

Spatial graphs cost and efficiency: exploring edges competition by MCMC

Guillaume Guex

Department of Geography,
University of Lausanne
guillaume.guex@unil.ch

Abstract. Recent models for spatial networks have been built by determining graphs minimizing some functional F composed by two antagonist quantities. Although these quantities might differ from a model to another, methods used to solve these problems generally make use of simulated annealing or operations research methods, limiting themselves to the study of a single minimum and ignoring other close-to-optimal alternatives. This contribution considers the arguably promising framework where the functional F is composed by a graph *cost* and a graph *efficiency*, and the space of all possible graphs on n spatially fixed nodes is explored by MCMC. Covariance between edges occupancy can be derived from this exploration, revealing the presence of cooperative and competition regimes, further enlightening the nature of the alternatives to the locally optimal solution.

Keywords: Spatial graph models, Efficiency, Cost, MCMC, Principal component analysis

1 Introduction

Spatial networks constitute a particular case in networks studies, where nodes and edges are embedded in a metric space. The study of these networks received a special attention in the recent years, as they model a large quantity of complex geographic systems, such as transportation networks (road, railroad and airlines networks), power grids networks and internet ([1, 2, 4–6, 8, 9, 11, 12, 16, 19, 21]). The particularity of these networks is that the underlying space directly controls the cost of edges, thus impacting their topology. Previous empirical studies have examined different spatial network structures and demonstrated that the effect of space greatly differs, depending on the nature of networks (reviewed extensively in [4]). Nevertheless, their designs typically attempt to maximize some utility function while minimizing some kind of cost function, making abstraction of other geographical or economical constraints encountered in real-world situations.

This article attempts to study a particular class of models of optimal networks defined as networks minimizing some functional F specified below. These models exhibit a great variety of interesting results, depending on the ingredients

entering the composition of F , and are aimed at modelling numerous different geographic systems of interest. Here, we will consider the case where $F = C - I \cdot E$, where C is the *cost* of the network (the sum of all edges length) and E the *efficiency* (the mean length of shortest-paths between all pairs of nodes), while the parameter I , the *investment*, acts as a balance between those quantities. This simple and intuitive model, already studied in [1, 2, 4, 11], gives results similar to our railroads, highways or power grid networks. Previous researches concentrated on finding a single graph minimizing F , discarding the study of the nature of the space of all possible graphs on n fixed nodes, controlled by F . By contrast, we attempt here to explore this space with a *Monte Carlo Markov Chain* (MCMC) algorithm ([3, 7, 10, 14, 17, 18]) or more precisely, a variant of simulated annealing model, implying heating as well as cooling schedules (see section 2.4). By examining the history of the algorithm, edge competition and synergies can be revealed, enabling the design of close-to-optimal graphs.

This article is divided in two parts. The first one sets the formalism and the mathematical tools needed to perform the algorithm and the second one examines a few case studies in more detail.

2 Formalism

2.1 Generalities and notations

A *graph* is a couple $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} are the vertices (or nodes) set of size n and \mathcal{E} the edges set of size m . A graph is said to be *spatial* when all vertices are embedded in a Euclidean space. A spatial graph is entirely defined by two matrices: X the matrix of vertex coordinates in space and the $n \times n$ symmetric adjacency matrix $A = (a_{ij})$, where $a_{ij} = 1$ if $\{i, j\} \in \mathcal{E}$, $a_{ij} = 0$ otherwise.

This article considers simple unoriented spatial graphs in \mathbb{R}^2 equipped with the Euclidean distance d^E . In this context, every edge $e = \{i, j\}$ possesses a *length* l corresponding to the Euclidean distance between nodes composing it, i.e. $l(\{i, j\}) = d^E(i, j)$. Edge lengths permit to define an alternative version of the well-known shortest-path distance, referred to, in the literature, as the *weighted shortest-paths distance* d^{wsp} (or route distance in [1, 2, 4, 11]):

$$d^{wsp}(i, j) = \min_{\xi \in \mathcal{P}(i, j)} \sum_{e \in \xi} l(e)$$

where $\mathcal{P}(i, j)$ is the set of all paths between i and j .

