .) Moreover, in order to be a shortest path to I , a path must first be a shortest path to one of F or G , and then take this last step to I . It follows that the number of shortest paths from A to I is precisely the number of shortest paths from A to F , plus the number of shortest paths from A to G .

We can use this as a general method to count the number of shortest paths from A to all other nodes, as depicted in Figure 3.19. Each node in the first layer is a neighbor of A , and so it has only one shortest path from A : the edge leading straight from A to it. So we give each of these nodes a count of 1. Now, as we move down through the BFS layers, we apply the reasoning discussed above to conclude that the number of shortest paths to




Figure 3.20: The final step in computing betweenness values is to determine the flow values from a starting node A to all other nodes in the network. This is done by working up from the lowest layers of the breadth-first search, dividing up the flow above a node in proportion to the number of shortest paths coming into it on each edge.

each node should be the *sum* of the number of shortest paths to all nodes directly above it in the breadth-first search. Working downward through the layers, we thus get the number of shortest paths to each node, as shown in Figure 3.19. Note that by the time we get to deeper layers, it may not be so easy to determine these number by visual inspection — for example, to immediately list the six different shortest paths from A to K — but it is quite easy when they are built up layer-by-layer in this way.

Determining Flow Values. Finally, we come to the third step, computing how the flow from A to all other nodes spreads out across the edges. Here too we use the breadth-first search structure, but this time working up from the lowest layers. We first show the idea in Figure 3.20 on our running example, and then describe the general procedure.

- Let's start at the bottom with node K . A single unit of flow arrives at K , and an equal number of the shortest paths from A to K come through nodes I and J , so this unit

of flow is equally divided over the two incoming edges. Therefore we put a half-unit of flow on each of these edges.

- Now, working upward, the total amount of flow arriving at I is equal to the one unit actually destined for I plus the half-unit passing through to K , for a total of $3/2$. How does this $3/2$ amount of flow get divided over the edges leading upward from I , to F and G respectively? We see from the second step that there are twice as many shortest paths from A through F as through G , so twice as much of the flow should come from F . Therefore, we put one unit of the flow on F , and a half-unit of the flow on G , as indicated in the figure.
- We continue in this way for each other node, working upward through the layers of the breadth-first search.

From this, it is not hard to describe the principle in general. When we get to a node X in the breadth-first search structure, working up from the bottom, we add up all the flow arriving from edges directly below X , plus 1 for the flow destined for X itself. We then divide this up over the edges leading upward from X , *in proportion to the number of shortest paths coming through each*. You can check that applying this principle leads to the numbers shown in Figure 3.20.

We are now essentially done. We build one of these breadth-first structures from *each* node in the network, determine flow values from the node using this procedure, and then sum up the flow values to get the betweenness value for each edge. Notice that we are counting the flow between each pair of nodes X and Y twice: once when we do the breadth-first search from X , and once when we do it from Y . So at the end we divide everything by two to cancel out this double-counting. Finally, using these betweenness values, we can identify the edges of highest betweenness for purposes of removing them in the Girvan-Newman method.

Final Observations. The method we have just described can be used to compute the betweennesses of nodes as well as edges. In fact, this is already happening in the third step: notice that we are implicitly keeping track of the amounts of flow through the nodes as well as through the edges, and this is what is needed to determine the betweennesses of the nodes.

The original Girvan-Newman method described here, based on repeated removal of high-betweenness edges, is a good conceptual way to think about graph partitioning, and it works well on networks of moderate size (up to a few thousand nodes). However, for larger networks, the need to recompute betweenness values in every step becomes computationally very expensive. In view of this, a range of different alternatives have been proposed to identify similar sets of tightly-knit regions more efficiently. These include methods of approximating the betweenness [34] and related but more efficient graph partitioning approaches using

divisive and agglomerative methods [35, 321]. There remains considerable interest in finding fast partitioning algorithms that can scale to very large network datasets.

3.7 Exercises

1. In 2-3 sentences, explain what triadic closure is, and how it plays a role in the formation of social networks. You can draw a schematic picture in case this is useful.
2. Consider the graph in Figure 3.21, in which each edge — except the edge connecting b and c — is labeled as a strong tie (S) or a weak tie (W).

According to the theory of strong and weak ties, with the strong triadic closure assumption, how would you expect the edge connecting b and c to be labeled? Give a brief (1-3 sentence) explanation for your answer.




Figure 3.21:

3. In the social network depicted in Figure 3.22, with each edge labeled as either a strong or weak tie, which nodes satisfy the Strong Triadic Closure Property from Chapter 3, and which do not? Provide an explanation for your answer.
4. In the social network depicted in Figure 3.23 with each edge labeled as either a strong or weak tie, which two nodes violate the Strong Triadic Closure Property? Provide an explanation for your answer.
5. In the social network depicted in Figure 3.24, with each edge labeled as either a strong or weak tie, which nodes satisfy the Strong Triadic Closure Property from Chapter 3, and which do not? Provide an explanation for your answer.




Figure 3.22:




Figure 3.23: A graph with a strong/weak labeling.




Figure 3.24: