

Graph animals, subgraph sampling, and motif search in large networksKim Baskerville,¹ Peter Grassberger,^{2,3} and Maya Paczuski²¹*Perimeter Institute for Theoretical Physics, Waterloo, Canada N2L 2Y5*²*Complexity Science Group, University of Calgary, Calgary, Canada T2N 1N4*³*Institute for Biocomplexity and Informatics, University of Calgary, Calgary, Canada T2N 1N4*

(Received 15 February 2007; published 11 September 2007)

We generalize a sampling algorithm for lattice animals (connected clusters on a regular lattice) to a Monte Carlo algorithm for “graph animals,” i.e., connected subgraphs in arbitrary networks. As with the algorithm in [N. Kashtan *et al.*, *Bioinformatics* **20**, 1746 (2004)], it provides a weighted sample, but the computation of the weights is much faster (linear in the size of subgraphs, instead of superexponential). This allows subgraphs with up to ten or more nodes to be sampled with very high statistics, from arbitrarily large networks. Using this together with a heuristic algorithm for rapidly classifying isomorphic graphs, we present results for two protein interaction networks obtained using the tandem affinity purification (TAP) method: one of *Escherichia coli* with 230 nodes and 695 links, and one for yeast (*Saccharomyces cerevisiae*) with roughly ten times more nodes and links. We find in both cases that most connected subgraphs are strong motifs (Z scores >10) or antimotifs (Z scores <-10) when the null model is the ensemble of networks with fixed degree sequence. Strong differences appear between the two networks, with dominant motifs in *E. coli* being (nearly) bipartite graphs and having many pairs of nodes that connect to the same neighbors, while dominant motifs in yeast tend towards completeness or contain large cliques. We also explore a number of methods that do not rely on measurements of Z scores or comparisons with null models. For instance, we discuss the influence of specific complexes like the 26S proteasome in yeast, where a small number of complexes dominate the k cores with large k and have a decisive effect on the strongest motifs with 6–8 nodes. We also present Zipf plots of counts versus rank. They show broad distributions that are not power laws, in contrast to the case when disconnected subgraphs are included.

DOI: [10.1103/PhysRevE.76.036107](https://doi.org/10.1103/PhysRevE.76.036107)

PACS number(s): 89.75.Fb, 02.70.Uu, 05.10.Ln, 87.10.+e

I. INTRODUCTION

Recently, there has been an increased interest in complex networks, partly triggered by the observation that naturally occurring networks tend to have fat-tailed or even power law degree distributions [1,2]. Thus real-world networks tend to be very different from the completely random Erdős-Renyi [3] networks that have been much studied by mathematicians, and which give Poissonian degree distributions. In addition, most networks have further significant properties that arise either from functional constraints, from the way they have grown (fat tails, e.g., are naturally explained by preferential attachment), or for other reasons. As a consequence, a large number of statistical indicators have been proposed to distinguish between networks with different functionality (neural networks, protein transcription networks, social networks, chip layouts, etc.) and between networks which were specially designed or which have grown spontaneously (such as, e.g., the world wide web), under more or less strong evolutionary pressure. These observables include various centrality measures [4], assortativity (the tendency of nodes with similar degree to link preferentially) [4], clustering [5,6], different notions of modularity [2,7–10], properties of loop statistics [11], the small world property (i.e., slow increase of the effective diameter of the network with the number of nodes) [12], bipartivity (the prevalence of even-sized closed walks over closed walks with an odd number of steps) [13], and others.

The frequency of specific subgraphs form a particular class of indicators. Subgraphs that occur more frequently

than expected are referred to as motifs, while those occurring less frequently are antimotifs [14–17]. Typically, motif search requires a null model for deciding when a subgraph is overabundant or underabundant. The most popular null model so far has been the ensemble of all random graphs with the same degree sequence. This popularity is largely due to the fact that it can be simulated easily by means of the so-called “rewiring algorithm” [18,19]. As we shall see, however, in the present analysis its value is severely limited, because it gives predictions that are too far from those actually observed. Other null models that retain more properties of the original network have been suggested [14,21], but have received much less attention. Analytic approaches to null models are discussed in Refs. [22–24].

A. Motifs and the search for structure

Up to now, motif search has been mainly restricted to small motifs, typically with three or four nodes. Certain specific classes of larger subgraphs have been examined in Refs. [16,20,32]. With the exception of Ref. [31], few systematic attempts have been made to learn about significant structures at larger scale, by counting all possible subgraphs. (For a different approach to the discovery of structure than discussed here, see the work on inference of hierarchy in Ref. [25].)

One reason for this is that the number of nonisomorphic (i.e., structurally different) subgraphs in any but the most trivial networks increases extremely fast (superexponentially) with their size. For instance, the number of different

undirected graphs with 11 nodes is $\approx 10^9$ [26]. Thus exhaustive studies of all possible subgraphs with >10 nodes becomes virtually impossible with present-day computers. But just because of this inflationary growth, counts at intermediate sizes contain an enormous amount of potentially useful information. Another obstacle is the notorious graph isomorphism problem [27,28], which is in the NP class (though probably not NP complete [29]). Existing state of the art programs for determining whether any two graphs are isomorphic [30] remain too slow for our purpose. Instead, we shall use heuristics based on graph invariants similar to those put forward in Ref. [31], where intermediate size motifs and antimotifs in the protein interaction network of *Escherichia coli* were detected.

The last problem when studying larger motifs, and the main one addressed in the present work, is the difficulty of estimating how often each possible subgraph appears in a large network, i.e., of obtaining a “subgraph census.” Most studies so far were based on exact enumeration. In a network with N nodes, there are $\binom{N}{n}$ subgraphs of size n . With $N=500$ and $n=6$, say, this number is $\approx 5 \times 10^{11}$. In addition, most of the subgraphs generated this way on a sparse network would be disconnected, while connected subgraphs are of more intrinsic interest. Thus some statistical sampling is needed. If one is willing to generate disconnected as well as connected subgraphs, then uniform sampling is simple: Just choose random n -tuples of nodes from the network [31]. Uniform sampling connected subgraphs is less trivial. To our knowledge, the only work which addressed this systematically was Kashtan *et al.* [32] (for a less systematic approach, see also [33]). There, a biased sampling algorithm was put forward. While generating the subgraphs is fast, computing the weight factor needed to correct for the bias is $\exp[O(n)]$, making their algorithm inefficient for $n \geq 7$.

B. Graph animals

In the present paper we exploit the fact that sampling connected subgraphs of a finite graph resembles sampling connected clusters of sites on a regular lattice. The latter is called the *lattice animal* problem [34], whence we propose to call the subgraph counting problem that of *graph animals*. It is important to recognize obvious differences between the two cases. In particular, lattices are infinite and translationally invariant, while networks are finite and heterogeneous (disordered). For lattice animals one counts the number of configurations up to translations (i.e., per unit cell of the lattice), while on a network the quantity of immediate interest is the absolute number of occurrences of particular subgraphs. Still, apart from these issues, the basic operations involved in both cases coincide.

Algorithms for enumerating lattice animals exactly exist and have been pushed to high efficiency [35], but are far from trivial [36]. Due to disorder, we should expect the situation to be even worse for graph animals. Algorithms for stochastic sampling of lattice animals are divided into two groups: Markov chain Monte Carlo (MCMC) algorithms take a connected cluster and randomly deform it while preserving connectivity [37–39], while “sequential” sampling

algorithms grow the cluster from scratch [40,41,43,44]. Even for regular lattices, MCMC algorithms seem less efficient than growth algorithms [41]. For networks, this difference should be even more pronounced, since MCMC algorithms would dwell in certain parts of the network, and averaging over the different parts costs additional time. Thus we shall in the following concentrate only on growth algorithms.

All growth algorithms similar to those in [40,41,43,44] produce unbiased samples of *percolation* clusters. As explained in Sec. II, this means that they sample clusters or subgraphs with nonuniform probability (for an alternative algorithm, see [45]). Consequently, computing graph animal statistics requires the computation of weights to be assigned to the clusters, in order to correct for the bias. In contrast to the algorithm in Ref. [32], the correct weights are easily and rapidly calculated in our graph animal algorithm. This is its main advantage.

