IOPscience iopscience.iop.org

Home Search Collections Journals About Contact us My IOPscience

Shape and efficiency in spatial distribution networks

This content has been downloaded from IOPscience. Please scroll down to see the full text.

J. Stat. Mech. (2006) P01015

(http://iopscience.iop.org/1742-5468/2006/01/P01015)

View the table of contents for this issue, or go to the journal homepage for more

Download details:

IP Address: 150.146.205.185
This content was downloaded on 10/07/2014 at 12:18

Please note that terms and conditions apply.

Journal of Statistical Mechanics: Theory and Experiment
An IOP and SISSA journal

Shape and efficiency in spatial distribution networks

Michael T Gastner[1,2] **and M E J Newman**[1]

[1] Department of Physics, University of Michigan, 450 Church Street, Ann Arbor, MI 48109-1040, USA
[2] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA
E-mail: mgastner@santafe.edu and mejn@umich.edu

Received 17 October 2005
Accepted 9 January 2006
Published 30 January 2006

Online at stacks.iop.org/JSTAT/2006/P01015
doi:10.1088/1742-5468/2006/01/P01015

Abstract. We study spatial networks that are designed to distribute or collect a commodity, such as gas pipelines or train tracks. We focus on the cost of a network, as represented by the total length of all its edges, and its efficiency in terms of the directness of routes from point to point. Using data for several real-world examples, we find that distribution networks appear remarkably close to optimal where both these properties are concerned. We propose two models of network growth that offer explanations of how this situation might arise.

Keywords: dendritic growth (theory), growth processes, communication, supply and information networks

©2006 IOP Publishing Ltd 1742-5468/06/P01015+9$30.00

Contents

1. Introduction

A network is a set of points or *vertices* joined together in pairs by lines or *edges*. Networks provide a useful framework for the representation and modelling of many physical, biological, and social systems, and have received a substantial amount of attention in the recent physics literature [1]–[3]. In this paper we study 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. The study of spatial networks has a long history in the sciences [4]–[6], and has in particular attracted considerable attention from physicists in recent years [7]–[12].

In this paper we study 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. In most cases the network also has 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 [13]–[15]; in this paper we follow an idea put forward by Stevens [16]. 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 be not 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. These two criteria are often at odds with each other, but, as we show in section 2, real networks nonetheless manage to find solutions to the distribution problem that come remarkably close to being optimal in both senses. We suggest possible explanations for this observation in the form of two models for geographic networks in section 3 that generate networks of comparable efficiency to

J. Stat. Mech. (2006) P01015

our real-world examples. The models we propose are based on growth processes in which edges are added one by one to connect new vertices to the network. This contrasts with the approach taken in [17, 12], in which networks are created by globally optimizing the arrangement of edges joining a given set of vertices. While both growing networks and globally optimized ones are suitable as models of spatial networks in appropriate circumstances, our focus in this paper is on the growing case.

2. Efficiency in real networks

We begin our study by looking at the properties of some real-world distribution networks. We consider four examples as follows.

Our 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 m 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 Australia (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 10 km (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 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 (figure 1(a)). 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 the 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.

We wish to quantify the efficiency of these networks in terms of path lengths and combined edge length, as described above. To do this, we compare our measurements of the networks to two theoretical models that are each optimal by one of these two criteria. If one is interested solely in short, efficient paths to the root vertex then the optimal network is the 'star graph', in which every vertex is connected directly to the root by a single straight edge (see figure 1(b)). Conversely, if one is interested solely in minimizing total edge length, then the optimal network is the minimum spanning tree (MST) (see figure 1(c)). (Given a set of n vertices at specified points on a flat plane, the MST is the

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

J. Stat. Mech. (2006) P01015

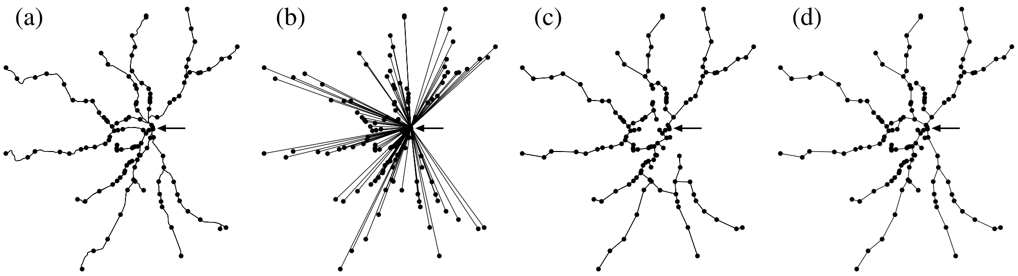

Figure 1. (a) Commuter rail network in the Boston area. The arrow marks the assumed root of the network. (b) Star graph. (c) Minimum spanning tree. (d) The model of equation (3) applied to the same set of stations.

