Shape and efficiency in growing spatial distribution networks

Michael T. Gastner
Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501

April 7, 2006

Abstract

We study spatial distribution networks, such as gas pipelines or train tracks, which grow from an initial source or sink of the commodity transported by the network. The efficiency depends on two properties. First, the paths to the root are ideally not much longer than the "crow flies" distance. Second, the length of all connections in the network should be low. Even though these two criteria cannot be optimized simultaneously, real networks are nevertheless nearly optimal in both respects. We propose two models to explain how this situation can arise and analyze the fractal properties of the resulting networks.

1 Introduction

A network (or graph) is a set of points or *vertices* joined together in pairs by lines or *edges*. Networks provide a useful framework for the representation and modeling of many technological, biological, and social systems, and have received a substantial amount of attention in the past few years. Reviews of recent developments in network research can be found in [1, 2, 3]. In this paper we look at the special case of networks in which the vertices occupy particular positions in geometric space. Not all networks have this property—web pages on the world wide web, for example, do not live in any particular geometric space—but many others do. Examples include transportation networks, communication networks, and power grids. Although the study of spatial networks has a history of several decades [4, 5, 6], it has only come back into the limelight during the latest surge of network research [7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23].

Here we focus on the spatial layout of man-made distribution or collection networks, such as oil and gas pipelines, sewage systems, and train or air routes. The vertices in these networks represent, for instance, households, businesses, or train stations and the edges represent pipes or tracks. The networks considered here also have a "root node," a vertex that acts as a source or sink of the commodity distributed—a sewage treatment plant, for example, or a central train station.

Geography clearly affects the efficiency of these networks. There are various possible definitions of efficiency and optimality [24, 25, 26]; in this paper we follow an idea put forward by Stevens [27]. A "good" distribution network, as we will consider it, has two definitive properties. First, the network should be efficient in the sense that the paths from each vertex to the root vertex are relatively short. That is, the sum of the lengths of the edges along the shortest path through the network should not be much longer than the "crow flies" distance between the same two vertices: if a subway track runs all around the city before getting you to the central train station, the train is probably not of much use to you. Second, the sum of the lengths of all edges in the network should be low so that the network is economical to build and maintain. In Sec. 2 we consider networks which are optimal in this second sense, but, as we will show in Sec. 3, perform poorly in the first sense. Real networks, however, manage to find solutions to the distribution problem that come remarkably close to being optimal in both senses. In Sec. 4 and 5 we suggest possible explanations for this observation in the form of two growth models for geographic networks that generate networks whose efficiency and shape are comparable to real-world examples.

2 A network growth model with minimum total length

To begin our analysis we will assume that the cost in building a distribution network is proportional to the total length of its edges. In reality the situation is certainly more complicated, but this assumption is a plausible approximation. The cheapest of all possible networks connecting all vertices together is then by definition the minimum spanning tree (MST).[1] MSTs are usually considered to be static objects where all vertices to be included in the network are known. However, if the network forms by growing outward from a root vertex as the population swells and infrastructure is extended and improved, the number of vertices in the final network cannot be predicted from the start. Instead we now introduce a dynamic version of the static MST which might be called a "growing" or "invading" minimum spanning tree [30].

We assume that at the beginning we are given the positions of the root vertex, e.g. an oil well or a central train station, and of several other points, e.g. houses or towns, that are candidates to join the network. A cluster connected to the root is built up by repeatedly adding the shortest edge that joins one unconnected vertex to another that is part of the cluster. In Fig. 1, the first few steps of this growth process are shown. Fig. 2 shows a large growing MST on a random point distribution where different colors were used to indicate when the edges were added.

The construction of this network is similar to that of invasion percolation clusters on Delaunay graphs as studied by McCarthy [31]. In his model random weights are assigned to each edge, and the edge with the smallest weight is added at each step. All of the edges in the growing MST are also edges in the Delaunay graph, but the weights on the edges are now equal to their geometric length, so the weights are, unlike in [31], neither uniformly nor independently distributed. This fact makes a rigorous mathematical treatment difficult,

[1]If we are not restricted to the specified vertex set, but are allowed to add vertices freely, then the optimal solution is the Steiner tree. The construction of a Steiner tree is an NP-hard problem [28], and hence not feasible for all but very small networks. The total length of the MST can be proved to be in the worst case only a factor $2/\sqrt{3} \approx 1.15$ more than that of the Steiner tree [29], and practically all of the results presented here remain unchanged if Steiner trees were used instead of MSTs.

