Small Worlds: The Dynamics of Networks between Order and Randomness by Duncan J. Watts (Princeton University Press, 1999, xvi + 262 pp., $39.50), reviewed by Jerrold W. Grossman (grossman@oakland.edu), Department of Mathematics and Statistics, Oakland University, Rochester, MI 48309-4485. Appeared in *The American Mathematical Monthly* **107** (2000), no. 7 (August-September), 664–668.

Networks are in! From the most obvious example that has forever changed our access to information—the Internet—to the pop cultural folklore that you can reach anyone in the world through a short string of acquaintances ("six degrees of separation"), networks have been forced into our thoughts, often without our realizing that they involve nontrivial mathematics. Even casual users of personal computers understand the tree structure in the organization of their files, folders, and directories; and the ease of clicking to new sites on the World Wide Web or hitting the "Back" button to retreat has provided a gentler introduction to discrete mathematics for the masses than we might have imagined possible.

Clearly the study of networks has assumed a new importance. In mathematics, of course, graph theory has been around for at least 264 years, ever since Leonhard Euler showed that the citizens of Königsberg couldn't traverse their seven bridges without backtracking. (We'll use the terms *graph* and *network* interchangeably.) Some early progress in graph enumeration grew out of chemists' desires to classify molecular structures, and biologists use coupled dynamical systems to study everything from neural networks to the synchronous chirping of crickets. Computer scientists and communications researchers have tackled complex problems in graph theory for decades, and discrete mathematical models have become firmly entrenched in the social sciences, for example in mapping the spread of infectious diseases. In this climate it is not surprising, therefore, that we find the book under review, which grew out of the author's recent PhD dissertation in Theoretical and Applied Mechanics at Cornell University, pitched to a lay audience. Glimpses of the ideas contained here have previously appeared in the journal *Nature* ([4], [7]) and attracted considerable media attention ([2], [6]).

Let's try to model the most basic of social structures—the acquaintanceship relation among people—with a graph: vertices represent people, and two vertices are joined by an edge if the two people know each other, say on a first-name basis. This network is quite sparse; that is, each of the $n \approx 6 \times 10^9$ vertices is joined to just a handful of others, say k on the average. Here k is called the *average degree* of each vertex and is surely orders of magnitude smaller than n. Note that of the $C(n, 2) = n(n-1)/2$ possible edges in this graph, only $kn/2$ are present (we have to divide by 2 because each edge has two endpoints).

At least two aspects of our network deserve attention: clustering, the propensity of the edges to occur in dense subgraphs, which is a local property; and the lengths of paths between pairs of vertices, which is a global property. Because of the nature of social interactions, one should expect this graph to be rather highly clustered—many of your friends tend to know one another. One might also expect most paths in this graph to be fairly long. Psychologist Stanley Milgram in the 1960s studied path lengths by having

random people in Kansas try to get a package delivered to random people in Massachusetts by sending it through a sequence of acquaintances. Surprisingly, Milgram found that the median path length was just six or seven edges.

Now let's see whether we can construct a graph abstractly that will reflect this situation: fairly high clustering and remarkably short paths. Suppose that we want to build a graph with $n = 1000$ vertices and 5000 edges (no loops or parallel edges are allowed); on average, then, every vertex will be adjacent to $k = 10$ others (these are the parameters that Watts uses in most of his simulations). We could arrange the vertices in a circle and put an edge between every vertex and its five nearest neighbors on each side. The result, which graph theorists would call a *power of a cycle* and denote C_{1000}^5, will be highly structured. Because of the "local" nature of this construction, the vertices adjacent to a given vertex will have many edges connecting them; in fact it's easy to count that of the $C(10, 2) = 45$ possible edges that could join the vertices in each such neighborhood, 30 of them are in our graph. Thus the *clustering coefficient* γ is 2/3. (This isn't the only way to achieve high clustering, of course, and Watts gives several other constructions, in which γ is even closer to 1.) On the other hand, vertices in our graph are rather far apart, about a quarter of the circle away on average ($246\frac{10}{37}$ edges to be precise); thus what Watts calls the *(characteristic) path length* L is very large. So this model achieves the high γ we were seeking, but it fails to give the desired low L.

Well, what about a random graph? We can start with the cycle C_{1000}—a circle of 1000 vertices, each joined to its nearest neighbor on each side—and then add 4000 more edges at random, taking care not to double up any edges. The result is a "random" graph with the desired parameters. (We start with the cycle for technical reasons, simply to ensure that the resulting random graph is connected.) Random graphs have been extensively studied for decades, and it is known that processes such as this result in graphs with very short paths. Both theory and Watts' empirical work show that almost surely we get an average path length of $L \approx 3$ in this case. On the other hand, we added only 4000 of the 498,500 possible new edges, so each edge has only a 0.008 chance of occurring. Thus we'd expect to find almost none of the edges joining vertices in the neighborhood of a typical vertex, so our graph has a very low average clustering coefficient. Thus this graph, too, poorly models acquaintanceship.