Table 1. Number of vertices n, route factor q, and total edge length 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. (Note that the route factor for the star graph is always 1 and so has been omitted from the table.)

		Route factor		Edge length (km)		
Network	n	Actual	MST	Actual	MST	Star
Sewer system	23 922	1.59	2.93	498	421	102 998
Gas (WA)	226	1.13	1.82	5578	4374	245 034
Gas (IL)	490	1.48	2.42	6547	4009	59 595
Rail	126	1.14	1.61	559	499	3 272

set of $n-1$ edges joining them such that all vertices belong to a single component and the sum of the lengths of the edges is minimized[4].)

To make the comparison with the star graph, 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 [18], we refer to this quantity as the network's *route factor*, and denote it q:

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

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. If there is more than one path through the network to the root, we take the shortest one. 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.

The route factors for our four networks are shown in table 1. As we can see, the networks are remarkably efficient in this sense, with route factors quite close to 1. Values

[4] 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; in practice we find that there is little difference between results for minimum spanning and Steiner trees in the present context.

J. Stat. Mech. (2006) P01015

range from $q = 1.13$ for the Western Australian gas pipelines to $q = 1.59$ for the sewer system.

We also show in table 1 the total edge lengths for each of our networks, along with the edge lengths for the MST on the same set of vertices and, as the table shows, we again find that our real-world networks are competitive with the optimal model, the combined edge lengths of the real networks ranging from 1.12 to 1.63 times those of the corresponding MSTs.

But now consider the remaining two columns in the table, which give the route factors for the MSTs and the total edge lengths for the star graphs. As the table shows, these figures are for all networks much poorer than the optimal case and, more importantly, much poorer than the real-world networks too. Thus, although the MST is 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[5]. Neither of these model networks 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, but without the drawbacks. In the remainder of the paper we consider mechanisms by which this might occur.

3. Network models with low route factors

The networks we are dealing with are not, by and large, designed from the outset for global optimality (or near-optimality) of either their total edge length or their route factors. Instead, they form by growing outward from the root, as the population they serve swells and infrastructure is extended and improved. To explore this process we consider a situation in which the positions of vertices (houses, towns, etc) are given and we are to build a network connecting them. For simplicity we will assume that the vertices are randomly distributed in two-dimensional space with unit mean density, with one vertex designated as the root of the network. A cluster connected to the root is built up by repeatedly adding an edge that joins one unconnected vertex i to another j that is part of the cluster. The question is how these edges are to be chosen. Our proposal is to use a simple 'greedy' optimization criterion [19] that always adds the current best-choice edge.

We specify a weight for each edge (i, j) thus:

$$w_{ij} = d_{ij} + \alpha \frac{d_{ij} + l_{j0}}{d_{i0}}, \tag{2}$$

where α is a non-negative independent parameter. As before, d_{ij} is the direct Euclidean distance between vertices i and j and l_{ij} the distance along the shortest path in the network. The first term in (2) 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

[5] It is crucial to remember that route factors are measured relative to the value $q = 1$, not $q = 0$, since by definition no network has q less than 1. Thus, while upon initial inspection it may appear that the values 1.13 for our gas pipeline and 1.82 for the corresponding MST are of the same order of magnitude, in fact, when compared with the fundamental value of 1 the former is very significantly better than the latter. Similar observations apply to the other networks as well.

J. Stat. Mech. (2006) P01015

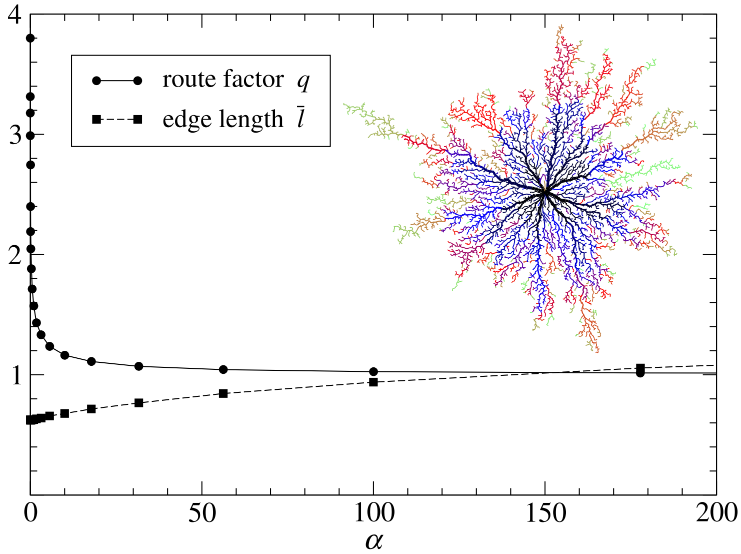

Figure 2. Simulation results for the route factor q and average edge length $\bar{l}$ as a function of α for our first model with $n = 10\,000$ vertices. The length scale is normalized by setting the mean density equal to one. Inset: an example model network with $\alpha = 12.0$. Colours indicate the order in which edges were added to the network.