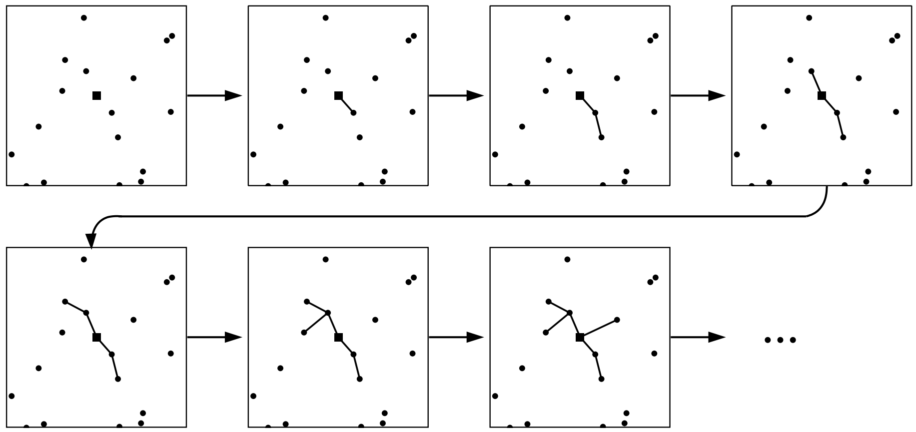

Figure 1: The first few steps during the construction of a growing minimum spanning tree. At the beginning, only the root vertex (square) is part of the tree. Then we repeatedly add the shortest edge between one connected and one unconnected vertex.

but we conjecture that the growing MST falls in the same universality class as ordinary percolation.

To support this conjecture, we have calculated the fractal dimension of the set of vertices in the tree as well as the fractal dimension of its external perimeter. The perimeter is defined by constructing the Voronoi cells for all vertices, including those that are unconnected.[2] Every cell containing a connected vertex is considered as filled, all others as empty. The external perimeter divides the filled Voronoi cells from the unbounded empty exterior, as in Fig. 3(a). A typical arrangement of filled Voronoi cells for a large network on a random point pattern is shown in Fig. 3(b). It has the typical appearance of a fractal — the interior contains numerous holes of various sizes and is bounded by a complex curve with "fjords" reaching deep into the cluster.

One of many (more or less equivalent) ways to define a fractal dimension is based on density correlations. First, we fix a number of vertices $v_1, v_2, ..., v_k$ belonging to the tree. Then we count the number of connected vertices in the tree inside a ball of radius ϵ about v_i, $N_i(\epsilon)$. If d is the fractal dimension of the point pattern, the sum of these numbers is expected to grow as ϵ^d. Hence, we can calculate d as the derivative

$$d = \frac{\mathrm{d}\log\left(\sum_{i=0}^{k} N_i(\epsilon)\right)}{\mathrm{d}\log(\epsilon)}. \tag{1}$$

The value of d obtained in this manner is the *correlation dimension* [32].

Strictly speaking, the above derivative depends on ϵ. For very small ϵ, the only point in each ϵ-ball is the center itself, hence the slope becomes zero. The same happens for large

[2]The Voronoi cell of a vertex is defined as the region in space closer to this vertex than to any of the other vertices.

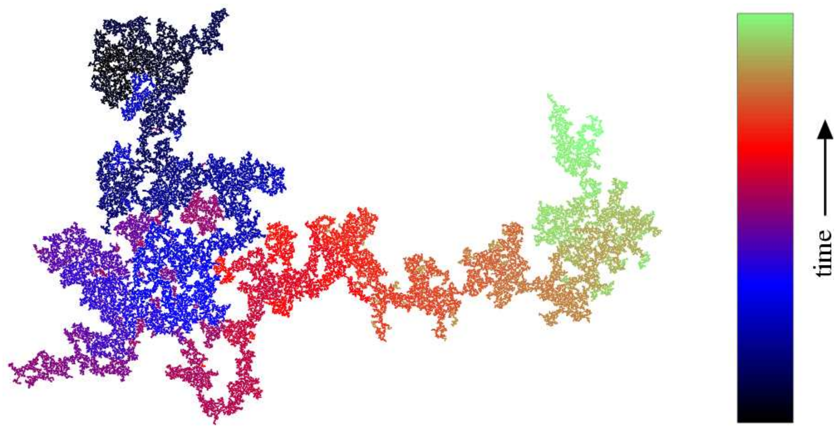

Figure 2: Growing minimum spanning tree with 100 000 vertices.

ϵ because each ϵ-ball encloses the whole tree. But for intermediate values of ϵ the slope is known to be a good estimate of the fractal dimension. We have therefore calculated d as the slope in the region of the steepest descent. Based on networks with 100 000 vertices each, we find a fractal dimension of $D_{netw} = 1.89 \pm 0.02$ which agrees with the value found for invasion percolation clusters [33].

Applying the same technique to the external perimeter, we determined its fractal dimension to be $D_{per} = 1.36 \pm 0.03$. Hence, the perimeter is not a simple one-dimensional line because it extends into all the fjords along the boundary. The fractal dimension is again consistent with that of the external accessible perimeter of invasion percolation clusters, 4/3, which supports our conjecture that growing MSTs belong to the same universality class. In Sec. 4 and 5 we will analyze how simple modifications of the presented growth model lead to quite different fractal geometries.

3 The efficiency of real networks

Let us now compare the properties of MSTs with some real-world distribution networks. We consider four examples as follows. The first network is the sewer system of the City of Bellingham, Washington. From GIS data for the city, we extracted the shapes and positions of the parcels of land (roughly households) into which the city is divided and the lines along which sewers run. We constructed a network by assigning one vertex to each parcel whose centroid was less than 100 meters from a sewer. The vertex was placed on the sewer at the point closest to the corresponding centroid and adjacent vertices along the sewers were connected by edges. The city's sewage treatment plant was used as the root vertex, for a total of 23 922 vertices including the root.