These failed attempts could be considered two extremes in a continuum of possibilities, from the highly structured to the highly random. Watts developed a mechanism—several mechanisms, actually—to interpolate between them. Here is his simplest one. Let $0 \leq \beta \leq 1$. Start with the highly structured $C_n^{k/2}$ from our first model. Then go through the edges one by one in a systematic way around the circle and, with probability β, "rewire" each edge by reassigning one of its endpoints to a randomly chosen vertex anywhere in the graph; with probability $1 - \beta$ the edge is left unchanged. Clearly if $\beta = 0$ then the result is the original graph, but if $\beta = 1$ then we have in effect a random graph with the same n and k. It is not surprising that once a few edges have been rewired, the path length is lowered dramatically, since these new edges provide significant shortcuts. Thus even small

nonzero values of β should help achieve our goal. (As the author points out, this basic idea was anticipated in a 1988 paper [3] by Béla Bollobás and Fan Chung, which analyzed the expected maximum path length for a cycle plus a random matching, giving each vertex degree 3.) On the other hand, if β isn't very large, then the few rewirings do not lower γ significantly. Watts found through extensive simulations that setting $\beta \approx 0.1$ brought L down to about 5 while keeping γ fairly high. This hybrid model does the trick.

So there do exist networks that have the low characteristic path lengths of random graphs but much larger clustering coefficients than a random graph would have. Watts dubs them *small-world graphs*, the subject of his book. He writes with the zeal of one who believes that he has discovered a key model that may unlock the secrets of relationships in every realm of human endeavor. Chapters 2–4 take us through the graph theoretical discussions and analyses of several variations on this model, both algebraically and empirically via computer simulations. In Chapter 5 Watts turns to some interesting real world situations.

Mathematicians have entertained themselves with a real small-world graph since the 1950s: the collaboration graph C. Here the vertices are mathematicians, and two vertices are joined by an edge if the mathematicians have done research together, resulting in a joint publication (for simplicity, we'll say with or without other coauthors). Paul Erdős (1913–1996) is usually viewed as the "center" of C, because he has over 500 collaborators, far more than any other mathematician. A person's *Erdős number* is just the length of a shortest path from that person to Erdős in C. Thus Paul's collaborators have Erdős number 1, their other collaborators (of which there are about 6000) have Erdős number 2, and so on. I have been studying research collaboration for the past five years, mostly by hand so far in the absence of accessible, reliable, and usable data in electronic form, and our Erdős number web site [5] contains lists of people with Erdős numbers less than 3, statistics on Erdős's neighborhood in C, preprints, miscellaneous mathematical and biographical information from the fanciful to the serious, and numerous links.

The available electronic evidence certainly leads one to conclude that the large connected component of C is a small-world graph. This component seems to have about $n \approx 200{,}000$ vertices and average degree about $k \approx 5$. (Another 100,000 vertices are in components of very much smaller sizes, including of course isolated vertices representing mathematicians with no joint papers.) We haven't yet been able to calculate average path lengths, but it appears that the typical mathematician is within about four steps of Erdős, so L can't be much more than 8. Since people tend to write papers more and more in larger and larger groups (less than half of all new papers listed in *Mathematical Reviews* have just one author) or with others working in their subspecialties, it is reasonable to expect that the average γ is significantly positive, although we haven't yet been able to calculate it, again for lack of reliable electronic data. (For the neighborhood of Paul Erdős, $\gamma \approx 0.012$, but this is not typical since his degree is so large. A more representative example might be my own neighborhood, for which $\gamma = 15/C(19, 2) \approx 0.088$. Someone with only one jointly authored paper, which is a three-author work, would have $\gamma = 1$. Watts seems to ignore

the problem of vertices of degree 1, for which γ would be undefined.)

The author acknowledges this primordial example of a small-world graph but needs to get his hands on actual data in electronic form to test his theories in the real world. The book contains three fascinating examples. The first is the collaboration graph of film actors, in which an edge joins two actors if they have appeared in a movie together. The producers of the "Kevin Bacon Game" web site (see [5] for a link to it) provided the author with the data, and Watts found that $n = 225{,}226$, $k \approx 61$, $L \approx 3.65$, and $\gamma \approx 0.79$, a small-world graph if there ever was one. His next example is based on the network used for electricity transmission in the western United States. For this power grid, $n = 4941$, $k \approx 2.67$, $L \approx 18.7$, and $\gamma \approx 0.08$. Again, L and γ fall into the small-world range. The third example comes from biology: the (almost) complete neural network of the worm *C. elegans*. Once again, it's a small world: $n = 282$, $k \approx 14$, $L \approx 2.65$, and $\gamma \approx 0.28$. The author encourages others to cast the net further and report back on the results: try science citations, word associations, organizational structures, or, of course, the World Wide Web. (A recent study of the Web reported in *Nature* [1] estimates the average path length at 19.)

