

CHAPTER THREE

Small Worlds

BACK WHEN STEVE STROGATZ AND I STARTED OUR WORK TOGETHER, we didn't know any of this. Neither of us had the foggiest idea about Rapoport or Granovetter, or really anything about social networks at all. We both knew some physics—in fact, I had majored in it at college. But college had been a military academy, and what little actual knowledge had seeped through the cracks of my officer training, outdoor adventures, and generally earthly preoccupations of a young man's life in the Navy seemed remote and of little immediate relevance. Graph theory was also a mystery. Effectively a branch of pure mathematics, graph theory can be divided roughly into two components—the almost obvious and the utterly impenetrable. I learned the obvious stuff out of a textbook, and after some futile struggling with the rest, convinced myself that it wasn't very interesting anyway.

All this profundity of ignorance left us in something of an awkward place. We were reasonably certain that someone must have thought about this problem before, and we worried that we might waste a lot of time reinventing the wheel. But we also thought that if we went out looking for it, we might get discouraged by how much had already been done, or else trapped into thinking about the problem from the same perspective and so get stuck on the very same things that other people had. After I spent a month at home in Australia to think about it, we met in Steve's office in January 1996 and made up our minds: we would go it alone. Telling almost nobody and reading virtually nothing, we would drop the crickets project and have a go at building some very simple models of social networks to look for features like the small-world phenomenon. No doubt feeling that he needed to protect me from myself, Steve insisted that we give it only four months—a single semester—after which, if we hadn't made some significant progress, we would concede defeat and return to the crickets. At worst, my graduation

would be delayed by a semester, and if it would make me happy, well, why not?

WITH A LITTLE HELP FROM MY FRIENDS

AT THAT POINT I HAD BEEN LIVING IN ITHACA FOR JUST OVER TWO years and was beginning to feel like I had a new home with new friends, but still felt closely connected to my old ones. It occurred to me, however, that if you asked the average Cornell student how close he or she felt to a random person in Australia, the answer would probably be “not very.” After all, most of my friends in America had never met another Australian before, and few of my Australian friends knew any Americans. The two countries are on virtually opposite sides of the planet, and despite a certain cultural similarity and a good deal of mutual fascination, are viewed by most of their inhabitants as being almost impossibly distant, even exotic. Nevertheless, at least some small group of Americans and some small group of Australians actually were very close to each other, although they might not have known it, by virtue of a single common friend—me.

A similar state of affairs applied on a smaller scale between my different groups of friends at Cornell. I belonged to the Department of Theoretical and Applied Mechanics, which was a small graduate department in which there were more foreign students than Americans. I spent an awful lot of time in that department and got to know the other graduate students pretty well. But I also taught rock climbing and skiing in Cornell’s outdoor education program, and most of the friends from Cornell that I still have today were fellow outdoor instructors or students. Finally, I had lived in a big dormitory in my first year and had made some good friends there. My classmates knew each other, my dorm friends knew each other, and my outdoor friends knew each other. But the different groups were all quite, well, *different*. Without me to come and visit, my climbing friends, for example, would have precious little reason ever to venture into my department in Kimball Hall, and (with some justification) tended to regard engineering graduate students as something of a breed apart.

That two people can share a mutual friend whom each regards as “close,” but still perceive each other as being “far away” is a facet of social life at once commonplace and also quite mysterious. As we will see in chapter 5, this paradox lies at the heart of the small-world problem, and through its resolution we can understand not only Milgram’s results but also a number of

other network problems that on the surface have nothing to do with sociology at all. That will take some more work, however. For now, it's enough to say that we don't just have friends, rather we have groups of friends, each of which is defined by the particular set of circumstances—some *context*, like our college dorm or our current workplace—that led to our getting acquainted. Within each group there will tend to be a high density of interpersonal ties, but ties between different groups will typically be sparse.

The groups, however, *are* connected by virtue of individuals who belong to more than one group. In time, these *overlaps* between groups may grow stronger, and the boundaries between them blur, as people from one group start to interact with people from another via the intermediation of a mutual friend. Over the years that I spent at Cornell, my different groups of friends eventually met each other and occasionally became friends themselves. Even some of my Australian friends came to visit, and although they didn't stay long enough to form any lasting relationships, the boundary between the two countries is now, in some small way, less distinct than it was.