C. Summary

In Sec. II we present the graph animal algorithm in detail. The method used to handle graph isomorphism is briefly reviewed in Sec. III. Extensive tests, mostly with two protein interaction networks, one for *E. coli* with 230 nodes and 695 links [46], and one for yeast with 2559 nodes and 7031 links [47], are presented in Sec. IV [48]. Both networks were obtained using the tandem affinity purification (TAP) method. In particular, our algorithm involves as a free parameter a percolation probability p . For optimal performance, in lattice animals p should be near the critical value where cluster growth percolates [41]. We show how the performance for graph animals depends on p , on the subgraph size n , and on other parameters. In Sec. V we use our sampling method to study these two networks systematically. We verify that large subgraphs with high link density are overwhelmingly strong motifs, while nearly all large subgraphs with low link density are antimotifs [16,31]—although our data show much more structure than suggested by the scaling arguments of [16]. We also find striking differences in the strongest motifs for the two networks. Dominant motifs for the *E. coli* network are either bipartite or close to it (with many nodes sharing the same neighbors) while “tadpoles” with bodies consisting of (almost) complete graphs dominate for yeast. Our conclusions and discussions of open problems are given in Sec. VI.

The present work only addresses undirected networks, but the graph animal algorithm works without major changes also for directed networks. Due to the larger number of different directed subgraphs, an exhaustive study of even moderately large subgraphs is much more challenging [42].

II. ALGORITHM

In this section we explain how our algorithm achieves uniform sampling of connected subgraphs in undirected networks. The graph animal algorithm executes a generalization of the Leath algorithm for lattice animals. The observation central to the work in Refs. [40,41,44] is that the animal and percolation ensembles concern exactly the same clusters. The only difference between the two ensembles is that clusters in

the percolation ensemble have different weights, while all clusters with the same number of nodes (sites) have the same weight in the animal ensemble. We focus on site percolation [49]. Bond percolation could also be used [41], but this would be more complicated and is not discussed here.

A. Leath growth for graph animals

For regular lattices and undirected networks we use the following epidemic model for growing connected clusters of sites [40]:

(1) Choose a number $p \in [0, 1]$ and a maximal cluster size n_{\max} . Label all sites (nodes) as “unvisited.”

(2) Pick a random site (node) i_0 as a *seed* for the cluster, so that the cluster consists initially of only this site; mark it as “visited.”

(3) Do the following step recursively, until all boundary sites of the cluster have been visited, or until the cluster consists of n_{\max} sites, whichever comes first: (Note that a boundary site of a cluster C is a site which is not in C , but which is connected to C by one or more edges).

(A) Choose one of the unvisited boundary sites of the present cluster, and mark it as visited;

(B) with probability p join it to the cluster.

Once a boundary site has been visited, it cannot later join the cluster; it either joins the cluster when it is first visited (with probability p) or is permanently forbidden to join (with probability $1-p$).

The order in which the boundary (or “growth”) sites are chosen influences the efficiency of the algorithm, but this is irrelevant for the present discussion. The growth algorithm can be seen as an idealization of an epidemic process (“generalized” or SIR epidemic [51,52]) with three types of individuals (susceptible, infected, removed). Starting with a single infected individual with all others susceptible, the infected individual can infect neighbors during a finite time span. Everyone either gets infected or does not at his (her) first contact. The latter are removed, as are the infected ones after their recovery, and do not participate in the further spread of the epidemic.

Assume that for some fixed node i_0 , a connected labeled subgraph G^ℓ exists, which contains i_0 and has $n < n_{\max}$ nodes and b visited boundary nodes. The chance that precisely this particular labeled subgraph will be chosen using the algorithm is

$$P_{G^\ell}(p; i_0) \equiv P_{nb}(p; i_0) = p^{n-1}(1-p)^b. \quad (1)$$

Since an independent decision is made at each boundary site, this is indeed the probability for $n-1$ sites to be selected to join the cluster, while b sites are rejected.

Denote by $c(G^\ell)$ the indicator function for the existence of G^ℓ , i.e., $c(G^\ell)=1$ if the subgraph exists in the network, and $c(G^\ell)=0$ otherwise. Furthermore, denote by $c(G^\ell; i_0)$ the explicit indicator that G^ℓ exists and contains the node i_0 . Then the total number of occurrences of the *unlabeled* subgraph G is given by

$$c_G = n^{-1} \sum_{i=1}^N c_{G,i} = n^{-1} \sum_{i_0=1}^N \sum_{G^\ell \sim G} c(G^\ell; i_0), \quad (2)$$

where $c_{G,i}$ is the number of occurrences that contain node i , and where the last sum runs over all labeled subgraphs G^ℓ that are isomorphic to G . The factor n^{-1} takes into account that a subgraph with n nodes is counted n times.

If we repeat the epidemic process M times, always starting at the same node i_0 , then the expected number of times G^ℓ occurs is

$$\langle m(G^\ell; p, i_0) \rangle = M c(G^\ell; i_0) P_{G^\ell}(p; i_0). \quad (3)$$

Hence, an estimator for $c_{G,i}$ based on the actual counts $m(G^\ell; P, i_0)$ after M trials is

$$\hat{c}_{G,i_0}(M) = M^{-1} \sum_{G^\ell \sim G} m(G^\ell; p, i_0) [P_{G^\ell}(p; i_0)]^{-1}. \quad (4)$$

Here and in what follows carets always indicate estimators.

More generally, the starting nodes are chosen according to some probability Q_{i_0} . After $M \gg 1$ trials in total, site i_0 will have been used as a starting point on average $Q_{i_0}M$ times. This gives then the estimator for the total number of occurrences of G ,

$$\begin{aligned} \hat{c}_G(M) &= n^{-1} \sum_{i_0=1}^N \hat{c}_{G,i_0}(Q_{i_0}M) \\ &= (nM)^{-1} \sum_{i=1}^N Q_i^{-1} \sum_{G^\ell \sim G} m(G^\ell; p, i) [P_{G^\ell}(p; i)]^{-1}. \end{aligned} \quad (5)$$

It is simplest to take a uniform probability $Q_{i_0} = 1/N$. But a better alternative is to choose each node with a probability proportional to its degree, as nodes with larger degrees have more connected subgraphs attached to them. This is accomplished by choosing a link with uniform probability $1/L$, where L is the total number of links in the network, and then choosing one of the two ends of this link at random. This gives

$$Q_i = (2L)^{-1} k_i. \quad (6)$$

The algorithms of [40,44] are directly based on Eq. (5). Their main drawback is that all information from clusters that are still growing at size n is not used. Clusters whose growth had stopped at sizes $< n$ do not contribute to \hat{c}_G either, of course. Thus only those that stop growing exactly at size n are used in Eq. (5). This requires, among other things, a careful choice of p : If p is too large, too many clusters survive past size n , while in the opposite case too few reach this size at all. But even with the optimal choice of p , most of the information is wasted.

B. Improved Leath method

The major improvement comes from the following observation [41]: Assume that a cluster has grown to size n , and among the b boundary sites there are exactly g that have not yet been tested (“growth sites”). Thus growth has definitely

stopped at $b-g$ already visited boundary sites, while the growth on the remaining g boundary sites depends on future values of the random variable used to decide whether they are going to be infected. With probability $(1-p)^g$ none of them are susceptible, and the growth will stop at the present cluster size n . Thus we can replace the counts $m(G^\ell; p, i_0)$ in the estimator for c_G by the counts of “unfinished” subgraphs, provided we weigh each occurrence of a subgraph isomorphic to G with an additional weight factor $(1-p)^g$. Formally, this gives, with uniform initial link selection [Eq. (6)],

$$\hat{c}_G = \frac{2L}{nM} \sum_{i=1}^N k_i^{-1} \sum_{G^\ell \sim G} p^{1-n} (1-p)^{g-b} m_{\text{unfinished}}(G^\ell; p, i, g). \quad (7)$$

The quantity $m_{\text{unfinished}}(G^\ell; p, i, g)$ is the number of epidemics (with parameter p) that start at node i , give a labeled subgraph G^ℓ of infected nodes, and leave g unvisited boundary nodes. The factor $p^{1-n}(1-p)^{g-b}$ has a simple interpretation. In analogy to Eq. (1) it is the probability to grow a cluster with $n-1$ nodes in addition to the start node, g growth nodes, and $b-g$ blocked boundary nodes,

$$P_{\text{nb}g}(p; i_0) = p^{n-1} (1-p)^{b-g}. \quad (8)$$

Equation (7) is the number of generated clusters, reweighted with their inverse probabilities to be sampled, given they exist. It is the formula we use to estimate frequencies of occurrences of connected subgraphs in the protein interaction networks as discussed later in the text.