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 $\alpha = 0$ we always add the vertex that is closest to the connected cluster. This limit produces a graph akin to a grown version of the minimum spanning tree, and we find it to give very poor route factors. As α is increased from zero, however, the model becomes more and more biased in favour of making connections that give good values for the route factor.

Figure 2 shows results from simulations of this model. We plot the route factor q of the entire network and the average length of an edge $\bar{l}$ against α. As α is increased the route factor does indeed decrease in this model, just as we expect. Furthermore, it initially decreases very sharply with α, while at the same time $\bar{l}$, which is a measure of 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 inset to figure 2 shows an example network grown using this model. The network has a dendritic appearance, with relatively straight trunk lines and short branches, and bears a qualitative resemblance to diffusion-limited aggregation clusters [20] or dielectric breakdown patterns [21], which have also been used as models of urban growth [22] although they are based on entirely different mechanisms.

In some respects, however, this model is quite unrealistic. In particular, many vertices are never joined to the network, even ones lying quite close to the root, because to do so would simply be too costly in terms of the route factor. (This is the reason for the dendritic shape.) This is not the way the real world works: one does not decide not to provide a sewer service to some parts of a city just because there is no convenient straight

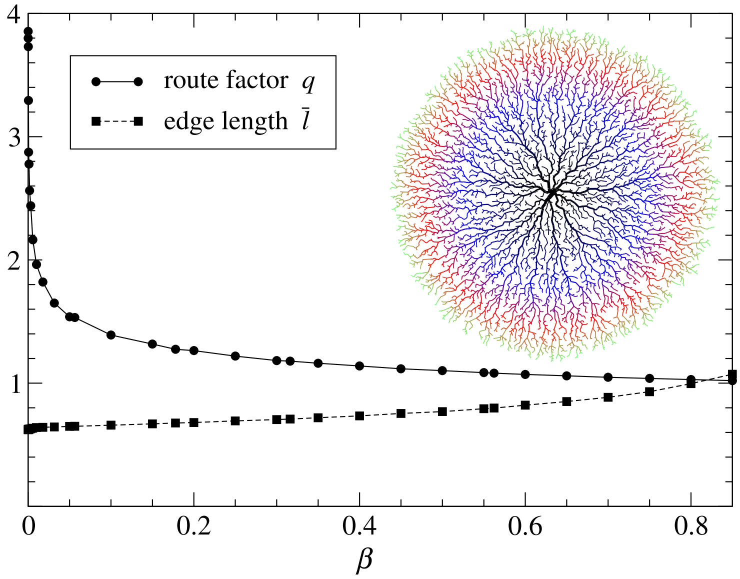

Figure 3. Route factor q and average edge length $\bar{l}$ as a function of β for our second model with $n = 10\,000$ vertices. Inset: an example model network with $\beta = 0.4$.

line for the sewer to take. Instead, connections seem to be 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 straighter line. As we now show, we can, by incorporating these considerations, produce a more realistic model that still generates highly efficient networks.

Let us modify equation (2) to give preference to short paths regardless of shape. To do this, we write the weight of a new edge (i, j) as simply

$$w'_{ij} = d_{ij} + \beta l_{j0}. \tag{3}$$

(A model with a similar weight function was studied previously by Fabrikant *et al* [23], but gives quite different results from ours because vertices were added to the network at the same time as the edges that connect them, where in our case all vertices are present from the outset. Our model and the model of Fabrikant *et al* can be thought of as two limiting cases, in which the positions of the vertices are initially entirely known or entirely unknown, respectively. Of course, the real world lies somewhere between these limits but, as is often the case, studies of the extremes can nonetheless be illuminating.) Note that there is now no explicit term that guarantees low route factors. Even so, the model self-organizes to a state whose route factor is small. Figure 3 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 behaviour 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 1.1–1.6 observed in the real-world networks are easily achieved.

When we look at the shape of the network itself however (see the inset of figure 3), we get quite a different impression. This model produces a symmetric network that fills space out to some approximately constant radius from the root, not unlike the clusters produced by the well known Eden growth model [24]. The second term in equation (3) 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, which is a more realistic situation than the dendritic network of figure 2.

And this 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 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.