2.2 Functional minimization

Some other quantities can be defined on spatial graphs. Define the *cost* C of a graph \mathcal{G} as the sum of all edge lengths:

$$C(\mathcal{G}) = \sum_{e \in \mathcal{E}} l(e)$$

Futhermore, define the *efficiency* E of graph \mathcal{G} as the mean, along all pairs of vertices, of the inverse of the weighted shortest-path distance:

$$E(\mathcal{G}) = \frac{1}{n(n-1)} \sum_{i \neq j \in \mathcal{V}} \frac{1}{d^{wsp}(i, j)}$$

Obviously, for any set of vertices, the empty graph yields a null cost and efficiency, while the complete graph gives their maximum. From a concrete point of view, the efficiency represents the ability of the network to effectively transport agents from any node to another, while the cost is self speaking. Therefore, an optimal network planning may seek to maximize the efficiency while minimizing the cost, leading to the minimization of the function F defined by:

$$F(\mathcal{G}) = C(\mathcal{G}) - I \cdot E(\mathcal{G})$$

where the parameter $I \geq 0$ is the *investment*, acting as an arbiter between the conflicting objectives. When $I \rightarrow 0$ the graph minimizing F is the empty graph, while $I \rightarrow \infty$ generates the complete graph. For carefully chosen intermediate values, depending in turn on several parameters such as the real cost of the edges and the insistence on the efficiency of the network, the solutions are similar to some real spatial networks, like railroad, highways or power grid networks ([2, 4, 11]). Note that, unless $I \rightarrow \infty$, the resulting graph may not be connected (see e.g. left plot in Fig. 1). If we replace the weighted shortest-path distance by the standard shortest-path distance in the formula for efficiency, optimal graphs will possess a structure of "Hub-and-spoke", similar to an airline network ([4, 11]).

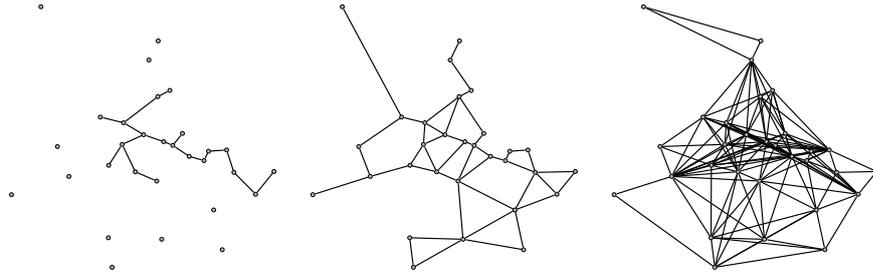


Fig. 1. Local minima for different investment values on 30 fixed points in \mathbb{R}^2 with abscissas and ordinates generated as $\mathcal{N}(0, 1)$. On the left $I = 1$, in the middle $I = 100$ and on the right $I = 10^6$.

2.3 MCMC exploration of the space of all graphs

For a fixed set \mathcal{V} of n vertices in space, the *space of all graphs* on these vertices is noted $\Gamma_{\mathcal{V}}$. This space is similar to the atomic spins space in the Ising model, where every possible edges can be in two states { presence, absence } ([10, 15]). Thus the size of $\Gamma_{\mathcal{V}}$ is $2^{n(n-1)/2}$. Let $W = (w_{kl}) = (w_{\mathcal{G}_k \mathcal{G}_l})$ be the transition matrix of a *Markov chain* on $\Gamma_{\mathcal{V}}$ defined by:

$$w_{kl} = w_{lk} = \begin{cases} \frac{2}{n(n-1)} & \text{if } \mathcal{G}_k \text{ and } \mathcal{G}_l \text{ differ exactly by one edge} \\ 0 & \text{otherwise} \end{cases}$$