Our next two examples are networks of natural gas pipelines, the first in Western Aus-

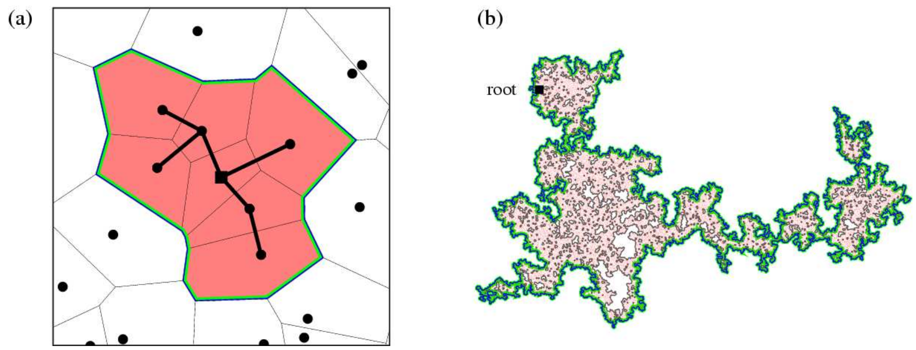

Figure 3: (a) The Voronoi cells of the tree in the last panel of Fig. 1. All cells containing a connected vertex are marked red and the perimeter is highlighted in green. (b) The interior and perimeter of the tree of Fig. 2.

tralia (WA) and the second in the southeastern part of the US state of Illinois (IL).[3] We assigned one vertex to each city, town, or power station within 10km (WA) or 10,000 feet (IL) of a pipeline. The vertex was placed on the pipeline at the point closest to each such place, and adjacent vertices were joined by edges. The root for WA was chosen to be the shore point of the pipeline leading to the Barrow Island oil fields and for IL to be the confluence of two major trunk lines near the town of Hammond, IL. The resulting networks have 226 (WA) and 490 (IL) vertices including the roots.

For our last example we take the commuter rail system operated by the Massachusetts Bay Transportation Authority in the city of Boston, MA (Fig. 4a). In this network, the 125 stations form the vertices and the tracks form the edges. In principle, there are two components to this network, one connected to Boston's North Station and the other to South Station, with no connection between the two. Since these two stations are only about one mile apart, however, we have, to simplify calculations, added an extra edge between the North and South Stations, joining the two halves of the network into a single component. The root node was placed halfway between the two stations for a total of 126 vertices in all.

In Table 1, we show the total edge lengths for each of our networks, along with the edge lengths for the MST on the same set of vertices. We find that the real-world networks, although not strictly optimal, are quite competitive with MSTs. The combined edge lengths of the real networks range from 1.12 to 1.63 times those of the corresponding MSTs. But even though the MST on the set of stations in the Boston commuter rail network, Fig. 4(b), is merely 11% shorter, it does not resemble the real network much. Looking at some of the paths in Fig. 4(a) and (b), we find that the paths to the root tend to be more circuitous in the MST than in reality. Hence, passengers would have to take quite long journeys in the MST, even from stations whose "crow flies" distance to the root is rather short. These circuitous routes would of course be annoying, and a good reason why one would not build

[3] South of 41.00°N and east of 89.85°W We consider only the largest component within this region.

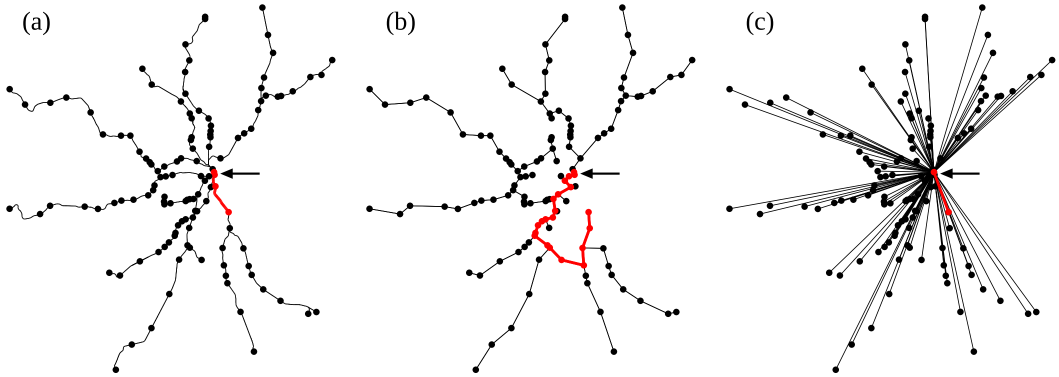

Figure 4: (a) Commuter rail network in the Boston area. (b) Minimum spanning tree. (c) Star graph. The arrow marks the assumed root of the network. Paths in the real network, like the one highlighted in red, are more direct than in the minimum spanning tree. In fact they are almost as short as the straight line connections of the star graph.

the MST, but some different network.

