Conceptual framework

Compare possible networks on given n cities.

Optimize trade-off between
- cost to build/operate network
- benefit (to operator, or cost to users).

Usually studied as algorithmic question. This talk focusses on theoretical understanding of properties of optimal networks in the $n \to \infty$ limit.

Spatial networks arise in many disciplines

Telephone (landline; cell)

Transportation (road, rail)

Distribution (electricity grid; Walmart)

Regional (spatial) economics

Biological (e.g. blood circulation to cells)

But what does this have to do with probability?

Analogy: the Galton-Watson branching process provides a mathematically simple "toy model" for broad notion of "branching process".

Goal: find a collection of mathematically simple "toy models" for spatial networks. Randomness can be introduced to model disorder (inhomogeneity) in space. One well known model is

Model 1: the geometric random graph (e.g. Penrose (2003) monograph).

Poisson point process; link two points if they are at distance $\leq c$ apart.

Ingredient in models for cell phone ("ad hoc") networks; much EE work over last 10 years, e.g.
P Gupta, PR Kumar (2000): The capacity of wireless networks. [Cited by 1496].

I'll show three snapshots of different models, involving

optimal design of networks (A, B)

optimal flow through a random network (C).

We study $n \to \infty$ asymptotics ($n =$ number of "cities"), which is a different (and less realistic?) methodology from what's done in other "spatial networks" disciplines.

Use a "density 1" convention: n cities in square of area n.

(A): Hub-and-spoke networks

(passenger air travel; package delivery)

Seek to model the situation where the time to travel a route depends on route length and number of hops/transfers. Introduce a weighting parameter Δ and define (for a network $\mathcal{G}_n$ linking n cities $\mathbf{x}_n$ in square of area n)

time to traverse a given route from x_i to x_j

$= n^{-1/2}$(route length)$+\Delta$(number of transfers).

$$\text{time}(i,j) = \text{ min. time, over all routes}$$

$$\begin{aligned} \text{time}(\mathcal{G}_n) &= \text{ave}_{i,j}\text{time}(i,j) \\ &\geq n^{-1/2}\text{ave}_{i,j} d(i,j) := \text{dist}(\mathbf{x}_n). \end{aligned}$$

This set-up leads to a 2-parameter question. What network $\mathcal{G}_n$ over cities $\mathbf{x}_n$ minimizes time($\mathcal{G}_n$) for a given value of total length and Δ?

Some numerical solutions from Gastner - M. Newman (2006).

Let's think about designing a network where routes involve 3 hops (2 transfers). Here's one scheme.
Divide area-n square into subsquares of side L.
Put a **hub** in center of each subsquare.
Link each pair of hubs.
Link each city to the hub in its subsquare (a **spoke**).

Cute freshman calculus exercise: what total network length do we get by optimizing over L?

[length of short edges]: order nL

[length of long edges]: order $(n/L^2)^2 n^{1/2}$.

Sum is minimized by $L =$ order $n^{3/10}$ and total length is order $n^{13/10}$.

This construction gives a network such that (even for worst-case configuration $\mathbf{x}_n$)

$$(i) \qquad \text{time}(\mathcal{G}_n) - \text{dist}(\mathbf{x}_n) \to 2\Delta$$

$$\text{len}(\mathcal{G}_n) = O(n^{13/10}).$$

Theorem 1 *For "really 2-dimensional" $\mathbf{x}_n$, no networks satisfying (i) can satisfy*

$$\text{len}(\mathcal{G}_n) = o(n^{13/10}).$$

Idea of 2-page hack proof: the only way to improve the construction would be to have shorter "short edges", implying more hubs and hence more "long edges".

Can you find a 1/2-page proof?

Schematic: length of network required for a given average number of hops and given weight parameter Δ.