With this simple transition matrix, the Markov chain will jump from the graph \mathcal{G}_k to any graph \mathcal{G}_l having exactly one more or one less edge with equal probability. It is obvious that this Markov chain can reach any nodes from any starting point, and therefore the chain is *irreducible*. The MCMC *Metropolis-Hasting* (MH) algorithm ([3, 14, 17]), is designed to create a new Markov chain having a desirable stationary distribution p_k on the states from any irreducible Markov chain:

1. From the state k , generate a new state l with probability w_{kl}
2. Jump to l with probability \tilde{w}_{kl} defined by:

$$\tilde{w}_{kl} = \min \left(1, \frac{p_l w_{lk}}{p_k w_{kl}} \right)$$

otherwise stay in k

3. Iterate

Since $w_{kl} = w_{lk}$, one has that $\tilde{w}_{kl} = \min(1, \frac{p_l}{p_k})$.

For any initial configuration, the algorithm will converge to the invariant distribution p_k , itself determined so as to favor near-optimal graphs, as in the *Gibbs sampling* of p_k ([3, 7]):

$$p_k = \frac{1}{Z} \exp(-\beta F(\mathcal{G}_k))$$

where $Z = \sum_k \exp(-\beta F(\mathcal{G}_k))$ is the standardization constant and $\beta = \frac{1}{T}$ is the *inverse temperature* parameter ($T \geq 0$ is the *temperature*). In fact, the value of β controls the randomness of the MH algorithm jumps, as seen by:

$$\tilde{w}_{kl} = \min \left(1, \frac{\exp(-\beta F(\mathcal{G}_l))}{\exp(-\beta F(\mathcal{G}_k))} \right) = \min \left(1, e^{\beta(F(\mathcal{G}_k) - F(\mathcal{G}_l))} \right)$$

If $\beta \rightarrow 0$, then $\tilde{w}_{kl} = 1$, i.e. the MH algorithm will jump to any candidate state l , while $\beta \rightarrow \infty$ implies that $\tilde{w}_{kl} = 1$ iff $F(\mathcal{G}_l) \leq F(\mathcal{G}_k)$ and $\tilde{w}_{kl} = 0$ otherwise.

2.4 Cooling schedule and exploration history

While a local minimum can easily be obtained with the MH algorithm and a simulated annealing cooling schedule (left plot in Fig. 2, as seen in [13, 20]), we are more interested here by the history of the exploration of space $\Gamma_{\mathcal{V}}$. Indeed, local minima are arguably often not really compatible with some real life constraints and we would be interested in finding alternative, but still efficient, ways to build the network. That is why we need our cooling schedule to be reheated periodically (right plot of Fig. 2) in order to avoid to be stuck in the same local minimum and to explore different parts of the space.

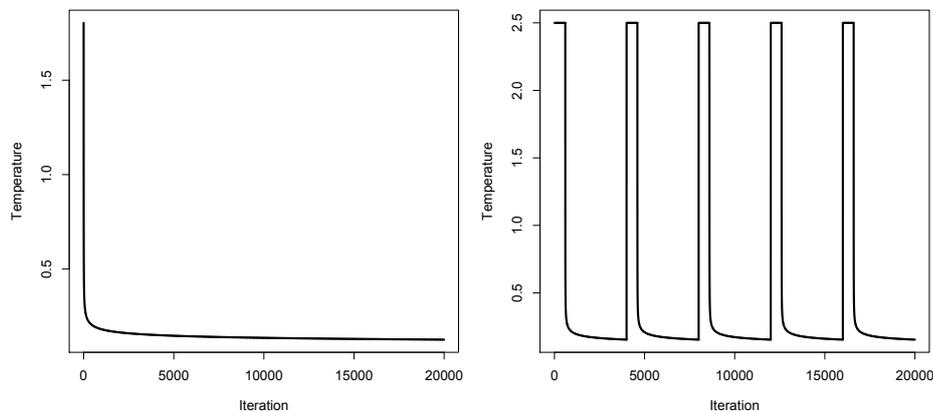


Fig. 2. Two types of cooling schedule: the first case represents a classical simulated annealing cooling schedule designed to find only one local minima: $T(t) = c/\log(1+t)$ with $c = 1.25$. In the second case, we periodically set a high temperature during 400 iterations followed by a similar cool-down, in the hope of finding another minimum.