The last third of *Small Worlds*—Part II—deals with something quite different. Until now we have looked at networks as static: the edges represent certain connections or relationships among the vertices, but we don't expect the network to *do* anything. But graphical models also provide ways to analyze dynamical systems, such as the spread of a disease in a population or the changing strategies of players in a game in which other players' strategies change as well. Watts has fewer results and more speculations here than in Part I, but the potential impact of his model may be far greater.

As an example, consider this density classification problem in the field of cellular automata. We are given a network containing a large number of processors, each of which can be either "on" or "off" at time t, for $t = 0, 1, 2, \ldots$. Each processor decides what state to enter at time $t+1$ based on the states of its neighbors and itself at time t. The problem is to design the processors so that if the system is started at time 0 with more processors on than off [respectively, more off than on], then at some future time all the processors will have entered and forever remain in the on [respectively, off] state. Thus we want to solve a global computational problem with a locally connected system. Research on this problem and similar ones has used tools such as genetic algorithms. Watts provides some experimental evidence that a more naive approach works as well if the underlying network is a small-world graph.

Should you read this book, or is reading this review and scanning the references cited here enough? Certainly if you are interested in the details or subtleties of the various graph constructions (we've only brushed the surface here), or full discussions of the numerous examples, you'll need to go to the source. As a graph theorist, I found *Small Worlds* quite interesting for its blend of theoretical and experimental graph theory and its treatment of such diverse applications. Although the book could be read with comprehension by someone without a lot of advanced mathematical training, some of it will come off as

either too dry or too technical. For example, I counted 120 half-page plots, showing the behavior of the parameters mentioned here and several more, and Chapter 4 contains some tedious elementary algebraic derivations of the expectations associated with the various constructions. The book suffers just a bit from the annoying common errors that scientists and engineers make when writing about our subject: occasionally sloppy mathematics, too many definitions and subscripted symbols, imprecision when precision is called for, the overuse of what appear to be technical terms in nontechnical ways, and misuse of some standard mathematical conventions (here big-O notation, for example). But these are minor gripes, and the book is basically well written.

Watts' book seems like an uneasy compromise between two other books that could have been written: a serious mathematical treatise, in collaboration with an expert in random graph theory, about the properties of these fascinating graphs that interpolate between order and randomness; and a lush account of the applications for a general audience, with the technicalities drastically pared down. Nevertheless, the compromise seems to work on its own terms. There is a nice bibliography: the author has clearly done his homework, and the apparent breadth and depth of his knowledge in so many different disciplines is impressive. (Since leaving Cornell, Duncan Watts has worked with sociologists and economists, among others, in postdoctoral positions at the Center for Social Sciences at Columbia University, the Santa Fe Institute, and the Laboratory for Financial Engineering in the MIT Sloan School of Management.) We should be thankful for the author's enthusiasm for his work and the publicity that such a book generates. One thing is clear: the subject has only begun to be studied, and it will be an exciting and fruitful area for research at all levels. On the one hand, it's something undergraduates can sink their teeth into, and on the other hand, it is full of profound open questions for specialists in mathematics, physics, biology, computer science, and myriad other fields.

REFERENCES

1. Réka Albert, Hawoong Jeong, and Albert-László Barabási, Diameter of the World-Wide Web, *Nature* **401** (9 September 1999) 130–131.
2. Sandra Blakeslee, Mathematicians show how big world becomes a small one, *The New York Times*, 16 June 1998, F3.
3. B. Bollobás and F. R. K. Chung, The diameter of a cycle plus a random matching, *SIAM J. on Discrete Math.* **1** (1988) 328–333.
4. James J. Collins and Carson C. Chow, It's a small world, *Nature* **393** (4 June 1998) 409–410.
5. Jerrold W. Grossman, The Erdős Number Project, World Wide Web site, January, 2000, `http://www.oakland.edu/~grossman/erdoshp.html`.
6. Ivars Peterson, Close connections: It's a small world of crickets, nerve cells, computers, and people, *Science News* **154** (22 August 1998) 124–126.
7. Duncan J. Watts and Steven H. Strogatz, Collective dynamics of 'small-world' networks, *Nature* **393** (4 June 1998) 440–442.