C. Resampling

In principle, Eq. (7) can be improved further. Reference [41] shows how to use the equivalents of Eqs. (7) and (8) for lattice animals as a starting point for a resampling scheme. For completeness, resampling for graph animals is briefly explained, even though it is not used in this work.

For each cluster that is still growing a *fitness function* is defined as

$$f_{\text{nb}g}(p) = p^{1-n} (1-p)^{-b} = [P_{\text{nb}g}(p; i_0)]^{-1} / (1-p)^g. \quad (9)$$

Clusters with too small fitness are killed, while clusters with too large fitness are cloned, with both the fitness and the weight being split evenly among the clones. The first factor in the fitness is just proportional to the weight, while the second factor takes into account that clusters with larger g have more possibilities to continue their growth, and thus should be more “valuable.” The precise form of Eq. (9) is purely heuristic, but was found to be near optimal in fairly extensive tests.

This resampling scheme was found to be essential, if one wants to sample clusters of sizes $n > 100$. In [41], the emphasis was on very large clusters (several thousand sites), and thus resampling was a necessity. Here, in contrast, we concentrate on subgraphs with ≈ 10 nodes or less, and stick to the simpler scheme without resampling. With respect to graph animals, we point out that optimal fitness thresholds for pruning and cloning depend in an irregular network on

the start node, i_0 , and have to be learned for each i_0 separately. Although a similar strategy achieves success for dealing with self-avoiding walks on random lattices [50], this is much more time consuming than for regular lattices.

D. Implementation details

For fast data access, we used several redundant data structures. The adjacency matrix was stored directly as a $N \times N$ matrix with elements 0/1 and as a list of linked pairs (i, j) , i.e., as an array of size $L \times 2$. The first is needed for fast checking of which links are present in a subgraph, while the second is the format in which the networks were downloaded from the web. Finally, for fast neighbor searches, the links were also stored in the form of linked lists. To test whether a site was visited during the growth of the present (say k th, $k = 1 \dots M$) subgraph, an array $s[i]$ of size N and type unsigned int was used, which was initiated as $s[i]=0$, $i=0, \dots, N-1$. Each time a site i was visited, we set $s[i]=k$, and $s[i] < k$ was used as indicator that this site had not been visited during the growth of the present cluster.

In Leath-type cluster growth, there are two popular variants. Untested sites in the boundary can be written either into a first-in first-out queue, or into a stack (first-in last-out queue). It was found in [41] that these two possibilities, whose efficiency is roughly the same when Eq. (5) is used, give vastly different efficiency with Eq. (7), in particular (but not only), in combination with resampling. In that case, the first-in first-out queue gives much better results, and we use this method to get the numerical results shown later.

III. SUBGRAPH CLASSIFICATION

After sampling a labeled subgraph G_ℓ , one has to find its isomorphism class G (i.e., $G_\ell \sim G$), by testing which of the representatives for isomorphism classes it can be mapped onto by permuting the node labels. State-of-the-art computer programs for comparing two graphs, such as NAUTY [30], proceed in two steps. First, some invariants are calculated such as the number of links, traces of various powers of the adjacency matrix, a sorted list of node degrees, etc. In most cases, this shows that the two graphs are not isomorphic (if any of these invariants disagree), but obviously this does not resolve all cases. When ambiguities remain, each graph is transformed into a standard form by a suitable permutation, and the standard forms are compared. The standard form is, of course, also a special invariant, so the distinction between “invariants” and “standard form” might seem arbitrary. It becomes relevant in practice, since the user of the package can specify which invariants (s)he deems relevant, while the calculation of the standard form is at the core of the algorithm and cannot be changed.

It is mostly the second step in this scheme, which is time limiting and which renders it useless for our purposes—although some invariants suggested, e.g., by NAUTY are also quite demanding in CPU time. Thus we skip the second step and only use invariants that are fast to compute. All these invariants, except for the number n of nodes and the number ℓ of links in the subgraph, are combined into a single index

I , which is intended to be a good discriminator between all nonisomorphic subgraphs with the same n and ℓ . Whenever a new subgraph is found, the triplet (n, ℓ, I) is calculated and compared to triplets that have already appeared. If the triplet appeared previously, the counter for this triplet is increased by 1; if not, a new counter is initiated and set to 1.

Since no known invariant (other than standard form) can discriminate between any two graphs, any method not using it is necessarily heuristic. Some of the invariants we used are those defined in Ref. [31]. In addition, we use invariants based on powers of the adjacency matrix and of its complement. More precisely, if A_{ij} is the adjacency matrix of a subgraph, then we define its complement by $B_{ij}=1-A_{ij}$ for $i \neq j$ and $B_{ij}=0=A_{ij}$ for $i=j$. Any trace of any product $A^{a_1}B^{b_1}A^{a_2}\dots$ is invariant, and can be computed quickly. The same is true for the number of nonzero elements of any such product, and for the sum of all its matrix elements. The index I is then either a linear combination or a product (taken modulo 2^{32}) of these invariants. The particular choices were *ad hoc* and there is no reason to believe they are optimal; hence those details are not given here.

With the indices described in [31], all undirected graphs of sizes $n \leq 8$ and all directed graphs with up to five nodes are correctly classified. In this work, a faster algorithm for counting loops is used; hence loop counting is always included, in contrast to the work of [31]. Index calculation based on matrix products is even faster but less precise: only 11 112 out of all 11 117 nonisomorphic connected graphs with $n=8$ were distinguished, and for directed graphs with $n=5$ just four graphs out of 9608 [53] were missed. For larger subgraphs we were not able to test the quality of the indices systematically, but we can cite some results for $n=9$. Using indices based on matrix products, we found 239 846 different connected subgraphs with $n=9$ in the *E. coli* protein interaction network [46] and its rewirings. Given the fact that there are only 261 080 different connected graphs with $n=9$ [53], that many of them might not appear in the *E. coli* network, and that our sampling was not exhaustive, our graph classification method failed to distinguish at most 9% of the nonisomorphic graphs—and probably many fewer.

IV. NUMERICAL TESTS OF THE SAMPLING ALGORITHM

To test the graph animal algorithm, we first sampled both $n=4$ and $n=5$ subgraphs of the *E. coli* network, as well as $n=4$ subgraphs of the yeast network. In these cases exact counts are possible, and we verified that the results from sampling agreed with results from exact enumeration within the estimated (very small) errors. To obtain these results we used crude estimates for optimal p values, namely, $p=0.11$ for *E. coli* and $p=0.03$ for yeast. For larger subgraphs more precise estimates for the optimal p are required.

Optimal values for p

When p is too small, only small clusters are regularly encountered. If p is too large, performance decreases because




FIG. 1. (Color online) Root mean square relative errors of connected subgraph counts, Eq. (10), for the yeast ($n=5$ to 8) and *E. coli* ($n=7$) networks. In most cases, clear minima indicate roughly the optimum value for p , with caveats as explained in the text. Each data point is based on 4×10^9 generated subgraphs. Smaller values of $\sigma_n(p)$ indicate that the census for subgraphs with n nodes is on average more precise.

the weight factors in Eq. (7) depend too strongly on the number of blocked boundary sites $b-g$. The latter varies from instance to instance, and this can create huge fluctuations in the weights given to individual subgraphs.