Recording the graph history $\{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_t\}$ of a MH run by keeping track of the states of every edges modified at least once $\{e_1, e_2, \dots, e_p\}$, permits to obtain statistics on the behavior of the MH algorithm. Let the *history* matrix $H = (h_{r,s})$ defined as followed:

$$h_{r,s} = \begin{cases} 1 & \text{if the edge } e_s \text{ was present in the graph } G_r \\ 0 & \text{otherwise} \end{cases}$$

This matrix can be viewed as an usual “individuals \times variables” matrix, enabling the computation of various indices. For instance, we can calculate the probability of the appearance of an edge as $p(e_s) = h_{\bullet,s}/t$, its variance $\text{var}(e_s) = p(e_s)(1 - p(e_s))$ and the variance-covariance matrix between edges as $\Sigma = \frac{1}{t}H^c H^c$ (where H^c is the matrix H after column centration). This variance-covariance matrix permits in turn to apply a *principal component analysis*, where the factor scores

of all observed graphs in the history will underline recurrent configurations in $T_{\mathcal{Y}}$ and the saturations between edges will highlight the competition or the cooperation existing between them.

3 Case studies

3.1 Randomly located nodes

Let us first analyse the behavior of the MH algorithm on small sized graphs. 30 nodes in \mathbb{R}^2 with coordinates following a $\mathcal{N}(0,1)$ are drawn, I is set to 50 and the temperature follows the cooling schedule exhibited on the right in Fig. 2 during $t = 20'000$ iterations.

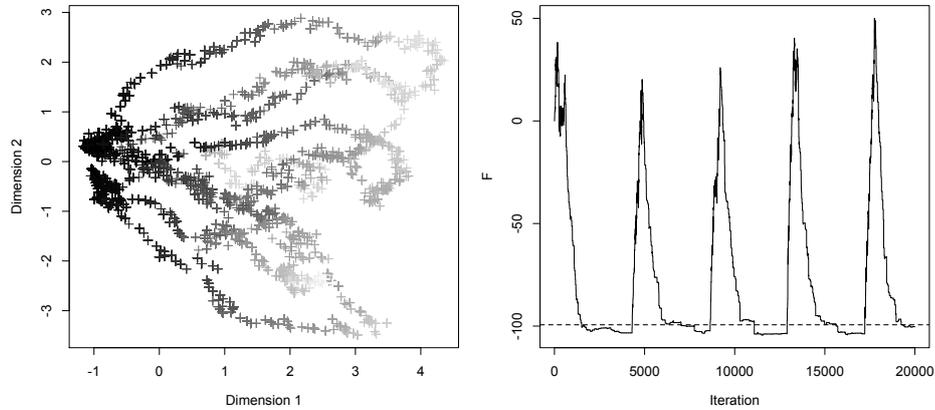


Fig. 3. Results for the graph history in the MH run. Left: first two dimensions of the factor scores, each point represents an iteration (proportion of variance explained: 13%, respectively 9%). The lower the value of F is for each iteration, the darker the point. Right: value of the functional F versus iteration.

As apparent on Fig. 3, the algorithm does not explore very efficiently the graph space. Each time we reheated the system it escaped from a local minimum before converging again on each cool down. The different minima seem to be close to each other, at least according to what appears in the first two factor scores (explaining only 13% and 9% of the variance). Fig. 4 confirms the closeness among the different minima, since some critical edges appear more frequently than others. The saturation plot shows that edges appearing frequently are correlated positively between themselves.