After thinking through these ideas a number of times, and a good deal of wandering around a very cold Cornell campus, Steve and I decided that there were four elements we wanted to capture in our model. The first was that social networks consist of many small overlapping groups that are densely internally connected and that overlap by virtue of individuals having multiple affiliations. The second was that social networks are not static objects. New relationships are continually being forged, and old ones abandoned. Third, not all potential relationships are equally likely. Who I know tomorrow depends at least to some extent on who I know today. But the final feature was that we occasionally do things that *are* derived entirely from our intrinsic preferences and characteristics, and these actions may lead us to meet new people who have no connection to our previous friends at all. My decision to move to America was driven solely by my desire to go to graduate school, and I didn't know a soul when I got there, nor did anyone else that I knew. Likewise, my decision to teach climbing was unaffected by my choice of department, as was the dorm I lived in.

In other words, we do what we do in part because of the position we occupy in our surrounding social structure and in part because of our innate preferences and characteristics. In sociology, these two forces are called *structure* and *agency*, and the evolution of a social network is driven by a trade-off between the two. Because agency is the part of an individual's

decision-making process that is *not* constrained by his or her structural position, actions derived from agency appear as random events to the rest of the world. Of course, decisions like moving to another country or going to graduate school are derived from a complicated mixture of personal history and psychology and so are not random at all. But the point is that as long as they aren't determined explicitly by the current social network, we can treat them *as if* they are random.

Once these apparently random affiliations have been made, however, structure reenters the picture, and the newly created overlaps become the bridges over which other individuals can cross and form additional affiliations of their own. The dynamic evolution of relationships in a social network, therefore, is driven by a balance of conflicting forces. On the one hand, individuals make what seem like random decisions to launch themselves into new social orbits. And on the other, they are constrained and enabled by their current friendships to reinforce the group structures that already exist. The \$64,000 question was, How important is one relative to the other?

Obviously we didn't know. And furthermore, we were pretty sure that nobody else would either. The world, after all, is a complicated place, and it is precisely this kind of uncertain, hard-to-measure balance between conflicting forces that makes it so. Fortunately, empirical tangles like this are precisely the places where theory comes into its own. Instead of trying to establish the balance between individuals willfulness and social structure—between randomness and order—that exists in the *actual* world, we could ask the question, What can we learn by looking at *all possible* worlds? In other words, think of the relative importance of order and randomness as a parameter that we can tune in order to move through a space of possibilities, in much the same way that the tuning knob on an old-fashioned radio allows us to scan the spectrum of radio frequencies.

At one end of the spectrum, individuals *always* make new friends through their current friends, and at the other end, they *never* do. Neither extreme is very realistic, but that was the point—by picking unreasonable extremes we hoped that somewhere in the messy middle ground we would find a believable version of reality. And even if we couldn't specify exactly where that point might be, our hope was that much of what lay in between the two extremes would be, in some well-defined sense, the same. What we were looking for was not a single kind of network to hold up as a model of a social

network, but in the spirit of universality, a *class* of networks, each of which might differ in detail from all the others but whose essential properties didn't depend on those details.

Casting around for the right kind of model took some time. The notion of group structure that we had started off with turned out to be harder to capture in a precise manner than we had expected. But eventually the breakthrough came. As always, I immediately rushed down the corridor to Steve's office and hammered on the door until he gave up trying to do whatever it was he was trying to do, and let me in.

FROM CAVEMEN TO SOLARIANS

PERHAPS IT'S NOT SURPRISING, BUT AS A BOY I WAS A KEEN Isaac Asimov fan. In particular I read and reread his two most famous series: the *Foundation* trilogy and the *Robot* series. Oddly enough, the psychohistory of Hari Seldon, the chief protagonist in *Foundation*, was probably my earliest exposure to the idea of emergence in social systems. As Seldon puts it, although the behavior of individuals is hopelessly complex and unpredictable, the behavior of mobs, and even civilizations, is amenable to analysis and prediction. As fantastical as it was when conceived in the early 1950s, Asimov's vision is remarkable in its anticipation of much that the study of complex systems attempts to do today. It was, however, the *Robot* series that I wanted to talk to Steve about.

In *The Caves of Steel*, the first book of the series, detective Elijah Baley sifts through a murder mystery on a futuristic Earth built entirely underground. While doing so, he also contemplates the mysteries of his own life and his relations with his fellow man. Amid the teeming masses of humanity packed into their caves of steel, Baley knows a small, tight-knit group of people very well and almost nobody else. Strangers do not talk to one another, and interactions between friends are physical and personal. In the sequel, *The Naked Sun*, Baley is sent on assignment to the colonized planet of Solaria, which much to his discomfort lies at the opposite end of the social interaction spectrum. Unlike native Earthmen, Solarians live on the surface of a sparsely populated planet. They inhabit enormous estates in solitary isolation, accompanied only by robots, and interact with each other (even their spouses) virtually, via what amounts to a global teleconferencing facility. Back on Earth, life is lived in the security of interlocking and mutually reinforcing ties, and initiating a relationship with a random stranger

would be inconceivable. But on Solaria, all interactions are equally accessible, and prior relationships are relatively unimportant to the establishment of new ones.