To make a comparison between MSTs and real networks, we consider the distance from each non-root vertex to the root first along the edges of the network and second along a simple Euclidean straight line, and calculate the mean ratio of these two distances over all such vertices. Following Ref. [34], we refer to this quantity as the network's *route factor*, and denote it q:

$$q = \frac{1}{n-1} \sum_{i=1}^{n-1} \frac{l_{i0}}{d_{i0}}, \tag{2}$$

where l_{i0} is the distance along the edges of the network from vertex i to the root (which has label 0), and d_{i0} is the direct Euclidean distance. (A similar measure "tortuosity" was used in [30].) If there is more than one path through the network to the root, we take the shortest one, but, since our examples are all nearly trees, most paths are unique. Thus, for example, $q = 2$ would imply that on average the shortest path from a vertex to the root through the network is twice as long as a direct straight-line connection.

The smallest possible value of the route factor is 1, which is achieved by the "star graph", Fig. 4(c), in which every vertex is connected directly to the root by a single straight edge. The route factors for the four real networks are shown in Table 1. As we can see, paths in the networks, although not perfectly straight, are in most cases not far away from simple straight lines, with route factors quite close to 1. Actual values range from $q = 1.13$ for the Western Australian gas pipelines to $q = 1.59$ for the sewer system. Furthermore, the route factors in all real networks are consistently better than those in the MSTs.

But now consider the column in the table, which gives the the total edge lengths for the star graphs. These figures are for all networks much larger than the optimal case and, more importantly, much poorer than the real-world networks too. Thus, although the MST is

network	n	edge length (km) actual	MST	star	route factor actual	MST	star
sewer system	23 922	498	421	102 998	1.59	2.93	1.00
gas (WA)	226	5 578	4 374	245 034	1.13	1.82	1.00
gas (IL)	490	6 547	4 009	59 595	1.48	2.42	1.00
rail	126	559	499	3 272	1.14	1.61	1.00

Table 1: Number of vertices n, total edge length, and route factor q for each of the networks described in the text, along with the equivalent results for the star graphs and minimum spanning trees on the same vertices.

optimal in terms of total edge length, it is very poor in terms of route factor, and the reverse is true for the star graph. Neither of these models would be a good general solution to the problem of building an efficient and economical distribution network. Real-world networks, on the other hand, appear to find a remarkably good compromise between the two extremes, possessing simultaneously the benefits of both the star graph and the minimum spanning tree, without any of the flaws. In the remainder of this paper we consider two mechanisms by which this might occur.

4 A network growth model with low route factor

Our first attempt to explain these observations is a modification of the growing MST of Sec. 2. Like the growing MST, this modified model will build a growing network based on a "greedy" optimization criterion that always adds the current best candidate edge. However, now the best candidate is not simply the shortest edge; instead the route factor will become part of the decision about which vertex to add next.

This is simply done by specifying a weight for each edge (i, j) thus:

$$w_{ij} = d_{ij} + \alpha \frac{d_{ij} + l_{j0}}{d_{i0}}, \tag{3}$$

where α is a non-negative independent parameter, d_{ij} is the direct Euclidean distance between vertices i and j, l_{ij} the distance along the shortest path in the network, and the root vertex has label 0. The first term in (3) is the length of the prospective edge, which represents the cost of building the corresponding pipe or track, and the second term is the contribution to the route factor from vertex i. At every step we now add to the network the edge with the global minimum value of w_{ij}. The single parameter α controls the extent to which our choice of edge depends on the route factor. For $\alpha = 0$, the network is a growing MST, which we found to give unrealistically high route factors. As α is increased from zero, however, the model becomes more and more biased in favor of making connections that give good values for the route factor.

For simplicity we will, as we already did in Sec. 2, assume that the vertices are randomly distributed in two-dimensional space with unit mean density and with one vertex randomly

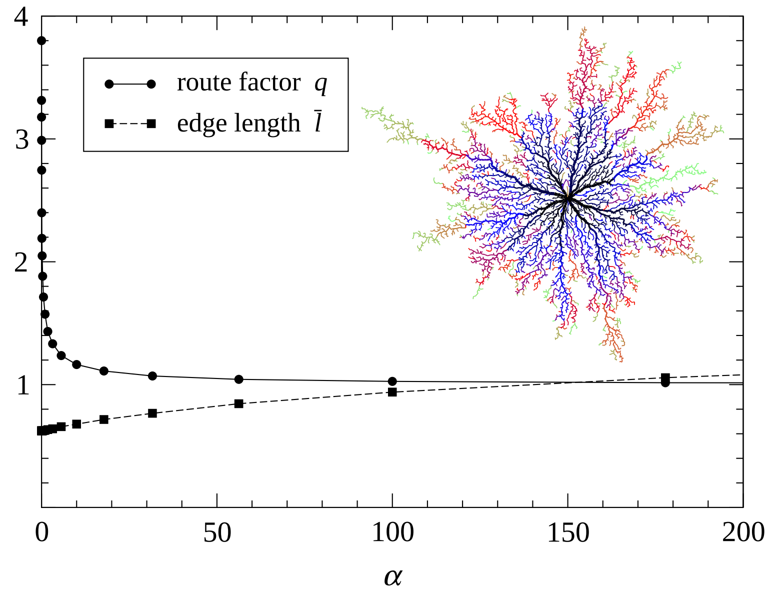

Figure 5: Simulation results for the route factor q and average edge length $\bar{l}$ as a function of α for our first model with 10 000 vertices. The length scale is normalized by setting the mean density equal to one. Inset: an example model network with $\alpha = 12.0$. Colors indicate the order in which edges were added to the network.

designated as the root of the network. The inset of Fig. 5 shows a network grown in this manner for $\alpha = 12$. The network has a dendritic appearance, with relatively straight trunk lines and short branches, bearing a superficial resemblance to diffusion-limited aggregation clusters [35], although they are based on entirely different mechanisms.