These results, while interesting, are a bit tarnished by the presence of high temperature states. While the presence of these states is essential to escape local

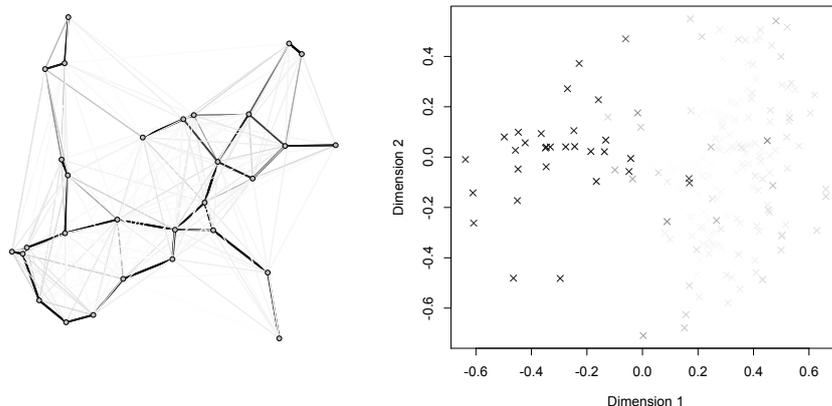


Fig. 4. Results for the edges in the MH run. Left: all edges created at least once during the process. Right: the saturation plot, where each points representing an edge and proximity capture correlation, i.e. two edges appearing frequently together will be close to each other. On both graphics, the darkness of an element is proportional to its apparition frequency during the run.

minima, they bear very little information on optimal and alternative solutions to the efficient network building. Therefore, a second analysis is performed after removing all states having a functional value higher than -100 (corresponding to the dotted line on the Fig. 3, functional plot). By construction, the selected states constitute near-optimal solutions.

The graphic in Fig. 5 illustrates the emergence of five different "cold" temperature regimes during the MH run differ more than what it appear at first glance, showing that they indeed correspond to different local minima of F . Points in the middle yield the lowest value of the functional and correspond to the third cool-down. Graphics in Fig. 6 emphasize the edges created during the process. Here, we can observe some competition between edges. For example, edges numbered 7 and 36, 42 and 49, 20 and 39, 22 and 46, are placed on the opposite side one to another in the saturation plot. In the graph, we can see that both pairs represent building alternative to a close-to-optimal graph. On the other hand, edges 11, 10 and 26 are very centered, meaning that they have a very low variance and represent a kind of "backbone" appearing in any close-to-optimal graph. Iterations 8'643 and 12'903 in Fig. 7 exhibit some built variations. Note that state 12'903 to have a lower functional value than state 8'643 (-103.8 versus -100.4).

3.2 US cities

To study a real life case, the algorithm will be run on nodes representing US cities with more than 500'000 inhabitants (Fig. 8). Latitudes θ_i and longitudes

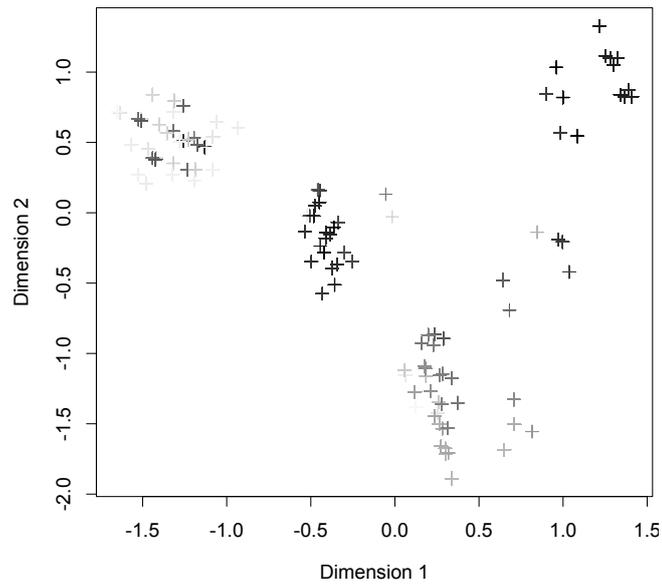


Fig. 5. The first two factorial coordinates of the states of the MH runs where states with high values of F have been removed. The five different cold temperature phases appear clearly, illustrating five different local minima.