Imagine then, two worlds—a world of caves and a world of random, independent relationships—and ask, how do new relationships get formed in each one? Specifically, think about the likelihood of meeting a particular, randomly selected person as a function of how many mutual friends you and that person currently possess. In a caveman world, an absence of mutual acquaintances suggests you live in different “caves,” so you will probably never meet. But if you have even one common friend, the implication is that you live in the same community, move in the same social circles, and so are extremely likely to become acquainted. Obviously, this would be a strange place in which to live, but again the point is to find the extremes. At the other extreme, akin to Solaria, your social history is irrelevant to your future. Even if two people happen to have many friends in common, they are no more or less likely to meet than if they have none.

Each of these general principles for choosing new friends can be expressed more precisely in terms of what we might call *interaction rules*. In the universe of our model, we could construct a network of nodes connected by social ties (let us imagine them to be friendships, although they need not be), and then let the network evolve in time as individuals make new friends according to a specified interaction rule. The two extreme types of worlds, the caveman world and Solaria, for example, can be captured by the rules in Figure 3.1. We can see that the tendency of two people to become friends is determined by the number of friends they currently have in common, but that the precise way in which it is determined varies dramatically from one rule to the other. The top curve corresponds to a caveman world because as soon as two individuals have even a single mutual friend, they immediately display a strong tendency to become friends themselves. The bottom curve, by contrast, corresponds to the Solaria world in which even a very large number of mutual friends has precious little effect on people’s tendency to interact. So under almost all circumstances, they interact randomly.




Figure 3.1. Two extreme kinds of interaction rules. In the top curve (caveman world), even a single mutual friend implies that A and B are highly likely to meet. In the bottom curve (Solaria world), all interactions are equally unlikely, regardless of how many friends A and B share.

The great advantage of formalizing the rules of network evolution in this way is that, as we can see in Figure 3.2, a continuum of *intermediate* rules can be defined as curves that lie between the two extremes. Each one of these rules expresses the tendency of two individuals to become friends as a function of how many mutual friends they have at that moment, but they vary in the extent to which mutual friends matter. Mathematically, this entire *family* of rules can be expressed in terms of an equation that contains a single *tunable parameter*. By adjusting, or *tuning*, the parameter between zero and infinity, we can pick out one of the interaction rules in Figure 3.2, and then build a network that evolves according to this rule. What we have created is a mathematical model of a social network. Because this was the first network model Steve and I ever built, we called it, for want of a better name, the *alpha model*, and so the parameter driving its behavior became *alpha*.

Although we didn't know it at the time, the alpha model was very close in spirit to Anatol Rapoport's random-biased nets. And like Rapoport, we quickly found it impossible to solve anything using only pencil and paper. Fortunately for us, five decades of technological development had finally yielded computers fast enough to do the job by brute force. Indeed, in many ways network dynamics problems are ideal grist for the mill of computer simulation. Very simple rules, at the level of individual actions, can generate bewildering complexity when many such individuals interact over time, each making decisions that necessarily depend on the decisions of the past. Often the results are highly counterintuitive, and pencil-and-paper math very rarely works on its own. Computers, however, love this stuff, the endless, mindless,

blindingly fast iteration of simple rules being precisely what every computer was born to do. In the same way that physicists do experiments in the lab, computers have enabled mathematicians to become experimentalists, testing their theories in a multitude of imaginary laboratories where the rules of reality can be adjusted at will.




Figure 3.2. Between the two extremes, a whole family of interaction rules exists, each one specified by a particular value of the tuneable parameter alpha (α). When $\alpha = 0$, we have a caveman world; when α becomes infinite, we have Solaria.

