Navigation in a Small World

Kevin Zhou*
MIT Department of Physics
(Dated: May 19, 2017)

Small-world networks are a class of random graphs which have both high clustering coefficients and low mean separation distance, serving as a model for complex systems such as social networks and the brain. In this paper, we investigate paths on Euclidean networks, a generalization of small-world networks characterized by an exponent δ, using scaling arguments and the renormalization group. We find that short paths exist on Euclidean networks for all $\delta < 2d$ and show that these paths can be efficiently found by a local search algorithm only when $\delta = d$. We then sketch how the structure of the small-world network increases its effective dimensionality, altering the behavior of spin systems placed on it.

I. INTRODUCTION

In 1967, the psychologist Stanley Milgram conducted a social experiment. Milgram delivered letters to strangers living in Nebraska, and asked for them to be forwarded to a stockbroker in Boston, stipulating that they could only exchange letters in person with people they already knew. When the letters arrived, it was found that on average only six exchanges were necessary to complete the journey, thus coining the term "six degrees of separation" [1, 2]. The resulting "small-world problem" was to explain the network structure that enabled such efficient communication. A solution to this problem would have applications in epidemiology [3] and the design of robust and efficient routing networks [4].

A social network with completely random edges would indeed have short separation distances, but such a network is unrealistic because friend groups come in clusters. More rigorously, one can compute a clustering coefficient for known social networks, which by far exceeds the coefficient for random graphs [5].

In 1998, Watts and Strogatz found a simple network with both a high clustering coefficient and low separation distance [6]. Their model is based on a random rewiring of a regular lattice. Starting with a cubic lattice with $N = L^d$ vertices, a fraction p of the edges are rewired so that one of their end vertices is replaced with a uniformly distributed random vertex on the lattice. Newman and Watts modified this definition so that instead of rewiring, new "shortcuts" are added to the lattice instead, ensuring the graph stays connected [7]. They then showed that the model indeed describes a small world, with the average separation distance scaling as $\ell \sim \log N$.

Shortly afterward, Kleinberg pointed out that while the average lengths of the shortest paths between vertices were small, it may be impossible to find these paths using only local information [8]. Kleinberg proposed a modified graph, called a Euclidean network, where the length distribution of the shortcuts follows a power law,

$$P(r) \sim r^{-\delta}, \tag{1}$$

and showed that paths using only local information are shortest when $\delta = d$ instead; the small-world graph is the case $\delta = 0$.

In this paper, we review the original analysis of the small-world graph's finite-size scaling properties. We then show that $\delta = d$ is optimal for local search, and additionally show by an RG analysis that for $\delta > 2d$, a Euclidean network does not have short paths. We then study Ising models defined on Euclidean networks, where the shortcut bonds act as a form of quenched disorder, modifying the critical exponents as a function of δ.

II. SHORTEST PATHS FOR $\delta = 0$

We define the mean separation ℓ to be the length of the shortest path between two vertices, averaged over all possible pairs of vertices. In a regular lattice, we have the scaling

$$\ell \sim L = N^{1/d} \tag{2}$$

while for a random graph we have

$$\ell \sim \log N \tag{3}$$

because the number of vertices reachable from a given start vertex increases roughly exponentially with the path length. In general, we expect $\ell \sim N^\theta$ and we define a graph to have the small-world property if $\theta = 0$, as in the random graph.

Numerical evidence shows that small-world graphs can have the small-world property even for low values of p [9]. However, for sufficiently small N at fixed p, the graph will always act like a regular lattice, as it is likely that no shortcuts will exist at all. Therefore the system has a characteristic length scale ξ at which this crossover behavior occurs; since p is the only parameter,

$$\xi \sim p^{-\tau}. \tag{4}$$

As long as $p \ll 1$, the lattice spacing is irrelevant and ξ is the only length scale, so ℓ obeys the finite-size scaling law

$$\ell = L f(L/\xi). \tag{5}$$

* knzhou@mit.edu

We now perform a simple position-space RG transformation by grouping pairs of vertices together, replacing all edges between vertices with edges between the corresponding pairs of vertices. Then

$$L' = L/2, \quad p' = 2^d p, \quad \ell' = \ell/2 \tag{6}$$

because the system is halved in size while the number of shortcuts stays constant. The renormalization of ℓ requires the limit $p \ll 1$, as the distance traveled along shortcuts is not halved; thus we require most edges in a route to be non-shortcuts.

