

Clustering and preferential attachment in growing networks

M. E. J. Newman

Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501

(Received 11 April 2001; published 26 July 2001)

We study empirically the time evolution of scientific collaboration networks in physics and biology. In these networks, two scientists are considered connected if they have coauthored one or more papers together. We show that the probability of a pair of scientists collaborating increases with the number of other collaborators they have in common, and that the probability of a particular scientist acquiring new collaborators increases with the number of his or her past collaborators. These results provide experimental evidence in favor of previously conjectured mechanisms for clustering and power-law degree distributions in networks.

DOI: 10.1103/PhysRevE.64.025102

PACS number(s): 89.75.Hc, 89.65.-s, 01.30.-y, 89.70.+c

Many systems take the form of networks: sets of nodes, or vertices, joined together by links, or edges. The Internet, the power grid, social networks, food webs, distribution networks, and metabolic networks are commonly cited examples. Investigations of networks within the physics community fall loosely into two categories: (i) studies of static network structure [1–6] and dynamical processes taking place on fixed networks [7–9]; (ii) studies of the dynamics of networks themselves: how and why their topology changes over time [1,2,10–12]. It is this second category that we address here, focusing on two properties which have received a large amount of attention in the literature, clustering and preferential attachment.

Sociologists have long known that social networks (networks of personal acquaintances, for example) display a high degree of transitivity, meaning that there is a heightened probability of two people being acquainted if they have one or more other acquaintances in common. In the physics literature this phenomenon is called “clustering.” Watts and Strogatz [1] measured clustering in a number of real-world networks, including both social and physical networks, by calculating a clustering coefficient C , equal to the probability that two vertices that are both neighbors of the same third vertex will be neighbors of one another. In mathematical notation,

$$C = \frac{3 \times (\text{number of triangles on the graph})}{(\text{number of connected triples of vertices})}. \quad (1)$$

Here a “triangle” is a trio of vertices each of which is connected to both of the others, and a “connected triple” means a vertex which is connected to an (unordered) pair of other vertices. Watts and Strogatz found that in many networks this clustering coefficient is much higher than its expected baseline value, which is set by comparison with a random graph.

It has also been pointed out by a number of authors [3–5,13], particularly in studies of the Internet and the World Wide Web, that real-world networks have highly skewed distributions of vertex degree. (The degree of a vertex is the number of other vertices to which it is connected.) In many cases, the degree distribution is found to follow a power law, a particularly telling functional form which often signifies an underlying process worthy of study.

Explanations have been put forward for both of these observations. In the case of clustering, it is conjectured that pairs of individuals with a common acquaintance (or several) are likely to become acquainted themselves through introduction by their mutual friend(s) [2]. In the case of degree distributions, it is conjectured that, for a variety of reasons, vertices accumulate new edges in proportion to the number they have already, leading to a multiplicative process which is known to give power-law distributions [10–12]. This process is often called “preferential attachment.” While both of these explanations are, in some contexts at least, perfectly plausible, there has been little if any empirical evidence in their favor. This is a glaring problem for two conjectures which have formed the foundation of a substantial body of research. The principal reason for this has been the lack of good time-resolved data on how networks grow.

In order to test a conjecture such as “people with many common friends are more likely to become acquainted than those with few or none,” one needs to watch a network grow and see if the process described by the conjecture does indeed happen with significantly heightened frequency. Although data on the structure of networks are quite plentiful, data on how they grow have proved harder to come by. Recently, however, the author conducted some empirical studies of collaboration networks of scientists: networks in which pairs of scientists are linked together if they have coauthored one or more papers [14–16]. These collaboration networks are true social networks, since two scientists who have coauthored a paper will normally be acquainted with one another. (There are occasional exceptions; see Ref. [15].) They are also well documented, since there exist extensive machine-readable bibliographies of the scientific literature. What’s more, as Barabási *et al.* have recently pointed out [17], these networks have excellent time resolution, because each paper comes with a publication or receipt date. As we now show, this allows us to test directly the clustering and preferential attachment conjectures.

In this study we look at collaboration networks derived from two bibliographic sources:

- (1) The Los Alamos E-print Archive, a database of preprints in physics, self-submitted by their authors;
- (2) Medline, a database of published papers in biology and medicine, whose entries are professionally maintained

by the National Institutes of Health.

While neither of these databases records the exact publication date of the papers they contain, both include a record of the sequence in which papers were added to the database. This is enough for our purposes: all that we need for our calculations is the order of the collaborations undertaken by each author in the database, and the order of the papers is a reasonable proxy for this (probably not correct in every case, but assumed to be correct in most). Two other databases that we studied previously [14] do not contain enough information to establish order of collaborations, recording publication or database entry of papers to the nearest year only. This creates ambiguity since many authors produce more than one paper a year, and so we did not use these databases for the current study.