The networks we are interested in are sparse ($L/N \approx \text{const} \ll N$) and approximately scale-free [47]. As a result, most nodes have only a few links, but some “hubs” have very high degree. In fact, the degrees of the strongest hubs may diverge in the limit $N \rightarrow \infty$. For such networks it is well known that the threshold for spreading of an infinite SIR epidemic is zero [54]. On finite networks this means that one can create huge clusters even for minute p , and this tendency increases as N increases. Thus, we anticipate the optimal p to be small, and to decrease noticeably in going from the *E. coli* ($N=230$) to the yeast network ($N=2559$). This is, in fact, what we find.

As a first test, we compute the root mean square relative errors of the subgraph counts, averaged over all subgraphs of fixed size n . Let γ_n be the number of different subgraphs of size n found, and let Δc_G be the error of the count for subgraph G . These errors were estimated by dividing the set of M independent samples into bins, and estimating the fluctuations from bin to bin. Then

$$\sigma_n(p) = \left[\frac{1}{\gamma_n} \sum_{j=1}^{\gamma_n} (\Delta c_{G_j} / \hat{c}_{G_j})^2 \right]^{1/2}. \quad (10)$$

Smaller values of $\sigma_n(p)$ indicate that the subgraph census is on average more precise. Figure 1 shows results for the yeast network, with various values of p and n . Also shown are data for the *E. coli* network, for $n=7$. Each simulation used for this figure (i.e., each data point) involved $M=4 \times 10^9$ generated clusters. Our first observation is that the results for *E. coli* are much more precise than those for yeast. This is mainly due to smaller hubs ($k_{\max}^{E. coli} = 36$, while $k_{\max}^{\text{yeast}} = 141$), so that much larger p values [55] could be used. Also in all other aspects, our algorithm worked much better for the *E.*




FIG. 2. (Color online) Histograms of $wP(\ln w) = w^2P(w)$ for connected $n=8$ subgraphs of the yeast network. Each curve corresponds to one run (4×10^9 generated subgraphs) with a fixed value of p , with p increasing from the lower curves to the higher ones. Results are more reliable, the further to the left the maximum of the curve is and the faster the decrease of its tail at large w is.

coli network than for yeast. Therefore we exhibit in the rest of this section only results for yeast, implying that whenever a test was positive for yeast, an analogous test had been made for *E. coli* with at least as good results.

Even with the large sample sizes used in Fig. 1, many $n=8$ subgraphs were found only once (in which case we set $\Delta c_G / \hat{c}_G = 1$), which explains the high values of $\sigma_g(p)$. This is also why we do not show any data for $n > 8$ in Fig. 1. The relative error $\sigma_n(p)$ for each $n < 8$ shows a broad minimum as a function of p . The increase in $\sigma_n(p)$ at small p is because of the paucity of different graphs being generated. This effect grows when n increases, explaining why the minimum shifts to the right with increasing n . The increase of $\sigma_n(p)$ for large p , in contrast, comes from large fluctuations of weights for individual sampled graphs. When p is large, the factor $(1-p)^{b-g}$ in Eq. (7) can also be large, particularly in the presence of strong hubs.

Unfortunately, if a subgraph is found only once, it is impossible to decide whether or not the frequency estimate is reliable. Even for strong outliers, when the frequency estimate is far too large, the formal error estimate cannot be larger than $\Delta c_G = O(\hat{c}_G)$. This underestimates the true statistical errors and is partially responsible for the fact that the curve for $n=8$ in Fig. 1 does not increase at large p [56].

A more direct understanding of the decreasing performance at large p comes from histograms of the (logarithms of) weight factors. Such histograms, for $n=8$ subgraphs in the yeast network, are shown in Fig. 2. From the results in Sec. II,

$$w = \frac{2L}{nMk} p^{1-n} (1-p)^{g-b} \quad (11)$$

is the weight for a subgraph with n nodes, b boundary nodes, and g growth nodes. The algorithm produces reliable estimates if $P(w)$ decreases for large w faster than $1/w^2$, since averages (which are weighted by w) are then dominated by subgraphs that are well sampled. If, in contrast, $P(w)$ de-




FIG. 3. (Color online) Scatter plots of $\hat{c}_G(p=0.025)$ and $\hat{c}_G(p=0.07)$ against $\hat{c}_G(p=0.04)$ for connected $n=8$ subgraphs of the yeast network. The clustering of the data along the diagonal indicates the basic reliability of the estimates, independent of the precise choice of p . Sample sizes were 4×10^{10} for $p=0.04$, 2.4×10^{10} for $p=0.025$, and 8×10^9 for $p=0.07$. The latter two correspond to roughly the same CPU time.

creases more slowly, then the tail of the distribution dominates, and the results cannot be taken at face value [57]. We observe from Fig. 2 that the data for $n=8$ is indeed reliable for $p < 0.07$ only. The curve for $p=0.09$ in Fig. 2 also bends over at very large values of w , indicating that even for this p our estimates should finally be reliable, when the sample sizes become sufficiently large. But this would require extremely large sample sizes.

As a last test we checked whether the estimates \hat{c}_G are independent of p as they should be. Figure 3 shows the estimates obtained for $n=8$ subgraphs in the yeast network with $p=0.025$ and $p=0.07$ against those obtained with $p=0.04$. Clearly, the data cluster along the diagonal—showing that the estimates are basically correct. They scatter more when the counts are lower (i.e., in the lower left corner of the plot). The asymmetries in that region result from the fact that rarely occurring subgraphs are completely missed for $p=0.04$ and even more so for $p=0.025$, cutting off thereby the distributions at small \hat{c}_G . For larger counts, the estimates for $p=0.025$ are more precise than those for $p=0.07$. The latter show high weight “glitches” arising from the tail of $P(w)$ discussed earlier in this section.

For increasing p , the numbers m_G of generated subgraphs of type G increase of course (as the epidemic survives longer), so that average weights, defined as $\langle w_G \rangle = \hat{c}_G M / m_G$, decrease. But this decrease is not uniform for all G . Rather, it is strongest for fully connected subgraphs [$\ell = n(n-1)/2$], and is weakest for trees. For the yeast network and $n=8$, e.g., $\langle w_G \rangle$ averaged over all trees decreases by a factor ~ 18 when p increases from 0.025 to 0.085, while $\langle w_G \rangle$ averaged over all graphs with $\ell \geq 25$ decreases by a factor ~ 1700 . Smaller values of $\langle w_G \rangle$ are preferable, as they imply smaller fluctuations. Thus it would be most efficient to use larger p values for highly connected subgraphs, and smaller p for treelike subgraphs. Counting very highly connected subgraphs—where every node has a degree in the subgraph $\geq k_0$, say—is also made easier by first reducing the network to its k core

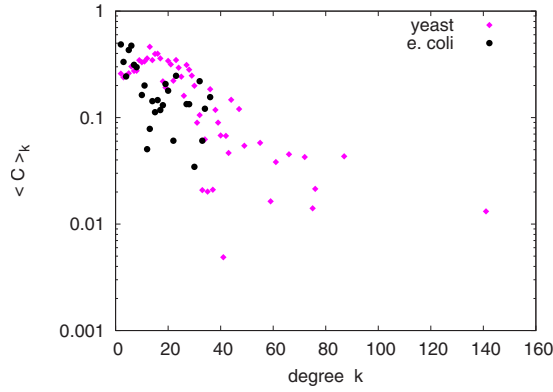


FIG. 4. (Color online) Average clustering coefficients for nodes with fixed degree k plotted versus the degree, for the giant component of the yeast and *E. coli* protein interaction networks. While the clustering coefficient decreases with k for *E. coli*, it attains a maximum at $k \approx 15$ for yeast.

with $k=k_0$, and then sampling from the latter.