Plugging our results into (5), we find $\tau = 1/d$. Intuitively, this is because ξ is the typical distance between nodes with shortcuts (as we require $L \gtrsim \xi$ for the network to have any shortcuts at all). We thus have the scaling form [7]

$$\ell = Lf(p^{1/d}L) \tag{7}$$

where we must have

$$f(x) \sim \begin{cases} (\log x)/x & x \gg 1 \\ d & x \ll 1 \end{cases} \tag{8}$$

to reproduce the random graph and regular graph results in the limits of high and lot p. In particular, this analysis shows that for *any* nonzero probability p, the small-world property is attained in the thermodynamic limit $N \to \infty$.

III. LOCAL SEARCH

Many of the shortest paths in the small-world network are unrealistic; a typical path could suddenly jump to the opposite end of the lattice, then jump back again, taking advantage of distant shortcuts. But in the Milgram experiment, we would expect participants to only have local information; for example, they might know where all of their friends are, but not the locations of any of their friends. In this case the optimal algorithm is to always reduce the Euclidean distance to the target; one should never move away from the target, by a shortcut or not, because it is impossible to know whether such a move will pay off [8]. As such, we will call these paths "naive".

There are two competing constraints on δ. For small δ, the typical shortcut length is large and paths can make fast initial progress, but once they get closer to their destination, the majority of shortcuts are useless as they point away. For large δ, shortcuts remain useful but their typical length is small.

To understand this quantitatively, we present a novel scaling argument for the typical length ℓ' of the shortest naive path between points, which includes known results in the literature as special cases [10, 11] and improves upon a known bound in $d = 2$ [8]. Getting the scaling correct is tricky, because there are many nontrivial length scales in the problem: the usefulness and typical length of a shortcut both depend on the distance r to the destination. To avoid this complication, we define $d(r)$ to be the expected number of edges required to reduce the separation from r to $r/2$.

Given a fixed r, there are two relevant length scales. The typical number of steps ξ_1 that must be taken between useful shortcuts is

$$\xi_1 \sim \frac{\int_1^L (r^{d-1}dr)/r^\delta}{\int_1^r (r^{d-1}dr)/r^\delta} \sim \left(\frac{L}{r}\right)^{d-\delta} \tag{9}$$

where we have set $p = 1$ for simplicity. We are suppressing all subleading powers; in particular, our notation implicitly means

$$x^n \sim \begin{cases} x^n & n > 0 \\ \log x & n = 0 \\ 1 & n < 0. \end{cases} \tag{10}$$

The other relevant length scale is the typical distance ξ_2 that a shortcut carries us,

$$\xi_2 \sim \frac{\int_1^r (r^d dr)/r^\delta}{\int_1^r (r^{d-1}dr)/r^\delta} \sim \frac{r^{d-\delta+1}}{r^{d-\delta}}. \tag{11}$$

From the point of view of the algorithm, there are no other relevant length scales: all the dependence on L is absorbed into ξ_1, and the dependence on r is absorbed into both. Therefore we have the simple scaling form

$$d(r) = rg(\xi_1(r)/\xi_2(r)) \sim rg(L^{d-\delta}/r^{d-\delta+1}) \tag{12}$$

In the case $\xi_1 \ll \xi_2$, most of the distance is traversed along shortcuts, and the total number of steps required is approximately $\xi_1(r/\xi_2)$. Otherwise, the shortcuts are not helpful and we require on the order of r steps. Then the scaling function g satisfies

$$g(x) \sim \begin{cases} x & x \ll 1 \\ 1 & x \sim 1. \end{cases} \tag{13}$$

Now, a typical pair of points lies a distance $r \sim L$ apart, which means the typical path length is

$$\ell' \sim \sum_{i=0}^{\log L} r_i g(L^{d-\delta}/r_i^{d-\delta+1}), \quad r_i = L/2^i. \tag{14}$$

For $\delta < d$, the crossover between the two regimes occurs at $r_c = L^{(d-\delta)/(d-\delta+1)}$. Splitting the sum at this distance, we find

$$\ell' \sim r_c + \sum_{i=0}^{r_i = r_c} \left(\frac{L}{r_i}\right)^{d-\delta} \tag{15}$$

The sum is dominated by its last element, which turns out to be proportional to r_c. By contrast, for $\delta > d$, the first term in the sum dominates, giving

$$\ell' \sim Lg(1/L^{d-\delta+1}) \sim L^{\delta-d} \tag{16}$$

for $\delta < d+1$, and L otherwise. We thus conclude that the lengths of naive paths scale as $\ell' \sim N^{\theta} = L^{d\theta}$ where

$$d\theta = \begin{cases} (d-\delta)/(d-\delta+1) & \delta < d \\ \delta - d & d < \delta < d+1 \\ 1 & d > d+1 \end{cases} \tag{17}$$

In particular, the small-world property is recovered exactly when $d = \delta$, where the distance scales roughly as a power of $\log L$.