Authors are identified by their full surname and all initials. As discussed previously [14,15], an author who gives their name differently on different papers may be confused for two people by this measure, while two people with identical surnames and initials may be confused for one. The error in the number of vertices in the network as a result of these problems was found to be on the order of 5%.

We study a six-year interval of time for both databases. (For the Los Alamos Archive we use 1995 to 2000 inclusive, for Medline 1994 to 1999.) Over this period the Los Alamos Archive records 58 342 distinct names, and Medline 1 648 660. In each of the calculations presented here, we use the first five of the six years to construct a collaboration network, and then examine how that network further changes in the remaining one year. Our assumption is that any scientist who is currently active will produce at least one paper during the initial five year period, as will any currently active collaboration between a pair of scientists, so that the network we have at the end of that period will be essentially complete. New vertices added in the sixth year represent, it is assumed, new individuals entering the field, and new edges represent genuine new collaborations. Of course there are some exceptions, such as established scientists who for one reason or another fail to publish anything for five years and then produce a paper in the sixth, and these will be misrepresented in our calculations. We assume these are a small fraction of the total. There will also be some scientists who leave the field during the six years, to go into different fields or professions, or because they retire. We make no attempt to guess which individuals leave in this way: everyone whose name appears even once is considered a member of the network for the entire period of study thereafter. This will introduce some error into our calculations. However, it is straightforward to convince oneself that the correlations we are looking for in the present study will only be weakened by this error, not strengthened, so there is no danger of false positive results.

Let us consider first the question of clustering in the network. We already know that the clustering coefficient is high in our collaboration networks: 0.45 for the Los Alamos Archive and 0.088 for Medline over a five-year period [14]. The calculation presented here improves on these results in two ways. First, the simple clustering coefficient includes

contributions from collaborations between authors which *preceded* their collaborations with any mutual acquaintances. By using time-resolved data we can exclude these collaborations from our measure of clustering. Second, we can determine whether the probability of two individuals collaborating increases as the number m of their previous mutual acquaintances goes up. If this is the case, then it suggests that the standard explanation of clustering, introduction of future collaborators to one another by common previous acquaintances, is correct, the probability of such an introduction presumably increasing with m . Other explanations, such as the institutional explanation proposed in Ref. [15], would be harder to justify.

Measuring the probability of collaboration between authors as a function of their number of mutual acquaintances is complicated by the fact that both the size of the graph and the numbers of mutual acquaintances themselves are changing over time. We consider the probability $P_m(t)$ that the two scientists connected by a link added at time t have m mutual acquaintances. (Time is somewhat arbitrary here. It can be real time, but it can also be any other function which increases monotonically as papers are added to the database; only the order of the papers matters, not their precise timing. The links created by a paper with three or more authors are all considered to be added at the same instant.) We have

$$P_m(t) = \frac{n_m(t)}{\frac{1}{2} N(t)[N(t) - 1]} R_m, \quad (2)$$

where $n_m(t)$ is the number of pairs with m mutual acquaintances immediately *before* the addition of the paper at time t , $N(t)$ is the current number of authors in the network, and R_m is the relative probability of collaboration between the two scientists connected by this link, i.e., the ratio between the actual probability of their collaborating and the probability of their collaborating in a network in which presence of mutual acquaintances makes no difference. We assume that the probability that two scientists with a given value of m collaborate at a particular time does not depend on the number of other scientists with that value of m , or on the size of the database in which the paper they write is archived, and hence that R_m is independent of t [18]. This makes it a suitable quantity to measure to test our clustering hypothesis. In a world with no clustering, we would have $R_m = 1$ for all m ; in a world in which clustering arises through introductions, as above, it should increase with increasing m .

To measure R_m , one simply constructs a histogram of the value of m for each link added to the graph in which each sample is weighted by a factor of $\frac{1}{2} N(t)[N(t) - 1]/n_m(t)$. In Fig. 1 we do this for the network of the Los Alamos Archive. As discussed above, we evaluate R_m for the last of our six years only, the previous five being used to establish the initial network for the calculation. As the figure shows, R_m does indeed increase with m , and is much greater than 1 for all $m > 0$. A pair of scientists who have five mutual previous collaborators, for instance, are about twice as likely to collaborate as a pair with only two, and about 200 times as likely as a pair with none. R_m increases roughly linearly for small m , perhaps indicating that each common collaborator