In Fig. 5 we also plot the route factor q of the network and the average length of an edge $\bar{l}$ against α. As α is increased, the route factor does indeed go down in this model, just as we expect. Furthermore, it decreases initially very sharply with α, while at the same time $\bar{l}$, which is proportional to the cost of building the network, increases only slowly. Thus, it appears to be possible to grow networks that cost only a little more than the optimal ($\alpha = 0$) network, but which have far less circuitous routes. This finding fits well with our observations of real distribution networks.

The transition from the rather amorphous shapes of a growing MST to a dendritic pattern is best studied by looking at the arrangement of filled Voronoi cells. Fig. 6 shows the patterns for different values of α. Clearly, the shape becomes more centered about the root (marked by the black square) and the holes become smaller and are pushed towards the periphery of the network as α increases. Assuming that only the filled Voronoi cells will be close enough to the network to attract human settlement, our findings for $\alpha = 2.4$ and $\alpha = 24$ are consistent with Benguigui's observation that cities are more compact around their centers than around their periphery [36]. One way to quantify how much better the

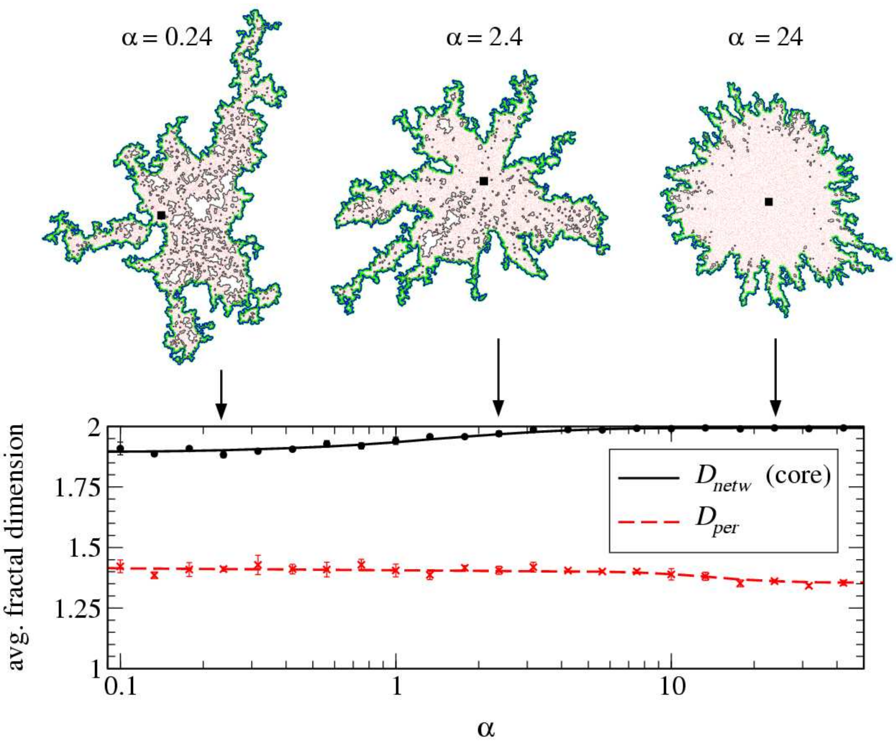

Figure 6: Top: The filled Voronoi cells of networks for three different values of α. The black square marks the position of the root. Bottom: The fractal dimension of the network D_{netw} and its perimeter D_{per} as a function of α measured in networks with 50,000 vertices each.

network fills the surrounding space is by looking at the fractal dimension. Since the holes tend to move to the periphery, the fractal dimension is not necessarily uniform everywhere, but we can obtain meaningful results by measuring the fractal dimension only for the core region. We define the core by first determining the maximum distance between the root and any point on the perimeter, and then eliminating everything farther away from the root than half that distance. The plot in Fig. 6 shows a smooth transition from a fractal dimension of 1.89 for the growing MST to the "trivial" value 2, indicating that the holes near the root gradually shrink and disappear. (This is different from the dendritic patterns of diffusion-limited aggregation whose fractal dimension is only 1.71 [37].)

The boundaries of the networks, on the other hand, remain complex geometric objects even for large α. This is reflected by the fractal dimension of the external perimeter whose value, although slightly decreasing, remains bigger than 1. Empirical measurements of city borders from satellite images yield fractal dimensions between 1.25 and 1.38 [38] similar to what we find in our model.