IV. SHORTEST PATHS FOR $\delta > 0$

It is also interesting to consider the behavior of the shortest path length ℓ as a function of δ. In this case a broader range of shortcut lengths is always better, since shortcuts will overlap less, so we expect that ℓ is minimum for $\delta = 0$. Since ℓ' is an upper bound on ℓ, we also know that the small-world property continues to hold up to $\delta = d$.

The result for higher values of δ can be found by a position space RG analysis[1], as detailed in [12]. We sketch the argument here; it is similar in spirit to our argument in Sec. III. We coarse grain by combining a block of b^d vertices into one, with edges combined as in Sec. II. Then if $q(r) = 1 - p(r)$ is the probability that vertices a distance r apart are not connected, the renormalized q is

$$\widetilde{q}(r/b) = [q(r)]^{b^{2d}}. \tag{18}$$

The probability that vertices a distance r apart are connected is

$$p(r) = \frac{Cp}{r^{\delta}} \tag{19}$$

where C is chosen to maintain normalization, i.e.

$$p = S_d \int_1^L p(r) r^{d-1}\, dr. \tag{20}$$

The three equations above can then be combined to find the renormalization of p, which can be shown to be

$$\widetilde{p} = b^{y_p} p \tag{21}$$

at leading order, where

$$y_p = \begin{cases} d & \delta \le d \\ 2d - \delta & \delta > d. \end{cases} \tag{22}$$

From the RG analysis, we see that the shortcuts are irrelevant for $\delta > 2d$, yielding $\theta = 1/d$ as in a regular

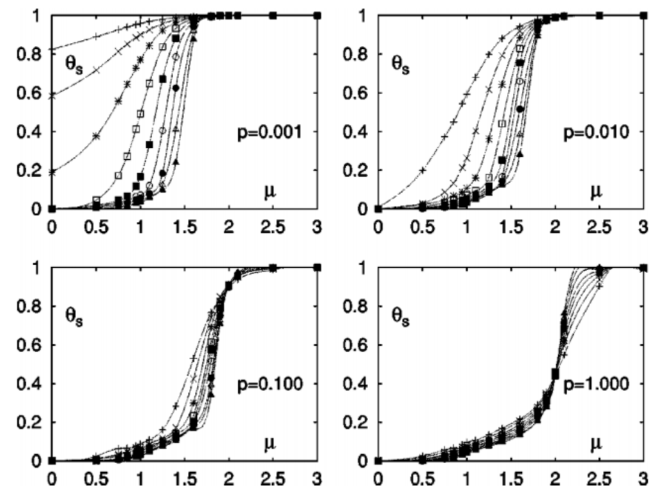

FIG. 1. Numerical computation of the exponent θ from [12] for a range of values of p, as a function of δ (here, called μ) at $d = 1$. The deviations from the theoretical predictions are due to finite size effects and decrease as larger samples are used.

lattice. Our previous section has also shown that for $\delta < d$, we must have $\theta = 0$.

The situation is more complicated for $d < \delta < 2d$; though the probability scales up under RG, the length scales up as well. Since we expect a scaling form for ℓ as in Sec. II (with a new, unknown exponent τ), θ may range between zero and $1/d$. As shown in Fig. 1, simulations indicate that θ smoothly increases from zero to $1/2d$ before discontinuously jumping to $1/d$ at $\delta = 2d$.

V. EFFECTIVE DIMENSIONS AND ISING MODELS

As shown in the previous sections, the addition of power-law distributed shortcuts affects the "effective size" of the network in a complex way. More generally, we can investigate how these shortcuts affect various measures of the dimensionality of the graph, and how this affects the dynamics of systems on these graphs.

We begin by returning to the case $\delta = 0$, following the derivation of [13]. One possible definition of the effective dimension is $N = \ell^D$, which implies

$$D = \frac{d \log N}{d \log \ell}. \tag{23}$$

This can be computed directly using our scaling form (7),

$$\frac{1}{D} = \frac{d \log \ell}{d \log L^d} = \frac{1}{d}\left(1 + \frac{d \log f(x)}{d \log x}\right). \tag{24}$$

Applying the known asymptotics of f from (8), we have

$$D = \begin{cases} d & \xi \gg L \\ d \log x & \xi \ll L. \end{cases} \tag{25}$$

[1] We attempted a similar analysis for the case of local search, but the result was more complicated due to the nontrivial dependence on r.