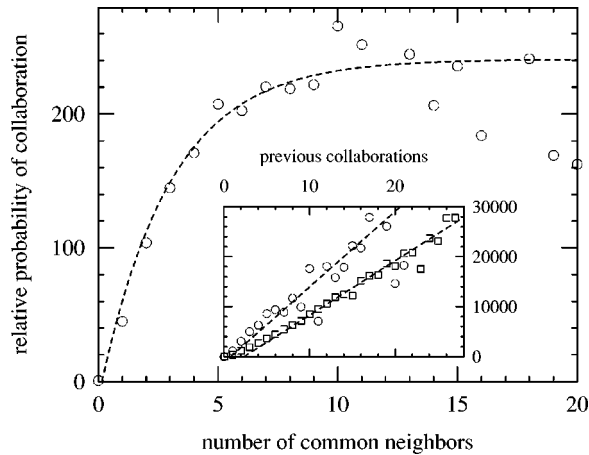


FIG. 1. Relative probability of collaboration between scientists in the Los Alamos Archive as a function of their number of mutual previous collaborators. The dotted line is the best fit of the form (3). Inset: the relative probability of collaboration as a function of number of previous collaborations of the same scientists, for the Los Alamos Archive (circles) and Medline (squares). The dotted lines are the best straight-line fits to the data. The data for Medline have been divided by a factor of 50 vertically to improve the clarity of the figure.

of a pair of scientists is equally likely to introduce them. The curve appears to flatten off for higher m , although the data become poor for $m \geq 8$, since the number of pairs of authors with this many common collaborators who have not already collaborated themselves is very small.

As well as supporting the standard explanation of clustering in social networks, our data for R_m might prove useful for modeling purposes. For example, in some models of the growth of social networks [2,19], a particular form is assumed for the probability of individuals becoming acquainted, as a function of their number of mutual friends. Figure 1 provides a rough empirical guide for what that functional form should be. In the figure we give a fit to the data of the form

$$R_m = A - B e^{-m/m_0}, \quad (3)$$

where A , B , and m_0 are constants. This form appears to fit reasonably well and might be suitable for use in the models.

In the calculation described above, we included only newly appearing edges in the network. Repeat collaborations between authors who had collaborated before were excluded; we assume that such collaborations are more likely to be a result of previous acquaintance than the result of network structure. This, however, raises another interesting question: does probability of collaboration also increase with the number of times one has collaborated before? The answer is yes, as shown in the inset of Fig. 1, which measures the relative probability S_n (defined similarly to R_m above) of two coauthors collaborating if they have collaborated n times previously within the period covered by our study. If collaboration probability were independent of previous collaboration, we would have $S_n = 1$ for all n , but as the figure shows, S_n increases roughly linearly with n , indicating that number of

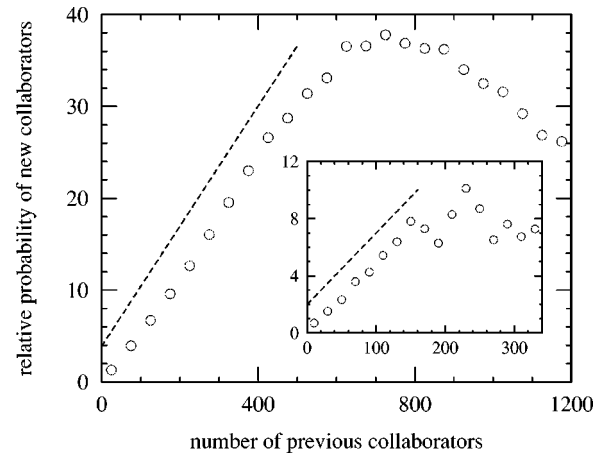


FIG. 2. Relative probability that a new edge in the collaboration network will connect to a vertex of given degree. The main figure shows data from the Medline database, the inset data from the Los Alamos E-print Archive.

past collaborations is a good indicator of the probability of future collaboration. However, one must bear in mind that this calculation may be influenced by varying frequencies of collaboration: regular collaborators who publish often will have more publications in the database as well as greater likelihood of publishing again in the last of our six years, producing a correlation just as seen in the figure. To eliminate this effect one would have to look at data for a longer period of time and compare collaborators with similar numbers of publications but different publication rates. Unfortunately, this is not practical with the data available to us at present.

We can also use our data to test for preferential attachment in the collaboration network. Barabási *et al.* [17] have previously looked for preferential attachment in two collaboration networks derived from data for publications in mathematics and neuroscience. Papers in their databases were dated only to the nearest year, making the order in which collaborations occur uncertain, as discussed above. To get around this, they restricted themselves to measuring the number of new papers each author in the network published in a single year, as a function of number of previous papers. This should be an increasing function if there is preferential attachment, or constant otherwise. Their results show a clear increase and hence favor preferential attachment.