Readers familiar with urban geography may argue that real networks, and the towns they serve, *are* dendritic in form. And this is true, but it is primarily a consequence of other factors, such as ribbon development along highways. In other words, the initial distribution of vertices in real networks is usually non-uniform, unlike our model. It is interesting, therefore, to see what happens if we apply our model to a realistic scatter of points, and in figure 1(d) 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 few exceptions, identical to the true rail network: only 13 out of the 125 edges present in the real network are absent in the model. Moreover, both the route factor of 1.11 and total edge length of 511 km for the model network are close to the values of 1.14 and 559 km for the real network. This is a nontrivial result: our first model, for example, does not reproduce the true network nearly so well.

4. Conclusion

To summarize, in this paper we have studied spatial distribution or collection networks such as pipelines and sewers, focusing particularly on their cost in terms of total edge length and their efficiency in terms of the network distance between vertices, as measured by the so-called route factor. While these two quantities are, to some extent, at odds with each other, the first normally being decreased only at the expense of an increase in the second, our empirical observations indicate that real-world networks find good compromise solutions giving nearly optimal values of both. We have presented two models of spatial networks based on greedy optimization strategies that reproduce this behaviour well, showing how networks possessing simultaneously good route factors and low total edge length can be generated by plausible growth mechanisms.

The results presented represent only a fraction of the possibilities in this area. Numerous other real-world networks fall into the class studied here, including utility, transportation, and shipping networks, as well as some biological networks, such as the circulatory system [25], fungal mycels [26], and others, and we hope that researchers will feel encouraged to investigate these interesting systems.

J. Stat. Mech. (2006) P01015

Acknowledgments

The authors thank Jonathan Goodwin and Sean Doherty for the pipeline network data and the staff of the University of Michigan's Numeric and Spatial Data Services for their help. This work was funded in part by the National Science Foundation under grant number DMS-0234188 and by the James S McDonnell Foundation.

References

[1] Albert R and Barabási A-L, 2002 *Rev. Mod. Phys.* **74** 47
[2] Dorogovtsev S N and Mendes J F F, 2002 *Adv. Phys.* **51** 1079
[3] Newman M E J, 2003 *SIAM Rev.* **45** 167
[4] Murray C D, 1926 *Proc. Nat. Acad. Sci.* **12** 207
[5] Horton R E, 1945 *Bull. Geol. Soc. Am.* **56** 275
[6] Edelstein-Keshet L and Ermentrout B, 1989 *SIAM J. Appl. Math.* **49** 1136
[7] Yook S H, Jeong H and Barabási A-L, 2001 *Proc. Nat. Acad. Sci.* **99** 13382
[8] Gorman S P and Kulkarni R, 2004 *Environment and Planning* B **31** 273
[9] Guimerà R, Mossa S, Turtschi A and Amaral L A N, 2005 *Proc. Nat. Acad. Sci.* **102** 7794
[10] Csányi G and Szendrői B, 2004 *Phys. Rev.* E **70** 016122
[11] Kaiser M and Hilgetag C C, 2004 *Phys. Rev.* E **69** 036103
[12] Gastner M T and Newman M E J, 2006 *Eur. Phys. J.* B at press
[13] Leopold L B, 1971 *J. Theor. Biol.* **31** 339
[14] McMahon T A and Kronauer R E, 1976 *J. Theor. Biol.* **59** 443
[15] Zamir M, 1976 *J. Theor. Biol.* **62** 227
[16] Stevens P S, 1974 *Patterns in Nature* (Boston, MA: Little, Brown and Company)
[17] Mathias N and Gopal V, 2001 *Phys. Rev.* E **63** 021117
[18] Black W R, 2003 *Transportation: A Geographical Analysis* (New York: Guilford Press)
[19] Cormen T H, Leiserson C E, Rivest R L and Stein C, 2001 *Introduction to Algorithms* 2nd edn (Cambridge, MA: MIT Press)
[20] Witten T A and Sander L M, 1981 *Phys. Rev. Lett.* **47** 1400
[21] Niemeyer L, Pietronero L and Wiesmann H J, 1984 *Phys. Rev. Lett.* **52** 1033
[22] Batty M, Longley P A and Fotheringham A S, 1989 *Environment and Planning* A **21** 1447
[23] Fabrikant A, Koutsoupias E and Papadimitriou C H, 2002 *Proc. Int. Colloquium on Automata, Languages and Programming* (*Springer Lecture Notes in Computer Science* vol 2380) ed P Widmayer, F T Ruiz, R M Bueno, M Hennessy, S Eidenbenz and R Conejo (Berlin: Springer) pp 110–2
[24] Eden M, 1961 *Proc. 4th Berkeley Symp. on Mathematical Statistics and Probabilities* ed F Neyman (Berkeley, CA: University of California Press) pp 223–39
[25] Nekka F, Kyriacos S, Kerrigan C and Cartilier L, 1996 *Bull. Math. Biol.* **58** 409
[26] Moore D, McNulty L J and Meskauskas A, 2005 *Branching Morphogenesis* ed J Davies (Georgetown: Landes Bioscience/Eurekah.com)