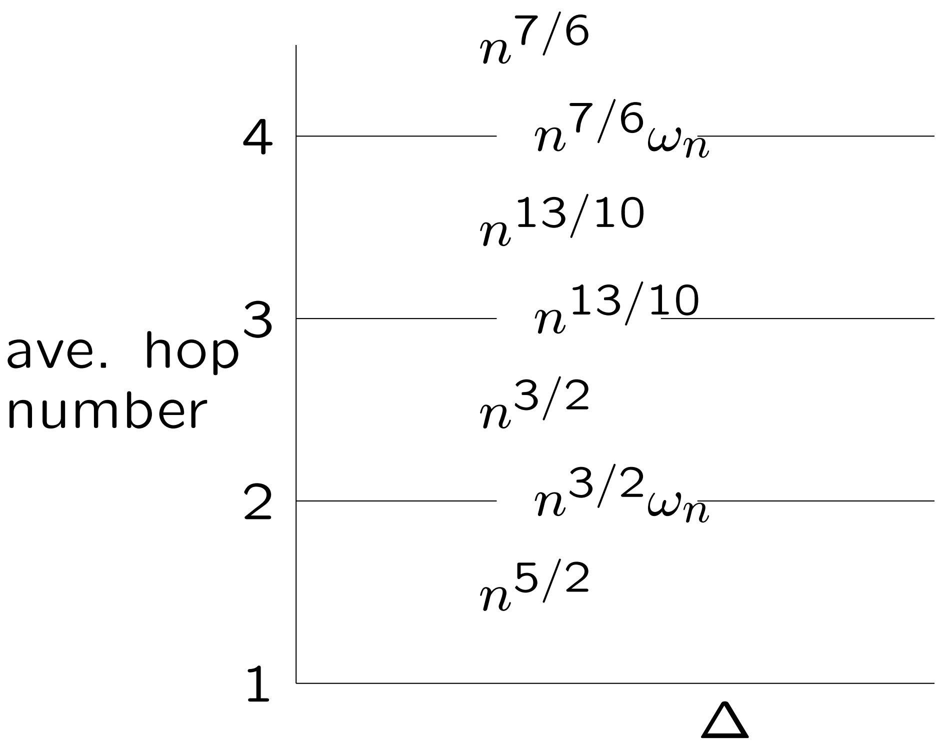

Our analysis is too crude to reveal how (for fixed large n) optimal network changes as we vary Δ.

(B): Optimal design of road/rail networks

Given n "cities" in square of area n. Want to create a network by adding edges. Study trade-off between **cost** and **benefit** of network. Take

cost = total length of network

and benefit (later) is some notion of "shortness of routes". First consider extreme case where we just minimize total length. The minimum-length connected network on cities $\mathbf{x}_n$ is by definition the **Steiner tree**, which has some length $ST(\mathbf{x}_n)$. But ST is clearly inefficient as a transportation network: routes are long.

For a network on $\{x_1, \ldots, x_n\}$ write
$d(i,j)$ = straight-line distance from x_i to x_j
$\ell(i,j)$ = route length from x_i to x_j.

First statistic to measure "benefit":

$$R := \text{ave}_{i,j} \frac{\ell(i,j)}{d(i,j)} - 1.$$

First guess at cost-benefit trade-off:

If network-length is constrained to be (say) 1.5 *times* $ST(\mathbf{x}_n)$ *then we can always make* R *less than (say)* 0.2.

It turns out that we can do much better than that. Recall that typically $d(i,j)$ is order $n^{1/2}$, so the first guess puts the "excess length" $\ell(i,j) - d(i,j)$ as order $n^{1/2}$. Recall typically len$(ST(\mathbf{x}^n))$ is order n.

Theorem 2 (with Wilf Kendall) *In worst case we can design networks $\mathcal{G}(\mathbf{x}_n)$ such that*

$$(i)\ \text{len}(\mathcal{G}(\mathbf{x}_n)) - \text{len}(ST(\mathbf{x}_n)) = o(n)$$

$$(ii)\ \text{ave}_{i,j}(\ell(i,j) - d(i,j)) = o(\omega_n \log n)$$

for $\omega_n \to \infty$ arbitrarily slowly.

(Preprint on Arxiv).

This rests upon a construction we'll show. There is a lower bound: under technical assumptions that the points are "truly 2-dimensional", if (i) holds then the average (ii) is at least order $\log^{1/2} n$.