The effective dimension of the graph evidently depends on the length scale, and increases without bound in the thermodynamic limit.

Given the above, we might expect that spin systems on a small-world network behave as if they live on a lattice of higher dimension. For concreteness, we consider the Ising model in $d = 1$ with coupling constant J. For small p, the network is tree-like, with branching factor $B \sim 1+p$. Thus, applying the Bethe-Peierls approximation [14], one can show that

$$T_c \sim \frac{J}{|\log p|} \tag{26}$$

giving a ferromagnetic phase transition for all nonzero p.

This is a sensible conclusion, as in the thermodynamic limit, our scaling hypothesis (7) shows that the network should behave essentially as a random network; equivalently, the divergence of the effective dimension D shows that a phase transition should appear with mean-field-like characteristics. Indeed, numerical simulations have shown that the critical exponents of the phase transition match those of mean field theory [15].

We can also consider Ising models on general Euclidean networks. Again, we expect the model to behave as if it has dimension greater than d, with the difference depending on the exponent δ.

One related model for this phenomenon is the Landau-Ginzburg model with a long range interaction. As shown on a problem set [16], the addition of an interaction $1/r^{d+\sigma}$ (with $\delta = d + \sigma$) modifies the lower and upper critical dimension when $\sigma < 2$, reducing them from 2 and 4 to σ and 2σ. As a result, we expect a phase transition to appear for $\delta < 2$ and to have mean field critical exponents for $\delta < 3/2$.

Numerical simulations show that while this idea is qualitatively correct, it is not correct in detail. In the Euclidean network, the shortcuts provide strong quenched random interactions between some pairs of spins, rather than a uniform power law interaction between all pairs of spins. As shown in [17], the transition is only mean field for $\delta < 1$, and appears for all $\delta < 2$.

VI. CONCLUSION

In this paper, we have discussed and derived some of the fundamental properties of small-world and Euclidean networks, and sketched the behavior of spin systems living on them. However, there are many directions this work could be extended.

One practical direction would be to further investigate the properties of routing algorithms on such networks. Can we significantly improve our results by giving our path finding algorithm a small amount of extra information, and if so, what is it? Are the paths found in the network robust to the destruction of a small number of vertices or edges? Can we get equally good performance with a smaller number of shortcuts? Such information could be important for the design of engineered networks, such as peer-to-peer networks or the Internet as a whole.

Conversely, an important empirical question is to establish how similar networks in the real world are to our models, and more importantly, in what fundamental ways they differ. We are sure there are many more relevant features to be discovered in our small world.

[1] J. Travers and S. Milgram, Phychology Today **1**, 61 (1967).

[2] J. Travers and S. Milgram, Sociometry , 425 (1969).

[3] R. Pastor-Satorras and A. Vespignani, Physical Review E **63**, 066117 (2001).

[4] V. Latora and M. Marchiori, Physical review letters **87**, 198701 (2001).

[5] M. E. Newman, Journal of Statistical Physics **101**, 819 (2000).

[6] D. J. Watts and S. H. Strogatz, nature **393**, 440 (1998).

[7] M. E. Newman and D. J. Watts, Physics Letters A **263**, 341 (1999).

[8] J. Kleinberg, in *Proceedings of the thirty-second annual ACM symposium on Theory of computing* (ACM, 2000) pp. 163–170.

[9] M. Barthélémy and L. A. N. Amaral, Physical Review Letters **82**, 3180 (1999).

[10] A. P. de Moura, A. E. Motter, and C. Grebogi, Physical Review E **68**, 036106 (2003).

[11] H. Zhu and Z.-X. Huang, Physical Review E **70**, 036117 (2004).

[12] C. F. Moukarzel and M. A. de Menezes, Physical Review E **65**, 056709 (2002).

[13] M. E. Newman and D. J. Watts, Physical Review E **60**, 7332 (1999).

[14] S. N. Dorogovtsev, A. V. Goltsev, and J. F. Mendes, Reviews of Modern Physics **80**, 1275 (2008).

[15] C. P. Herrero, Physical Review E **65**, 066110 (2002).

[16] M. Kardar, *Statistical physics of fields* (Cambridge University Press, 2007).

[17] A. Chatterjee and P. Sen, Physical Review E **74**, 036109 (2006).