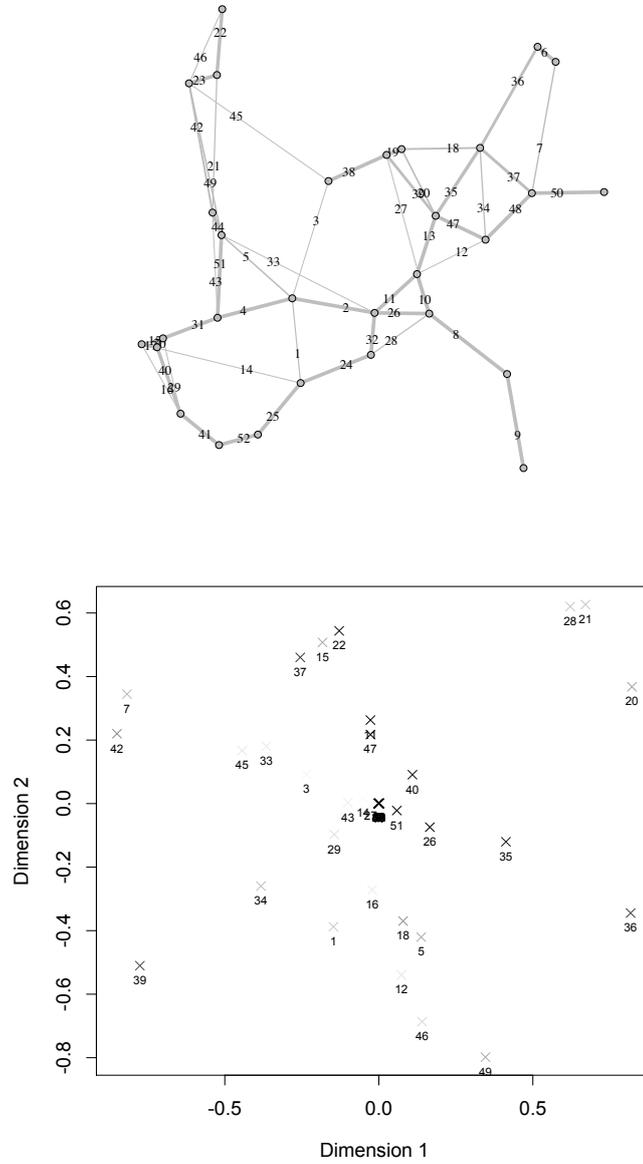


Fig. 6. Top: edges occupation frequencies. Bottom: saturations, with the same labeling.

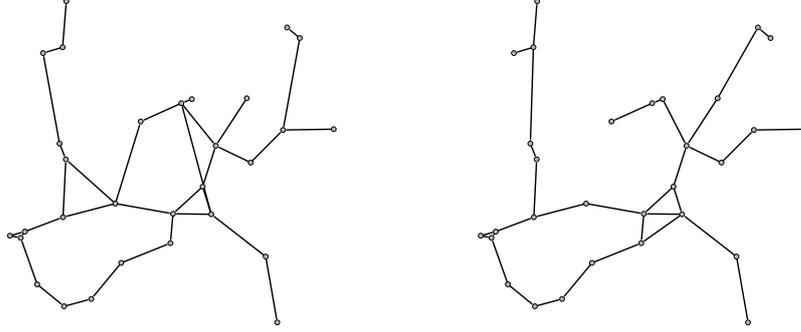


Fig. 7. Graph states at iteration number 8'643 (on the left) and 12'903 (on the right). Their F value are respectively -100.4 and -103.8.

α_i have been extracted from the R data `world.cities{maps}` and we consider the *geodesic dissimilarity* between those cities: $D_{ij} = \arccos^2(\kappa_{ij})$, where $\kappa_{ij} = \sin \theta_i \sin \theta_j + \cos \theta_i \cos \theta_j \cos(\alpha_i - \alpha_j)$. Again, 20'000 iterations of MH are run with an investment of $I = 50$ (distances have been multiplied by 30 to match distances of the previous example) and higher temperature states have been removed from analysis.