But what sorts of things were we supposed to test? Remember that the problem we wanted to understand—the origin of the small-world phenomenon—seemed to depend on the presence of two apparently contradictory properties of social networks. On the one hand, the network should display a large *clustering coefficient*, meaning that on average a person’s friends are far more likely to know each other than two people chosen at random. On the other hand, it should be possible to connect two people chosen at random via a chain of only a few intermediaries. Hence, even globally separated individuals will be joined by short chains, or *paths* in the network. Each of these properties is trivial to satisfy on its own, but it wasn’t at all clear how they could be combined. The caveman world of Elijah Baley, for example, is clearly highly clustered, but our intuition would suggest that if all the people we know tend to know only each other, it would be very hard to connect ourselves through them to the rest of the world in only a few steps. All that local redundancy might be good for group cohesion, but it’s clearly no help in fostering global connectivity. By contrast, the Solaria world is much more likely to display short network path lengths. In fact, when people interact purely at random, it is a standard result from graph

theory that on average the typical length of a path between any two of them will be short. However, it is also easy to show that in a random graph, the probability that any of our friends will know each other becomes negligible in the limit of a very large global population; hence, the clustering coefficient will be small. Our intuition, therefore, suggests that either the world can be small or it can be clustered, but it can't be both. Computers, however, are unconcerned with intuition.

SMALL WORLDS

USING PATH LENGTH AND CLUSTERING AS OUR PROBES, WE STARTED to build our *alpha networks* in the computer, first constructing them and then implementing some standard algorithms to measure the corresponding statistics. The programming required was mostly elementary, but I had to teach myself the language in the process, so the resulting code was ugly and sluggish, and I frequently spent many hours trying to track down some bug that killed my program after running happily for a day or more. Computer simulation may be less messy than the real world, but it can still be a pain. After a frustrating month or so, however, we at last had some results to ponder.

At first it seemed that our intuition had been correct. When the value for alpha was low, meaning that the nodes had a strong preference to connect only to friends of friends, the resulting graphs tended to be highly clustered. So much so, in fact, they actually fragmented themselves into many tiny components, or *caves*. Within each cave, everyone was well connected to one another, but between different caves no connections existed at all. This result was actually an inconvenience, because when networks are fragmented in this fashion, it is difficult to define the distance between nodes in different fragments. Fortunately it is possible to define a sensible notion of path length that can account for network disintegration. In the simplest such modification, one measures the shortest path length between pairs of nodes exactly as before, but only computes the average over pairs that lie in the same connected component. The result, as shown in Figure 3.3, is that the typical path length is small when the alpha value is low and it is small when the alpha value is high, but the path length spikes upward when the alpha value is somewhere in the middle. The explanation is that for low alpha, the graph is highly fragmented, but because the average is computed only over nodes in the same connected components (the caves), the small size of the

components themselves yields short path lengths. This is the *Caves of Steel* world—the people who can be reached can be reached easily, and the people who can't be reached easily can't be reached at all. When the alpha value is high, by contrast, the graph is more or less random. As a result, it is now connected into a single universal component, and the typical separation between any pair of nodes is small, just as we know is true for random graphs. This is the Solaria world, in which *everyone* can be reached with roughly the same ease.




Figure 3.3. Path length as a function of alpha (α). At the critical alpha value, many small clusters join to connect the entire network, whose length then shrinks rapidly.

The spike in the middle of Figure 3.3 is where all the interesting behavior is. To the left of the spike, as alpha increases, the fragments are rapidly joining together, resulting in a large increase in the apparent path length. The world is getting bigger, but it is doing so because the previously isolated components are starting to get connected. It is harder to reach people, on average, but more and more people can be reached. To the right of the spike, all the components in the network have been connected into a single entity, and now the average path length starts to collapse rapidly as the interaction rule becomes more random still. Right at the spike is a critical point, a phase transition very similar to the one discussed for random graphs, at which everyone becomes connected but the typical path length between pairs of individuals tends to be very large. At the peak of the spike, in a network of, say, a million people, each of whom has a hundred friends, the typical path length would be in the thousands. A network in which you are a few thousand handshakes from the president is clearly the antithesis of a small world. But—and this is important—such a world is inherently unstable. Almost as soon as the phase transition has occurred, and the network has become globally

connected, the average path length starts to drop like a stone, rapidly approaching its eventual minimum. Although it seemed mysterious at the time, it was this surprisingly rapid drop in the length that turned out to be critical.

The clustering coefficient also displayed some unexpected behavior, first rising to a maximum for some low value of α , and then also falling rapidly, just like the average path length. What was more interesting, however, was the location of this transition relative to the corresponding transition in path length. Because we expected, on the one hand, highly clustered graphs with large characteristic lengths and, on the other hand, poorly clustered graphs with small characteristic lengths, we expected that the transition of both statistics might also correspond to each other. Instead, as Figure 3.4 shows, the length started to plummet just as the clustering was reaching its maximum value.