V. RESULTS

A. Characterization of the networks

As already stated, both networks as we use them are fully connected [48]. The *E. coli* network has 230 nodes and 695 links, while the yeast network has 2559 nodes and 7031 links. Both networks show strong clustering, as measured by the clustering coefficients [5]

$$C_i = \frac{2}{k_i(k_i - 1)} \sum_{j < m} A_{jm}, \quad (12)$$

where k_i is the degree of node i and the sum runs over all pairs of nodes linked directly to i . In Fig. 4 we show averages of C_i over all nodes with fixed degree k . We see that $\langle C \rangle_k$ is quite large, but has a noticeably different dependence on k for the two networks. While it decreases with k for *E. coli*, it attains a maximum at $k \approx 15$ for yeast.

The unweighted average clustering $\bar{C} = N^{-1} \sum_{i=1}^N C_i$ is 0.1947 for yeast, and 0.2235 for *E. coli*. Due to the different behavior of $\langle C \rangle_k$, the ranking is reversed for the weighted averages

$$\langle C \rangle = \frac{\sum_{i=1}^N C_i k_i (k_i - 1)}{\sum_{i=1}^N k_i (k_i - 1)} = \frac{3n_{\Delta}}{3n_{\Delta} + n_{\vee}}, \quad (13)$$

where n_{Δ} is the number of fully connected triangles on the network and n_{\vee} is the number of triads with two links (see [6] for a somewhat different formula). Numerically, this gives $\langle C \rangle = 0.1948$ for yeast and 0.1552 for *E. coli*. This can be understood as a consequence of the fact that the relative frequency of fully connected triangles is higher in yeast than in *E. coli*: in yeast (*E. coli*) there are 6969 (478) triangles compared to 86 291 (7805) triads with two links.

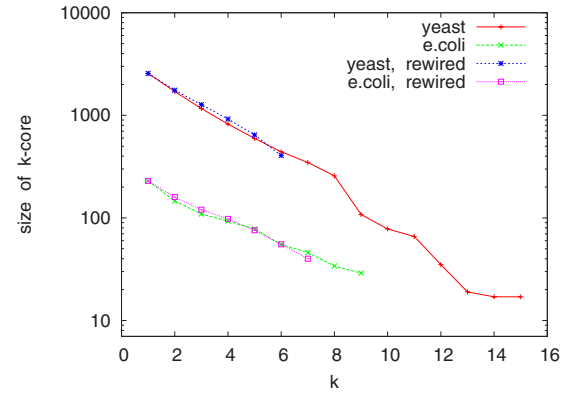


FIG. 5. (Color online) Sizes of the k cores for the two networks, plotted against k . Notice that the k cores for yeast contain a nearly fully connected cluster with 17 nodes. In addition to the core sizes for the original networks, the figure also shows average core sizes for rewired networks as discussed in Sec. V C.

Associated with this difference are distinctions between the k cores [58] of the two networks. Figure 5 shows the sizes of the k cores against k . We see that the yeast network contains nonempty cores with k up to 15. Moreover, the core with $k=15$ has exactly 17 nodes. It is a nearly fully connected subgraph with just one missing link. All 17 proteins in this core are parts of the 26S proteasome, which consists of 20 or 21 proteins [59,60]. All these proteins presumably interact very strongly with each other. When the interactions between the proteins within the 26S proteasome are taken out (the corresponding elements of the adjacency matrix are set to zero), the k core with highest k has $k=12$ and consists of 15 nodes. All its nodes correspond to proteins in the mediator complex of RNA polymerase II [59], which contains 20 proteins altogether. After eliminating all interactions between these, two 11 cores with, respectively, 13 and 14 nodes remain, the first corresponding to the 20S proteasome and the second corresponding to the RSC complex [59]. Again these particular complexes have only a few more proteins than those contained within their largest k cores, so they are very tightly bound together. All remaining complexes appear to be more loosely bound, so that much of the strong larger scale clustering in the yeast network (involving 7–10 nodes) can be traced to only a few tightly bound complexes. This has a big effect on the subgraph counts, as we shall see.

B. Trends in subgraph counts

Subgraph counts \hat{c}_G for the *Escherichia coli* and yeast networks, plotted against $n^2 + 2\ell$, are shown in Figs. 6 and 7. For large n we see a very wide range, with counts varying between 1 and $>10^8$. In general, counts decrease with an increasing number of links, i.e., trees are most frequent. This is a direct consequence of the fact that the networks are sparse. Even when n and ℓ are fixed, the counts c_G can range over six orders of magnitude (e.g., for yeast with $n=8$ and $\ell=17$).

For the yeast network, there are clear systematic trends for the counts at fixed n and ℓ . The most frequent subgraphs are those with strong heterogeneity, i.e., with a large varia-




FIG. 6. (Color online) Counts for connected subgraphs with fixed topology and with $n \leq 8$ in the *E. coli* network, plotted against $n^2 + 2\ell$. The variable $n^2 + 2\ell$ is used to spread out the data, so that the dependence on both n and ℓ (number of links) can be seen independently, without data points overlapping. For most of the points, the error bars are smaller than the sizes of the symbols.

tion of the degrees (within the subgraph) of nodes, while the most rare are those with minimal variation. Figure 8 shows the counts \hat{c}_G for $n=8$ and with four different values of ℓ plotted against the variance of the degrees of the nodes within the subgraph,

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n k_i^2 - \left[\frac{1}{n} \sum_{i=1}^n k_i \right]^2. \quad (14)$$

For all four curves we see a trend, where the count increases with σ , but hardly any trend like this is seen for the *E. coli* network (data not shown). The effect seen in the yeast data is probably related to the very strongly connected core in that network (see the last subsection). As we shall also see later in Sec. V D, subgraphs with high counts in yeast often have a tadpole form with a highly connected body (which is part of one of the densely connected complexes discussed in the last section) and a short tail attached to it. These cores may also be responsible for the main difference between Figs. 6 and 7, namely, the strong representation of very highly connected (large ℓ) subgraphs in the yeast network. Taking out all in-




FIG. 7. (Color online) Counts for subgraphs with fixed topology and with $n \leq 8$ in the yeast network, plotted against $n^2 + 2\ell$ as in Fig. 6.




FIG. 8. (Color online) Counts for $n=8$ subgraphs of the yeast network with $\ell=7, 10, 13$, and 18 , plotted against the variance of the node degrees within the subgraphs, as given by Eq. (14). Zero variance means that all nodes have exactly the same degree, whereas a higher variance indicates that the nodes differ more widely. Typically, subgraphs with more variation in their nodes (and thus with larger σ^2) have higher counts than those for which the degrees within the subgraph are more uniform.

teractions within the 26S and 20S proteasomes, within the mediator complex and within the RSC complex reduces substantially the counts for highly connected subgraphs. The count for the complete $n=7$ subgraph, e.g., is reduced in this way from $25\,164 \pm 68$ to 682 ± 23 . The removal of interactions within the 26S proteasome makes by far the biggest contribution.

C. Zipf plots

In [31] it was found that “Zipf plots” (subgraph counts versus rank) in the *E. coli* network exhibit power law behavior, whose origin is not yet understood. The essential difference between the subgraph counts in [31] and in the present paper is that we sample only connected subgraphs, while *all* subgraphs with given n were ranked in [31]. Also, noting that disconnected subgraphs are more likely to be sampled than connected ones when picking nodes at random (due to the sparsity of the networks), we can go to much higher ranks for the connected subgraphs.

Zipf plots for connected subgraphs in the *E. coli* network are shown in Fig. 9. Each curve is based on $4 \times 10^9 - 10^{10}$ generated subgraphs. Each is strongly curved, suggesting that there are no power laws—at least for subgraph sizes where we obtain reasonable statistics for the census. The curves show less curvature for larger n , but this is a gradual effect. It seems that the scaling behavior found in [31] was mainly due to the presence of disconnected graphs, although it is not immediately obvious why those should give scale-free statistics either. In addition, the right hand tails of the Zipf plots in [31] were cut off because of substantially lower statistics. In our case, apparently sharp cutoffs in the counts are observed for ranks $\approx 1.08 \times 10^4$ for $n=8$, $\approx 2.1 \times 10^5$ for $n=9$, and $\approx 2.9 \times 10^6$ for $n=10$. For $n \leq 9$ these are close to the total number of different connected subgraphs [26], suggesting that we have fairly complete statistics. For $n=10$ the




FIG. 9. (Color online) “Zipf” plots showing the counts for individual connected subgraphs with fixed n , plotted against their rank. Data are for the *E. coli* network.

cutoff is more affected by lack of statistics, but it is still within a factor of four of the upper limit.