While this is a pleasing result, another aspect of the model is quite unrealistic. Some vertices, even ones lying quite close to the root, are joined to the network very late, because connecting them is costly in terms of the route factor. Arguably, the real world does not work this way: one does not decide to leave parts of a city without sewer service just because there is no convenient straight line for the sewer to take. Instead, connections are presumably made to those vertices that can be connected to the root by a reasonably short path, regardless of whether that path is straight. In the case of trains, for instance, people will use a train service, and thereby justify its construction, if their train journey is short in absolute terms, and are less likely to take a longer journey even if the longer one is along a straight line. As we now show, we can, by incorporating these considerations, produce a different model that still generates highly efficient networks.

5 A network growth model with short connections to the root

Let us modify Eq. (3) to give preference to short paths regardless of their shape. To do this, we write the weight of a new edge (i, j) as simply

$$w'_{ij} = d_{ij} + \beta l_{j0}. \tag{4}$$

A model with a similar weight function was studied previously by Fabrikant *et al.* [7]. However, unlike there, we assume here that the vertex positions are specified from the outset, rather than being added to the network one by one. This corresponds to a situation where sites available for settlement are determined from the beginning, which we feel is appropriate for a model of urban networks since city ordinances typically determine potential sites decades before they are finally developed.

Note that, unlike in Eq. (3), there is now no explicit term in Eq. (4) that guarantees low route factors. Nonetheless, the model self-organizes to a state whose route factor is small. Figure 7 shows results from simulations of this second model. As the plot shows, the results are qualitatively quite similar to our first model: the high value of q seen for $\beta = 0$ drops off quickly as β is increased, while the mean edge length increases only slowly. Thus we can again choose a value for β that gives behavior comparable with our real-world networks, having simultaneously low route factor and low total cost of building the network. Values of q in the range 1.1 to 1.6 observed in the real-world networks are easily achieved.

When we look at the shape of the network itself, however, we get quite a different impression (see Fig. 8). This model produces more symmetric networks that fill space out to some approximately constant radius from the root. The second term in Eq. (4) makes it economically disadvantageous to build connections to outlying areas before closer areas have been connected. Thus, all vertices within a given distance of the root are served by the network, without gaps, increasing the fractal dimension of the network quickly from 1.89 to 2 already for small β. However, unlike the dendritic shapes of Fig. 5, the perimeter now becomes more circular and, hence, in the limit of large β, a one-dimensional object.

Radial growth in fact may be the secret of how low route factors are achieved in reality. Our second model, unlike our first, does not explicitly aim to optimize the route factor, but it does a creditable job nonetheless, precisely because it fills space radially. The main trunk

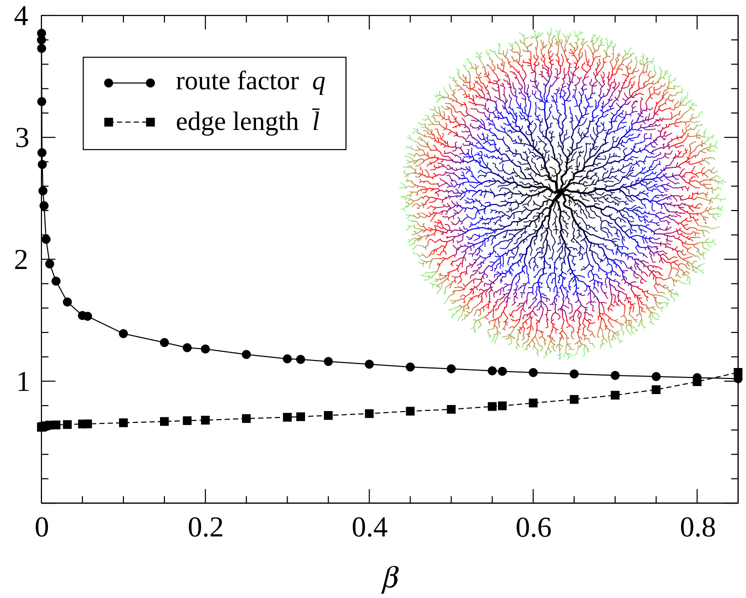

Figure 7: Route factor q and average edge length $\bar{l}$ as a function of β for our second model and 10 000 vertices. Inset: an example model network with $\beta = 0.4$.

lines in the network are forced to be approximately straight simply because the space to either side of them has already been filled and there is nowhere else to go but outwards.

How can this be reconciled with the observation that real networks, and the towns they serve, *are* dendritic in form? One might argue that it is primarily a consequence of other factors, such as ribbon development along rivers or highways. In other words, the initial distribution of vertices in real networks is usually non-uniform, unlike our model. It is interesting to see therefore what happens if we apply our model to a realistic scatter of points, and in Fig. 4b we have done this for the stations of the Boston rail system. The figure shows the network generated by our second model with $\beta = 0.4$ acting on the real-world positions of the stations. The result is, with only a couple of exceptions, identical to the true rail network, with a comparable route factor of 1.11 and total edge length 511km. This is a nontrivial result: our first model, for example, does not reproduce the true network nearly so well for any value of α.