Using our data we can measure preferential attachment in our networks directly by a method similar to the one we used to measure clustering above. We define a relative probability T_k that a link added at time t connects to a vertex representing a scientist who has collaborated previously with k others. By analogy with Eq. (2), the corresponding time-dependent absolute probability $P_k(t)$ that this link connects to a vertex with degree k is $P_k(t) = T_k n_k(t) / N(t)$, where $n_k(t)$ is the number of vertices with degree k immediately before addition of this link. Then T_k can be estimated by making a histogram of the degrees k of the vertices to which each link is added in which each sample is weighted by a factor of $N(t)/n_k(t)$. If there is no preferential attachment, T_k should equal 1 for all k . If there is preferential attachment, it should

be an increasing function of k , and the widely held belief is that it should in fact increase linearly with k . If it increases linearly, then the resulting degree distribution of the network will be a power law [10–12].

In Fig. 2 and its inset we show empirical results for T_k for the databases studied here. As the figure shows, the relative probability is in both cases close to linear in the initial part of the curve, but falls off once k becomes large. This is understandable: no one can collaborate with an infinite number of people in a finite period of time, so at some point T_k must start to decrease. This point appears to be around 150 collaborators in physics and 600 in biomedicine. Interestingly, these figures coincide roughly with the points at which the observed degree distributions in these networks start to deviate from the power-law form [15], lending support to the theory that preferential attachment is the origin of the power law.

Our results differ somewhat from those of Barabási *et al.* [17], who found preferential attachment for their networks, but did not find linear behavior. In the language used here, their finding was that $T_k \sim k^\nu$, with $\nu \approx 0.8$. This form does not fit our data very well. A power-law fit to the increasing part of T_k for our data gives $\nu = 1.04 \pm 0.04$ for Medline and $\nu = 0.89 \pm 0.09$ for the Los Alamos Archive, both of which

are roughly compatible with the conjecture of linear preferential attachment, while only the latter is compatible with $\nu = 0.8$. In practice however, this difference may have little effect. As Krapivsky *et al.* [11] have shown, sublinear preferential attachment gives rise to a stretched exponential cutoff in the resulting degree distribution, but we already have a similar cutoff in our distribution as a result of the deviation of T_k from linear behavior for large enough k .

To conclude, we have measured the probability of collaboration between scientists in two collaboration networks as a function of their number of mutual acquaintances in the network, their number of previous collaborations, and their number of previous collaborators. We find that the probability of collaboration is strongly positively correlated with each of these, and for the latter two that the relationship is close to linear over a large part of its range. These results lend strong support to previously conjectured theories about the way in which networks grow.

The author thanks László Barabási and Duncan Watts for helpful conversations, and Erzsébet Ravasz for providing an early preprint of Ref. [17]. This work was funded in part by the National Science Foundation and by Intel Corporation.

-
- [1] D. J. Watts and S. H. Strogatz, *Nature (London)* **393**, 440 (1998).
 - [2] D. J. Watts, *Small Worlds* (Princeton University Press, Princeton, NJ, 1999).
 - [3] R. Albert, H. Jeong, and A.-L. Barabási, *Nature (London)* **401**, 130 (1999).
 - [4] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener, *Comput. Netw.* **33**, 309 (2000).
 - [5] L. A. N. Amaral, A. Scala, M. Barthélémy, and H. E. Stanley, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 11 149 (2000).
 - [6] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, *Phys. Rev. E* **64**, 026118 (2001).
 - [7] S. H. Strogatz, *Nature (London)* **410**, 268 (2001).
 - [8] R. Monasson, *Eur. Phys. J. B* **12**, 555 (1999).
 - [9] C. Moore and M. E. J. Newman, *Phys. Rev. E* **61**, 5678 (2000).
 - [10] A.-L. Barabási and R. Albert, *Science* **286**, 509 (1999).
 - [11] P. L. Krapivsky, S. Redner, and F. Leyvraz, *Phys. Rev. Lett.* **85**, 4629 (2000).
 - [12] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, *Phys. Rev. Lett.* **85**, 4633 (2000).
 - [13] M. Faloutsos, P. Faloutsos, and C. Faloutsos, *Comput. Commun. Rev.* **29**, 251 (1999).
 - [14] M. E. J. Newman, *Proc. Natl. Acad. Sci. U.S.A.* **98**, 404 (2001).
 - [15] M. E. J. Newman, *Phys. Rev. E* **64**, 016131 (2001).
 - [16] M. E. J. Newman, *Phys. Rev. E* **64**, 016132 (2001).
 - [17] A.-L. Barabási, H. Jeong, Z. Néda, E. Ravasz, A. Schubert, and T. Vicsek, e-print cond-mat/0104162.
 - [18] This assumption is debatable. Certainly scientific collaboration patterns have changed over the last century, so that R_m probably does vary with t on long time-scales. Whether this variation is visible over the course of a single year is unclear.
 - [19] E. M. Jin, M. Girvan, and M. E. J. Newman (unpublished).