D. Null model comparison and motifs

One of the most striking results of [31] was that most large subgraphs were either strong motifs or strong antimotifs. However, this finding was based on rather limited statistics and on a single protein interaction network. One of the purposes of the present study is to test this and other results of [31] with much higher statistics and for a larger network, the protein interaction network of yeast.

To define a motif requires a null model. We take this to be the ensemble of networks with the same degree sequence, obtained by the rewiring method. The average subgraph counts in the null ensemble are denoted as $\langle c_G^{(0)} \rangle$. In Figs. 10 and 11 we plot the ratios $c_G / \langle c_G^{(0)} \rangle$ against the variable $n^2 + 2\ell$ for each connected subgraph that was sampled both in the original graph and in at least one of the rewired graphs. The error bars, which include both statistical errors from sampling and the ensemble fluctuations of the null model




FIG. 10. (Color online) Ratios between the count estimates \hat{c}_G for connected subgraphs in the *E. coli* network, and the corresponding average counts $\langle \hat{c}_G^{(0)} \rangle$ in rewired networks. The data are plotted against $n^2 + 2\ell$, again to spread the points out conveniently. Most error bars are smaller than the symbols.




FIG. 11. (Color online) Same as Fig. 10, but for the yeast network. Notice that most data points for large n and ℓ are missing. Indeed, for $n=7$ all (!) data points with $\ell > 16$ are missing, because no such subgraphs were found in the rewired ensemble.

estimated from several hundred rewired networks, are for most points smaller than the symbols. A subgraph is a motif (antimotif), if this ratio is significantly larger (smaller) than 1. Even without rigorous criteria to estimate significance, it is clear that most densely connected subgraphs are motifs in the yeast network. The fact that trees or subgraphs with few loops tend to be antimotifs might not be so evident from Fig. 11, since the ratios for trees and treelike graphs are close to one. Thus we have to discuss significance more formally.

1. Z scores

Usually [31], the significance of a motif (or antimotif) is measured by its Z score

$$Z = \frac{c_G - \langle c_G^{(0)} \rangle}{\sigma_G^{(0)}}, \tag{15}$$

where $\sigma_G^{(0)}$ is the standard deviation of c_G within the null ensemble. A subgraph is a motif (antimotif), if $Z \gg 1$ ($Z \ll -1$).

The eight strongest motifs with $n=7$ in the *E. coli* network according to this definition are shown in Fig. 12, together with their Z values. To name the strongest motifs in the yeast network is less straightforward, since many subgraphs did not show up in any rewired network at all. Assuming for those subgraphs $\sigma_G^{(0)} = \langle c_G^{(0)} \rangle = 0$ would give $Z = \infty$.




FIG. 12. The eight strongest motifs with $n=7$ in the *E. coli* protein interaction network. These tend to be almost bipartite graphs, and many pairs of nodes are linked to the same set of neighbors. Their Z scores, in order from left to right, first then second row, are 2.9×10^4 , 932, 885, 648, 595, 532, 516, and 377. Their estimated frequencies in the original *E. coli* network are, in the same order: 20936 ± 8 , 161521 ± 63 , 8312 ± 5 , 1331 ± 2 , 838 ± 2 , 5985 ± 5 , 5165 ± 4 , and 519 ± 1 .

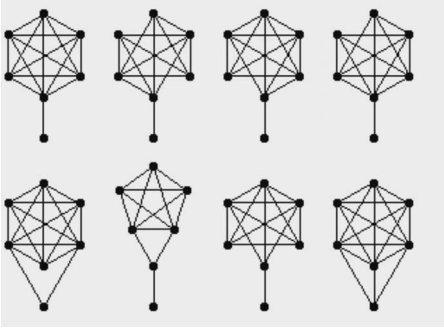


FIG. 13. Eight very strong motifs with $n=7$ for the yeast protein interaction network. These tend to be almost complete graphs with a single dangling node. Four of these graphs were not seen in any realization of the null model, so only lower bounds on their Z scores can be given. From left to right, first then second row, the estimated Z scores are $>3 \times 10^7$, 9×10^5 , $>8 \times 10^6$, 5×10^5 , $>4 \times 10^6$, 3×10^5 , 2.5×10^5 , and $>1.5 \times 10^6$. Estimated frequencies are, in the same order: $6.68(1) \times 10^5$, $9.27(5) \times 10^4$, $1.76(1) \times 10^5$, $4.84(1) \times 10^5$, $7.78(2) \times 10^4$, $3.13(6) \times 10^5$, $1.38(1) \times 10^5$, and $3.35(1) \times 10^4$.

Rough lower bounds on Z are obtained for them by assuming that $\langle c_G^{(0)} \rangle < 1/R$ and $\sigma_G^{(0)} < 1/\sqrt{R}$, where R is the number of rewired networks that were sampled, giving $Z \geq c_G \sqrt{R}$. Some of the strongest motifs in the yeast network, together with their estimated Z scores, are shown in Fig. 13. Note that no $n=7$ graphs with $\ell > 16$ were found in any of the realizations of the null model, while they were all found in the real yeast network. Hence these are all strong motifs. Those motifs in Fig. 13 for which only lower bounds for the Z score are given are the most frequent in the real network, hence they have the highest lower bound. It was pointed out in [33,61] that cliques (complete subgraphs) are in general very strong motifs. In yeast, the $n=7$ clique (with $\ell=21$) is indeed a very strong motif, but it does not have the largest lower bound on the Z score. In comparison, antimotifs have rather modest Z scores. The strongest antimotif with $n=7$ has $Z=-32.9$ ($Z=-24.7$) for *E. coli* (yeast).

With Z values up to 10^7 and more, as in Fig. 13, the motivation for using Z scores becomes suspect. On the one hand, the null model is clearly unable to describe the actual network, and has to be replaced by a more refined null model. This will be done in a future paper [42]. On the other hand, it suggests one to use instead a Z score based on *logarithms* of counts,

$$Z_{\ln} = \frac{\ln c_G - \langle \ln c_G^{(0)} \rangle}{\sigma_{\ln, G}^{(0)}}, \quad (16)$$

where $\sigma_{\ln, G}^{(0)}$ is the standard deviation of $\ln c_G^{(0)}$. An advantage of Eq. (16) would be that it suppresses $|Z|$ for motifs, but enhances $|Z|$ for antimotifs.

In general, strong yeast motifs have a tadpole structure with a complete or almost complete body, and a tail consisting of a few nodes with low degree. This agrees nicely with our previous observation that frequently occurring subgraphs in the yeast network have strong heterogeneity in the degrees




FIG. 14. (Color online) Count ratios $c_G / \langle c_G^{(0)} \rangle$ for individual subgraphs in the *E. coli* network, plotted against the count ratio for the same subgraph in the yeast network. To highlight the dependence on the number of twin nodes in the subgraph, subgraphs with $n_{\text{twin}} > 1$ ($n_{\text{twin}} > 3$) are marked by asterisks (bullets). Whereas almost all ratios are much higher in the yeast network, this is noticeably less true for subgraphs containing more than three pairs of twin nodes. These tend to fall on the diagonal indicated by the dashed line.

of their nodes. In contrast, strong *E. coli* $n=7$ motifs with not too many loops are all based on a 4-3 or 5-2 bipartite structure. When the number of loops increases, strictly bipartite structures are impossible, but the tendency towards these structures is still observed.