Here again, edges frequencies and saturations in Fig. 9 reveal the occurrence of competing edges in the construction of the network together with some robust edges. Note the possibility of weighting each node relatively to its population resulting in a weighted efficiency functional, currently under investigation. Nevertheless, the present result can constitute a good start to explore ways of building real networks, where particular edges can be discarded in a second time, due to some morphological or economical constraints.

4 Conclusion

Exploring the possible graphs states on n nodes by MCMC not only reveals alternatives to the optimal network, but also gives insights on the structure of this space as controlled by the functional F . In the present case, the functional makes the shortest-paths requirement conflicts with the length of the edges, and permits to preliminary explore how the shortest-paths distance is linked to the Euclidean distance in this context. The investment, the cooling schedule, the starting state and the number of iterations are shown to greatly affect this exploration, and a careful design should be made depending on what is searched. The question of how to precisely set the parameters according to spatial configuration in hand

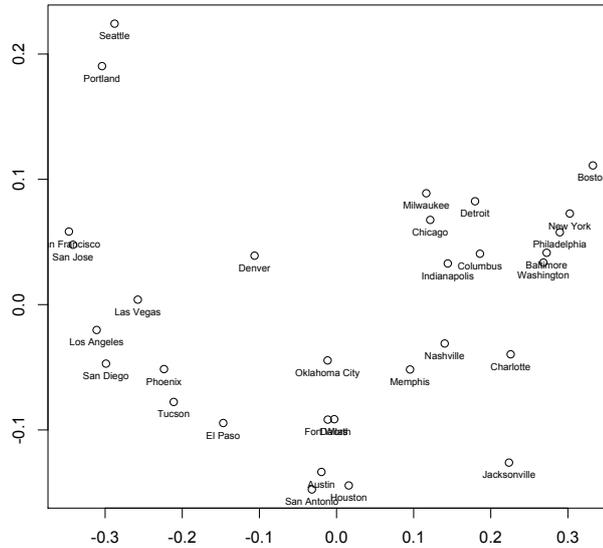


Fig. 8. Representation of the US cities with more than 500'000 inhabitants created by multidimensional scaling from their geodesic dissimilarities.

remains largely open, and a deeper study should be performed before implementing this algorithm in real life applications. The numerical complexity and computational demands of the algorithm are also quite heavy, requiring a way of optimizing the parameters before applying the algorithm to a larger set of nodes. Nevertheless, case studies already show promising results and, provided the procedure can be efficiently refined, its flexibility should permit numerous applications to a large variety of situations.

References

1. Aldous, D.J.: Optimal spatial transportation networks where link costs are sub-linear in link capacity. *Journal of Statistical Mechanics: Theory and Experiment* **2008**(03) (2008) P03006
2. Aldous, D.J., Shun, J., et al.: Connected spatial networks over random points and a route-length statistic. *Statistical Science* **25**(3) (2010) 275–288
3. Andrieu, C., De Freitas, N., Doucet, A., Jordan, M.I.: An introduction to mcmc for machine learning. *Machine learning* **50**(1-2) (2003) 5–43
4. Barthélemy, M.: Spatial networks. *Physics Reports* **499**(1) (2011) 1–101
5. Berg, J., Lässig, M.: Correlated random networks. *Physical review letters* **89**(22) (2002) 228701
6. Brede, M.: Coordinated and uncoordinated optimization of networks. *Physical Review E* **81**(6) (2010) 066104