6 Conclusion

In this paper, we have studied models of growing spatial distribution networks and compared our results with empirical data. We presented a growing version of the minimum spanning tree which keeps the cost in terms of edge length at a minimum. However, in terms of the network distance between vertices, as measured by the so-called route factor, this model is rather poor. Generally, short network length and small route factor are at odds with one

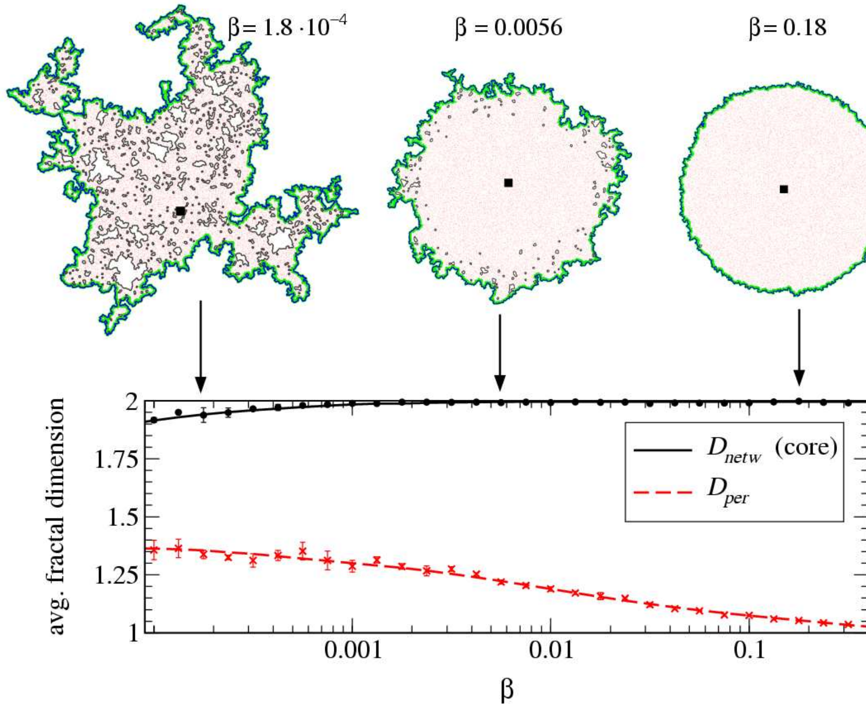

Figure 8: Top: The filled Voronoi cells of networks for three different values of β and $50,000$ vertices. Bottom: Fractal dimensions.

another, the first normally being decreased only at the expense of an increase in the second. Nonetheless, analyzing several real spatial distribution and collection networks, we found that the real world finds good compromise solutions giving nearly optimal values of both network length and route factor.

We have presented two models of growing networks based on greedy optimization strategies that show how this might occur. The first model produces networks of dendritic shape, the second model leads to radial growth. A comparison between fractal dimensions in model and reality gives some credibility to the first model. The second model, however, appears more plausible from first principle; it also is more successful at reproducing actual network structures on real point distributions.

Our focus was primarily on man-made networks, but it is plausible that our arguments are applicable to biological networks as well, such as the circulatory system [39, 40] or fungal mycels [41]. A more careful investigation, however, will be left for future research.

7 Acknowledgments

Helpful discussions with M. E. J. Newman are acknowledged.

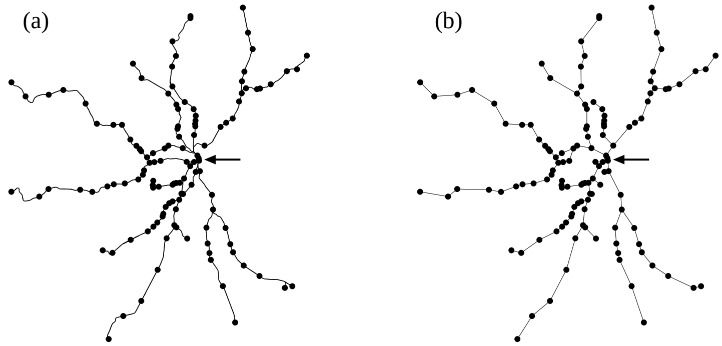

Figure 9: (a) Commuter rail network in the Boston area. (b) The model of Eq. (4) applied to the same set of stations. The arrow marks the assumed root of the network.

References

[1] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, p. 47, 2002.

[2] S. N. Dorogovtsev and J. F. F. Mendes, "Evolution of networks," *Advances In Physics*, vol. 51, p. 1079, 2002.

[3] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, p. 167, 2003.

[4] W. L. Garrison, "Connectivity of the Interstate Highway systerm," *Papers and proceesings of the Regional Science Association*, vol. 6, pp. 121–137, 1960.

[5] K. J. Kansky, *Structure of transportation networks: Relationships between network geometry and regional characteristics.* University of Chicago: Department of Geography, 1963.

[6] P. Haggett and R. J. Chorley, *Network analysis in geography.* New York: St. Martin's Press, 1969.

[7] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou, "Heuristically optimized tradeoffs: A new paradigm for power laws in the internet," in *ICALP*, vol. 2380 of *Lecture Notes in Computer Science*, pp. 110–112, Springer, 2002.

[8] S. S. Manna and P. Sen, "Modulated scale-free network in Euclidean space," *Physical Review E*, vol. 66, p. 066114, 2002.

[9] A. F. Rozenfeld, R. Cohen, D. ben Avraham, and S. Havlin, "Scale-free networks on lattices," *Physical Review Letters*, vol. 89, p. 218701, 2002.