Whether we use Z scores or the ratio $c_G / C_G^{(0)}$ to identify motifs makes very little difference. Using either criterion, the strengths of the strongest motifs skyrocket with subgraph size. This is most dramatically apparent for the yeast network. Indeed, correlations between Z scores of individual graphs in the yeast and *E. coli* networks (data not shown) are much weaker than correlations between count ratios. The latter are shown in Fig. 14 for $n=7$ subgraphs.

2. Twinning versus clustering

Another characteristic feature of strong motifs in the *E. coli* network is the tendency for “twin” nodes. We call two nodes in a subgraph twins if they are connected to the same set of neighbors in the subgraph. Otherwise said, nodes i and k are twins, iff the i th and k th rows of the subgraph adjacency matrix are identical. Notice that twin nodes can be created most naturally by duplicating genes. We found that subgraphs with many pairs of twin nodes are, in general, also motifs in the yeast network, but they do not stand out spectacularly from the mass of other motifs. They could be the “genuine” motifs also for yeast, but only a better null model where all subgraphs actually occur with reasonable frequency would be able to prove or disprove this.

In Fig. 14 we also indicated the dependence on the number n_{twin} of pairs of twin nodes, by marking subgraphs with $n_{\text{twin}} > 3$ ($n_{\text{twin}} > 1$) by bullets (asterisks). We see that all strong motifs in *E. coli* have multiple pairs of twin nodes. These subgraphs tend to be also motifs of comparable strength in yeast—the bullets in Fig. 14 tend to cluster on the diagonal $[c_G / \langle c_G^{(0)} \rangle]_{E. coli} = [c_G / \langle c_G^{(0)} \rangle]_{\text{yeast}}$. However, there are




FIG. 15. (Color online) Counts c_G , respectively, $\langle c_G^{(0)} \rangle$ for individual subgraphs in the *Escherichia coli* network, plotted against counts for the same subgraph in the yeast network. It can be seen that the two rewired networks are much more similar (display higher correlation) than the original networks.

even stronger motifs in yeast that have no twin nodes. These graphs are typically much weaker motifs or not motifs at all in *E. coli*.

As we have already indicated, many of the strong motifs in yeast seem to be related to a few densely connected complexes such as those discussed in Sec. V A. They are either part of their cores, or they have most of their nodes in the core, with one or two extra nodes forming the tail of what looks like a tadpole. This effect is even more pronounced for $n=8$ subgraphs. For instance, the three most frequent subgraphs with $n=8$ and $\ell=17$ all contained a six clique and two nodes connected to it either in chain or in parallel. None of them occurred even in a single rewired network.

The situation is different for the *E. coli* network. There, the three most frequent graphs with eight nodes and 17 edges also have a tadpole structure, few twin nodes, and low bipartivity. But they are not very strong motifs since they occur also frequently in the rewired networks. The three strongest motifs with $n=8$ and $\ell=17$, in contrast, have many twin pairs and high bipartivity. They have slightly lower counts (by factors 2–4), but occur much more rarely in the rewired networks.

3. Effects of rewiring on differences between networks

Finally, Fig. 15 shows counts for individual subgraphs in the *E. coli* network against counts for the same subgraph in yeast. This is done for all four combinations of original and rewired networks. We see that the correlation is strongest when we compare rewired networks of *E. coli* to rewired networks of yeast. This is not surprising. It means that a lack of correlations is mostly due to special features of one network, which are not shared by the other. Rewiring eliminates most of these features. The other observation is that rewiring, in general, reduces further the counts for subgraphs, which are already rare in the original networks. This is mainly due to the fact that such subgraphs are relatively densely connected, and appear in the original networks only because of the strong clustering. This effect is more pronounced for yeast than for *E. coli*, because it is more sparse and has more densely connected clusters or complexes.

VI. DISCUSSION

In this paper we have presented an algorithm for sampling connected subgraphs uniformly from large networks. This algorithm is a generalization of algorithms for sampling lattice animals, hence we refer to it as a “graph animal algorithm” and to the connected subgraphs as “graph animals.” It allowed us to obtain high statistics estimates of subgraph censuses for two protein interaction networks. Although the graph animal algorithm worked well in both cases, the analysis of the smaller network (*E. coli*) was much easier than that of the bigger (yeast). This was not so much because of the sheer size of the latter (the yeast network has about ten times more nodes and links than the *E. coli* network), but was mainly caused by the existence of stronger hubs. Indeed, the presence of hubs places a more stringent limitation on the method than the size of the network.

One of the main results is that many subgraph frequency counts are hugely different from those in the most popular null model, which is the ensemble of networks with fixed degree sequence. Based on a comparison with this null model, most subgraphs with size ≥ 6 in both networks would be very strong motifs or antimotifs. This clearly shows that alternative null models are needed, which take clustering and other effects into account.

While this was not very surprising (hints of it had been found in previous analyses), a more surprising result is the fact that the dominant motifs in the two protein interaction networks show very different features. Most of these seem to be related to the densely connected cores of a small number of complexes in the yeast network, which have no parallels in the *E. coli* network and which strongly affect the subgraph census. Further studies are needed to disentangle these effects from other—possibly biologically more interesting—effects.

Finally, a feature with likely biological significance is the dominance of subgraphs with many twin nodes. These are nodes which share the same list of linked neighbors within the subgraph. They correspond to proteins which interact with the same set of other proteins. The most natural explanation for them is gene duplication. Connected to this is a preference for (approximately) bipartite subgraphs. These two features are very clearly seen in the *E. coli* network, much less so in yeast. But it would be premature to conclude that gene duplication was evolutionarily more important in *E. coli* than in yeast. It is more likely that its effect is just masked in the yeast network by other effects, most probably by the densely connected complexes and other clustering effects, which do not show up to the same extent in *E. coli*.

Up to now, we know very little about the biological significance of our findings. One main avenue of further work could be to relate our results on subgraph abundances in more detail to properties of the network that are associated with biological function. Another important problem is the comparison between network reconstructions, which supposedly describe the same or similar objects. There exist, e.g., a large number of published protein-protein interaction networks for yeast. Some were obtained by means of different experimental techniques, either with conventional or with high throughput methods, while others were obtained by




FIG. 16. (Color online) Count ratios $c_G/\langle c_G^{(0)} \rangle$ for individual subgraphs in the yeast networks of Refs. [63,64], plotted against counts for the same subgraph in the network of [47]. If all three networks were identical, all points should lie on the diagonal (indicated by the straight dashed line), whereas in fact systematic deviations are observed.

comprehensive literature compilations. In a preliminary step, we compared three such networks: The network obtained by Krogan *et al.* [47] that was studied above, a somewhat older network downloaded from [62] and attributed to Bu *et al.* [63], and the “high confidence” (HC) network of Batada *et al.* [64]. The latter is the most recent. It was obtained by

extracting the most reliable interactions from a vast data base, which includes the data of both Bu *et al.* and Krogan *et al.* In Fig. 16 we plot the ratios between the actual counts and the average counts in rewired networks for Bu *et al.* and for the HC data set against the analogous ratios for the Krogan *et al.* networks. If the three data sets indeed describe the same yeast network—as they purport to do, within experimental uncertainties—the points should all fall onto the diagonal. Instead, we see systematic deviations. Surprisingly, these deviations are much stronger between the Krogan *et al.* and the HC networks than between the Krogan *et al.* and the Bu *et al.* networks. Clarifying these and other systematic irregularities should give valuable insight into the strengths and weaknesses of the methods used in constructing the networks as well as their biological reliability, and should lead to improved methods for network reconstruction.

In the present paper we have only dealt with undirected networks. The basic sampling algorithm works equally well for directed networks. The main obstacle in applying our methods to the latter is the huge number of directed subgraphs, even for relatively small sizes. Nevertheless, we will present an analysis of directed networks in forthcoming work, as well as applications to other undirected networks.

ACKNOWLEDGMENT

We thank Gabriel Musso for valuable information on the yeast network.