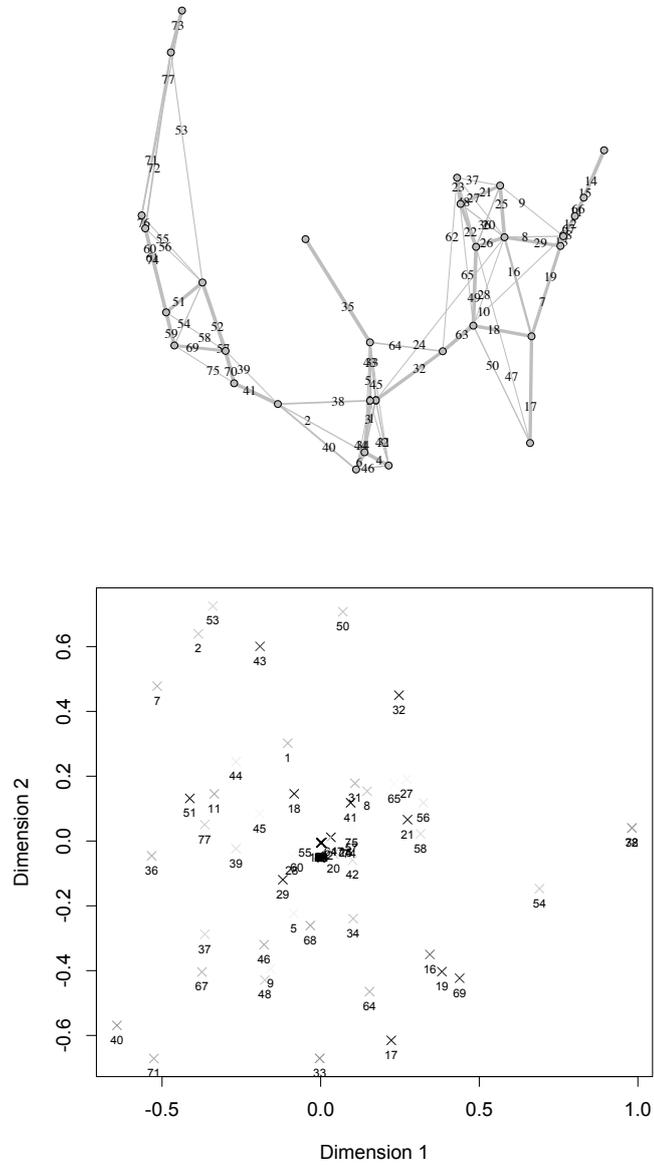


Fig. 9. Top: edges occupation frequencies. Bottom: saturations resulting from a complete MH run with 20'000 iterations.

7. Carter, C.K., Kohn, R.: On gibbs sampling for state space models. *Biometrika* **81**(3) (1994) 541–553
8. Courtat, T., Gloaguen, C., Douady, S.: Mathematics and morphogenesis of cities: A geometrical approach. *Physical Review E* **83**(3) (2011) 036106
9. Crucitti, P., Latora, V., Marchiori, M.: A topological analysis of the italian electric power grid. *Physica A: Statistical Mechanics and its Applications* **338**(1) (2004) 92–97
10. Ferrenberg, A.M., Swendsen, R.H.: New monte carlo technique for studying phase transitions. *Physical review letters* **61**(23) (1988) 2635
11. Gastner, M.T., Newman, M.E.: The spatial structure of networks. *The European Physical Journal B-Condensed Matter and Complex Systems* **49**(2) (2006) 247–252
12. Gendron, B., Crainic, T.G., Frangioni, A.: *Multicommodity capacitated network design*. Springer (1999)
13. Hajek, B.: Cooling schedules for optimal annealing. *Mathematics of operations research* **13**(2) (1988) 311–329
14. Hastings, W.K.: Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**(1) (1970) 97–109
15. Ising, E.: Beitrag zur theorie des ferromagnetismus. *Zeitschrift für Physik A Hadrons and Nuclei* **31**(1) (1925) 253–258
16. Mathias, N., Gopal, V.: Small worlds: How and why. *Physical Review E* **63**(2) (2001) 021117
17. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. *The journal of chemical physics* **21**(6) (1953) 1087–1092
18. Newman, M., Barkema, G.: *Monte carlo methods in statistical physics*. Oxford University Press (1999)
19. Valverde, S., Cancho, R.F., Sole, R.V.: Scale-free networks from optimal design. *EPL (Europhysics Letters)* **60**(4) (2002) 512
20. Van Laarhoven, P.J., Aarts, E.H.: *Simulated annealing*. Springer (1987)
21. Yang, H., Bell, M.G.: Models and algorithms for road network design: a review and some new developments. *Transport Reviews* **18**(3) (1998) 257–278