At first we assumed there was an error in the code, but after some careful checking and some vigorous head scratching, we realized that what we were staring at was the very small-world phenomenon we were looking for. In the universe defined by our model, here was a regime in which networks displayed the high local clustering of disconnected caves but were connected such that any node could be reached from any other in an average of only a few steps. We called this class of networks *small-world networks*, which wasn't perhaps the most scientific of labels but had the great advantage of being catchy. Small-world networks have since received a lot of attention and while somewhere in all the fuss the original α model has been largely forgotten, it still has a few things to teach us about the world.




Figure 3.4. Comparison between path length (L) and clustering coefficient (C). The region between the curves, where L is small and C is large (shaded), represents the presence of small-world networks.

The first thing the alpha model tells us is that either the world will be fragmented into many tiny clusters, like isolated caves, or it will be connected in a single giant component within which virtually anyone can be connected to anyone else. It is not possible, for example, to have two, or even a few, large components between which the world is evenly divided. This result may seem surprising because the world often appears divided along geographical, ideological, or cultural lines into a small number of large and incompatible factions—East and West, black and white, rich and poor, Jewish, Christian, and Muslim. Although these cleavages may drive our perceptions, and thereby affect our actions in important ways, what the alpha model tells us is that they do not apply to the network itself. We are either all connected or not at all connected—there really isn’t anywhere in-between.

Furthermore, it turns out that the highly connected state is overwhelmingly more likely than the highly fragmented one. Our parameter alpha, remember, represents a balance between the constraints of social structure and the freedom of individual agency. At the moment, alpha is a hard parameter to interpret, so it is unclear exactly what a particular value means in terms of the real world. Once we learn a little more about networks, however, it will become clear that just a little bit of agency goes a long way. The result is that the world we live in almost certainly lies to the right of the peak in Figure 3.4, implying that each of us can be connected to everyone else. In fact, the model makes an even stronger statement than this. Because the drop-off on the right side of the peak is so rapid, not only is it likely that the world is globally connected, but it is almost certainly the case that the world is *small* in the sense that almost any pair of individuals can be connected via a short chain of intermediaries. This result may also come as a surprise to many of us who spend most of our lives interacting with a relatively small group of people—friends, family, and coworkers—who are often quite similar to ourselves. Even educated, privileged people can feel isolated within their little communities. They may not be unhappy about it, but they are still likely to feel impossibly distant from the large majority of the world that is completely unlike the relatively tiny portion they actually know. So how is it that despite this (quite real) perception, we really *are* all connected?

The resolution of the paradox is that the clustering coefficient does not drop off nearly as rapidly as the path length. No matter what the network is like on a global scale—whether fragmented or connected, large or small—the clustering coefficient is almost certainly going to be high. Individuals, therefore, have severe limitations imposed on what they can deduce about the world based on what they can observe. A well-known aphorism contends that all politics is local, but really we should say that all *experience* is local—we only know what we know, and the rest of the world, by definition, lies beyond our radar screen. In social networks, the only information we have access to, and therefore the only data we can use to make assessments of the world, lies in our *local neighborhood*—our friends and acquaintances. If most of our friends know each other—if our local neighborhood is highly clustered—and if everyone else’s neighborhood is likewise clustered, then we tend to assume that not all these clusters can be connected.

But they can be, and that’s why the small-world phenomenon is so counterintuitive—it is a global phenomenon, yet individuals are capable only of local measurements. You only know who you know, and maybe *most* of the time, your friends know the same sort of people you do. But if just *one* of your friends is friends with just *one* other person who is friends with someone not like you at all, then a connecting path exists. You may not be able to use that path, you may not know it’s there, and finding it may be difficult. But it *is* there. And when it comes to the propagation of ideas, influence, or even disease, that path can matter whether you know about it or not. Just like Hollywood, who you know is important, but now there’s more to the story: who your friends know, and who *those* people know can be important as well.

AS SIMPLE AS POSSIBLE

THE ALPHA MODEL WAS AN ATTEMPT TO UNDERSTAND HOW SMALL-WORLD networks can be generated in terms of the rules that people follow when making new friends. But once we knew the small-world phenomenon was possible, we wanted to figure out exactly what generated it. It didn’t seem adequate simply to conclude that the effect we were seeing was a function of our parameter alpha, because we didn’t really know what alpha was, and therefore what any particular value of alpha might signify. As simple as the alpha model was, it was still too complicated, so we decided that if we really wanted to understand what was going on, we needed to follow Einstein’s