-
- [1] C. Faloutsos, M. Faloutsos, and P. Faloutsos, *Comput. Commun. Rev.* **29**, 251 (1999).
 - [2] A.-L. Barabasi and R. Albert, *Science* **286**, 509 (1999).
 - [3] B. Bollobas, *Random Graphs* (Academic Press, London, 1985).
 - [4] M. E. J. Newman, *SIAM Rev.* **45**, 167 (2003).
 - [5] D. J. Watts and S. H. Strogatz, *Nature (London)* **393**, 440 (1998).
 - [6] M. E. J. Newman, *Phys. Rev. E* **64**, 016131 (2001).
 - [7] E. Ravasz, L. Somera, D. A. Mongru, Z. N. Oltvai, and A.-L. Barabási, *Science* **297**, 1551 (2002).
 - [8] M. E. J. Newman and M. Girvan, *Phys. Rev. E* **69**, 026113 (2004).
 - [9] E. Ziv, M. Middendorf, and C. Wiggins, *Phys. Rev. E* **71**, 046117 (2005).
 - [10] M. Rosvall and C. T. Bergstrom, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 7327 (2007).
 - [11] K. Klemm and P. F. Stadler, *Phys. Rev. E* **73**, 025101(R) (2006).
 - [12] S. Milgram, *Psychol. Today* **2**, 60 (1967).
 - [13] E. Estrada and J. A. Rodríguez-Velázquez, *Phys. Rev. E* **72**, 046105 (2005); E. Estrada, *J. Proteome Res.* **5**, 2177 (2006).
 - [14] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon, *Science* **298**, 824 (2002).
 - [15] S. Shen-Orr, R. Milo, S. Managan, and U. Alon, *Nat. Genet.* **31**, 64 (2002).
 - [16] A. Vázquez, R. Dobrin, D. Sergi, J.-P. Eckmann, Z. N. Oltvai, and A.-L. Barabasi, *Proc. Natl. Acad. Sci. U.S.A.* **101**, 17940 (2004).
 - [17] N. Kashtan, S. Itzkovitz, R. Milo, and U. Alon, *Phys. Rev. E* **70**, 031909 (2004).
 - [18] J. Besag and P. Cliffors, *Biometrika* **76**, 633 (1989).
 - [19] S. Maslov and K. Sneppen, *Science* **296**, 910 (2002).
 - [20] M. Middendorf, E. Ziv, and C. H. Wiggins, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 3192 (2005).
 - [21] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat, e-print arXiv:cs.NI/0605007v2.
 - [22] J. Park and M. E. J. Newman, *Phys. Rev. E* **68**, 026112 (2003).
 - [23] J. Park and M. E. J. Newman, *Phys. Rev. E* **70**, 066117 (2004).
 - [24] J. Foster, D. Foster, P. Grassberger, and M. Paczuski, e-print arXiv:cond-mat/0610446.
 - [25] A. Clauset, C. Moore, and M. E. J. Newman, *Proceedings of the 23rd International Conference on Machine Learning (ICML)* (IMLS/ICML, Pittsburgh, 2006).
 - [26] K. Briggs, <http://keithbriggs.info/cgt.html> (2006).
 - [27] J. U. Köbler and J. T. Schöning, *The Graph Isomorphism Problem: Its Structural Complexity* (Birkhauser, Boston, 1993).
 - [28] J.-L. Faulon, *J. Chem. Inf. Comput. Sci.* **38**, 432 (1998).
 - [29] J. Torán, *SIAM J. Comput.* **33**, 1093 (2004).
 - [30] For the “nauty” program of B. MacKay, see <http://cs.anu.edu.au/people/bdm/nauty/>

- [31] K. Baskerville and M. Paczuski, Phys. Rev. E **74**, 051903 (2006).
- [32] N. Kashtan, S. Itzkovitz, R. Milo, and U. Alon, Bioinformatics **20**, 1746 (2004).
- [33] V. Spirin and L. A. Mirny, Proc. Natl. Acad. Sci. U.S.A. **100**, 12123 (2003).
- [34] R. C. Read, Can. J. Math. **14**, 1 (1962).
- [35] I. Jensen, J. Stat. Phys. **102**, 865 (2001).
- [36] S. Redner, J. Stat. Phys. **29**, 309 (1982).
- [37] D. Stauffer, Phys. Rev. Lett. **41**, 1333 (1978).
- [38] R. Dickman and W. C. Schieve, J. Phys. (Paris) **45**, 1727 (1984).
- [39] E. J. Janse van Rensburg and N. Madras, J. Phys. A **25**, 303 (1992).
- [40] P. Leath, Phys. Rev. B **14**, 5046 (1976).
- [41] H.-P. Hsu, W. Nadler, and P. Grassberger, J. Phys. A **38**, 775 (2005).
- [42] K. Baskerville, P. Grassberger, and M. Paczuski (unpublished).
- [43] P. Grassberger, H. Frauenkron, and W. Nadler, *PERM: A Monte Carlo Strategy for Simulating Polymers and Other Things*, in “*Monte Carlo Approach to Biopolymers and Protein Folding*,” edited by P. Grassberger *et al.* (World Scientific, Singapore, 1998).
- [44] C. M. Care, Phys. Rev. E **56**, 1181 (1997); C. M. Care and R. Ettelaie, *ibid.* **62**, 1397 (2000).
- [45] S. Redner, J. Phys. A **12**, L239 (1979).
- [46] G. Butland *et al.*, Nature (London) **433**, 531 (2005); <http://www.cosin.org>
- [47] N. J. Krogan *et al.*, Nature (London) **440**, 637 (2006).
- [48] The networks given in [46,47] are not connected. In the present paper we used only their largest connected components.
- [49] D. Stauffer and A. Aharony, *An Introduction to Percolation Theory*, 2nd ed. (Taylor and Francis, London, 1994).
- [50] P. Grassberger, J. Phys. A **26**, 1023 (1993).
- [51] D. Mollison, J. R. Stat. Soc. Ser. B (Methodol.) **39**, 283 (1977).
- [52] P. Grassberger, Math. Biosci. **63**, 157 (1983).
- [53] *The On-Line Encyclopedia on Integer Sequences*, <http://www.research.att.com/~njas/sequences> (AT&T Labs, 2006).
- [54] R. Pastor-Satorras and A. Vespignani, Phys. Rev. Lett. **86**, 3200 (2001).
- [55] We might try to estimate the optimal p by the threshold for an infinite SIR epidemic on an infinite treelike network with the same degree distribution, $p_c = \langle k \rangle / \langle k^2 \rangle$ [54]. For the two networks considered in this paper, this would give $p_c(\text{yeast}) = 0.062$, $p_c(E. coli) = 0.070$, i.e., a much smaller difference in the optimal p values. One reason why this is not observed might be the very strong clustering, in particular, in the yeast data, which is neglected in this argument.
- [56] Another reason why no minimum appears in the $n=8$ curve is that we kept the number of generated clusters fixed, not CPU time. Since larger p values also imply larger clusters in average, the CPU time per cluster increases sharply for larger p .
- [57] P. Grassberger and W. Nadler, “*Go With the Winners*”—*Simulations*, in “*Computational Statistical Physics: From Billiards to Monte Carlo*,” edited by K. H. Hoffmann *et al.* (Springer, Heidelberg, 2000).
- [58] S. B. Seidman, Soc. Networks **5**, 269 (1983).
- [59] MIPS data base: <http://mips.gsf.de/genre/proj/yeast/Search/Catalogs/catalog.jsp>
- [60] SGD data base: <http://www.yeastgenome.org/cgi-bin/GO/go.pl?>
- [61] I. Ispolatov, P. L. Krapivsky, I. Mazo, and A. Yuryev, New J. Phys. **7**, 145 (2005).
- [62] <http://vlado.fmf.uni-lj.si/pub/networks/data/bio/Yeast/Yeast.htm>
- [63] D. Bu *et al.*, Nucleic Acids Res. **31**, 2443 (2003).
- [64] N. N. Batada, T. Reguly, A. Breitkreutz, L. Boucher, B. J. Breitkreutz, L. D. Hurst, and M. Tyers, PLoS Biol. **4**, e317 (2006).