[10] P. Sen, K. Banerjee, and T. Biswas, "Phase transitions in a network with a range-dependent connection probability," *Physical Review E*, vol. 66, p. 037102, 2002.

[11] C. P. Warren, L. Sander, and I. M. Sokolov, "Geography in a scale-free network model," *Physical Review E*, vol. 66, p. 056105, 2002.

[12] R. Xulvi-Brunet and I. M. Sokolov, "Evolving networks with disadvantaged long-range connections," *Physical Review E*, vol. 66, p. 026118, 2002.

[13] S.-H. Yook, H. Jeong, and A.-L. Barabási, "Modeling the internet's large-scale topology," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, pp. 13382–13386, 2002.

[14] M. Barthélemy, "Crossover from scale-free to spatial networks," *Europhysics Letters*, vol. 63, pp. 915–921, 2003.

[15] P. Sen and S. S. Manna, "Clustering properties of a generalized critical Euclidean network," *Physical Review E*, vol. 68, p. 026104, 2003.

[16] J. I. Alvarez-Hamelin and N. Schabanel, "An internet graph model based on trade-off optimization," *European Physical Journal B*, vol. 38, pp. 231–237, 2004.

[17] M. Kaiser and C. C. Hilgetag, "Spatial growth of real-world networks," *Physical Review E*, vol. 69, p. 036103, 2004.

[18] A. Barrat, M. Barthélemy, and A. Vespignani, "The effects of spatial constraints on the evolution of weighted complex networks," *Journal of Statistical Mechanics*, p. P05003, 2005.

[19] T. Petermann and P. De Los Rios, "Spatial small-world networks: A wiring-cost perspective." Preprint cond-mat/0501420, 2005.

[20] M. Barthélemy and A. Flammini, "Optimal traffic networks." Preprint physics/0601203, 2006.

[21] A. D. Flaxman, A. M. Frieze, and J. Vera, "A geometric preferential attachment model of networks." Preprint http://www.math.cmu.edu/ af1p/GeoWeb.pdf, 2006.

[22] P. Sen, A. K. Chandra, K. B. Hajra, and P. K. Das, "Time evolution of link length distribution in PRL collaboration network." Preprint physics/0511181, 2006.

[23] M. T. Gastner and M. E. Newman, "The spatial structure of networks," *European Physical Journal B*, vol. 49, pp. 247–252, 2006.

[24] L. B. Leopold, "Trees and streams: The efficiency of branching patterns," *Journal of Theoretical Biology*, vol. 31, pp. 339–354, 1971.

[25] T. A. McMahon and R. E. Kronauer, "Tree structures: Deducing the principle of mechanical design," *Journal of theoretical biology*, vol. 59, pp. 443–466, 1976.

[26] M. Zamir, "Optimality principles in arterial branching," *Journal of Theoretical Biology*, vol. 62, pp. 227–251, 1976.

[27] P. S. Stevens, *Patterns in nature*. Boston: Little, Brown and Company, 1974.

[28] M. R. Garey, R. L. Graham, and D. S. Johnson, "Some NP-complete geometric problems," in *Proc. 8th Annual ACM Symp. on Theory of Computing*, pp. 10–22, 1976.

[29] F. K. H. Ding-Zhu Du, "A proof of the Gilbert-Pollak conjecture on the Steiner ratio," *Algorithmica*, vol. 7, pp. 121–135, 1992.

[30] A. Kansal and S. Torquato, "Globally and locally minimal weight spanning tree networks," *Physica A*, vol. 301, pp. 601–619, 2001.

[31] J. F. McCarthy, "Invasion percolation on a random lattice," *Journal of Physics A*, vol. 20, pp. 3465–3469, 1987.

[32] S. H. Strogatz, *Nonlinear Dynamics and Chaos*. Cambridge, MA: Westview Press, 1994.

[33] D. Stauffer and A. Aharony, *Introduction to percolation theory*. London: Taylor & Francis, 1994.

[34] W. R. Black, *Transportation: A Geographical Analysis*. New York, NY: Guilford Press, 2003.

[35] T. A. Witten and L. M. Sander, "Diffusion-limited aggregation, a kinetic critical phenomenon," *Physical Review Letters*, vol. 47, pp. 1400–1403, 1981.

[36] L. Benguigui, "A new aggregation model. application to town growth," *Physica A*, vol. 219, pp. 13–26, 1995.

[37] T. A. Witten and L. M. Sander, "Diffusion-limited aggregation," *Physical Review B*, vol. 27, pp. 5686–5697, 1983.

[38] http://www.uni-konstanz.de/zwn/winterschool2004/CONTR/Villamil_HAND.pdf.

[39] C. D. Murray, "The physiological principle of minimum work. I. The vascular system and the cost of blood volume," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 12, pp. 207–214, 1926.

[40] F. Nekka, S. Kyriacos, C. Kerrigan, and L. Cartilier *Bull. Math. Biol.*, vol. 58, pp. 409–426, 1996.

[41] D. Moore, L. J. McNulty, and A. Meskauskas, *Branching Morphogenesis*. Georgetown: Landes Bioscience/Eurekah.com, 2005.