The construction is simple: take the Steiner tree and superimpose a **Poisson line process** of small density $\eta > 0$.

Why does this network have short routes? Key is a cute calculation.

Lemma 3 *Take a PLP of rate* 1. *Erase the lines separating* $(0,0)$ *from* $(x,0)$. *Now these two points lie in a convex region* $R(\mathbf{x})$ *bounded by PLP lines.*

$$\mathbb{E}(\textit{boundary length of } R(x)) - 2x \sim \frac{8}{3}\log x.$$

So there is a route using PLP lines from near $(0,0)$ to near $(x,0)$ of length around $x + \frac{4}{3}\log x$.

Comment. The math is basically 100-year-old integral geometry.

The lower bound result is: under an "equidistribution" assumption on $\mathbf{x}_n$

For any network $\mathcal{G}(\mathbf{x}_n)$ whose length is $O(n)$,

$$\text{ave}_{i,j}(\ell(i,j) - d(i,j)) \geq c_* \log^{1/2} n.$$

7-page proof involves tension between two facts.

1. If there is a short route between x_i and x_j then a random orthogonal line (rooted where it crosses $\overline{x_i x_j}$) must cross a network line at some distance $\leq y_n$ from the root and at same angle $\frac{\pi}{2} \pm \delta_n$.

2. For any length Ln network in the square of area n, the positions and angles of intersections of a random line with the network have a mean intensity which just depends on L.

To relate these facts, need to know that the process

pick random x_i, x_j from $\mathbf{x}_n$, take random line orthogonal to $\overline{x_i x_j}$

is approximately the same as "take a random line". Here we need the "equidistribution" assumption on $\mathbf{x}_n$.

(C) Optimal flows through the disordered lattice. (Preprint on Arxiv). Here cities/roads correspond to vertices/edges of the two-dimensional grid.

Order-of-magnitude calculation on $N \times N$ grid. Send flow volume ρ_N between each (source,destination) pair. Average flow volume $\bar{f}$ across edges is

$$(N^2 \times N^2) \times \rho_N \times N \approx \bar{f} \times N^2$$

To make $\bar{f}$ be order 1 we take

$\rho_N = \rho N^{-3}$ where ρ is normalized demand.

Open Problem. Take i.i.d. random capacities (cap(e)) with $0 < c_- \leq$ cap(e) $\leq c_+ < \infty$. Obvious: a feasible flow with normalized demand ρ exists for $\rho < \rho_-$ and doesn't exist for $\rho > \rho_+$. Prove there is a constant ρ_* depending on distribution of cap(e) such that as $N \to \infty$

$$P(\exists \text{ feasible flow, norm. demand } \rho) \begin{array}{ll} \to 1, & \rho < \rho_* \\ \to 0, & \rho > \rho_*. \end{array}$$

Instead of focussing on capacities, let's focus on <u>congestion</u>. In a network without congestion, the cost (to system; all users combined) of a flow of volume $f(e)$ scales linearly with $f(e)$. With congestion, extra users impose extra costs on other users as well as on themselves. So cost scales super-linearly with $f(e)$.

Model: The cost of a flow $\mathbf{f} = (f(e))$ in an environment $\mathbf{c} = (c(e))$ is

$$\text{cost}_{(N)}(\mathbf{f}, \mathbf{c}) = \sum_e c(e) f^2(e).$$

Theorem 1. $N \times N$ <u>torus</u> (for simplicity)
Large constant bound B on edge-capacity (for simplicity)
i.i.d. cost-factors $c(e)$ with

$$0 < c_- \leq c(e) \leq c^+ < \infty.$$

Let Γ_N be minimum cost of flow with normalized intensity $\rho = 1$. Then

$$N^{-2} E \Gamma_N \to \text{constant}(B, \text{dist}(c(e))).$$

Comments. Methodology is to compare with flows across (boundary-to-boundary) $M \times M$ squares. Should work to prove existence of limits in other "optimal flows on $N \times N$ grid" models. But details are surprisingly hard to prove. Theorem 1 has 36 page proof!