Scale-invariant random spatial networks

David Aldous

September 8, 2010

Actual content is math theory: inventing and studying a (slightly) new class of random process, an abstraction of road networks.

- Axiomatic setup
- Math examples – particular processes within the class.
- Consequences – properties of such processes.
- Realism?

Framing the topic:

- How does your car GPS device find routes?
- Online road maps/routes differ from paper maps in 2 ways that influence the way we set up the model.
- (end of talk) reflections upon data vs specific models vs. classes of process.

How your car GPS device finds routes

You type street address ($\approx$ 100 million in U.S.).
Recognized as between two street intersections.
U.S road network represented as a graph on about 15 million street intersections (vertices).

Want to compute the shortest route between two vertices. Neither of the following two extremes is practical.

- pre-compute and store the routes for all possible pairs;
- or use a classical Dijkstra-style algorithm for a given pair without any preprocessing.

Key idea: there is a set of about 10,000 intersections (**transit nodes**) with the property that, unless the start and destination points are close, the shortest route goes via some transit node near the start and some transit node near the destination.
[Bast - Funke - Sanders - Schultes (2007); *Science* paper and patent]
Given such a set, one can pre-compute shortest routes and route-lengths between each pair of transit nodes; then answer a query by using the classical algorithm to calculate the route lengths from starting (and from destination) point to each nearby transit node, and finally minimizing over pairs of such transit nodes. Takes 0.1 sec.

Could regard this key fact (10,000 transit nodes such that . . . . . .) as merely an empirical property of real network. And in some qualitative sense it's obvious – there's a hierarchy of roads from freeways to dirt tracks, and "transit nodes" are intersections of major roads.

Is there some Theory? How do transit nodes arise in a math model? Why 10,000 instead of 1,000 or 100,000?

Abraham - Fiat - Goldberg - Werneck (SODA 2010) define **highway dimension** as the smallest integer h such that for every r and every ball of radius $4r$, there exists a set of h vertices such that every shortest route of length $> r$ within the ball passes through some vertex in the set.

They analyse algorithms exploiting transit nodes and other structure, giving performance bounds involving h and number of vertices and network diameter.

What kind of mathematical object should we model a road network as? Traditional paper maps suggest two possibilities. First, for intercity road networks.

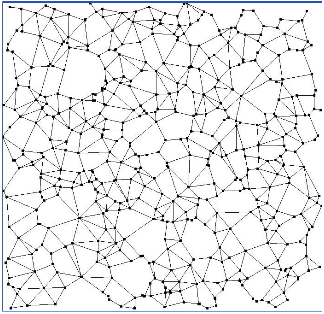

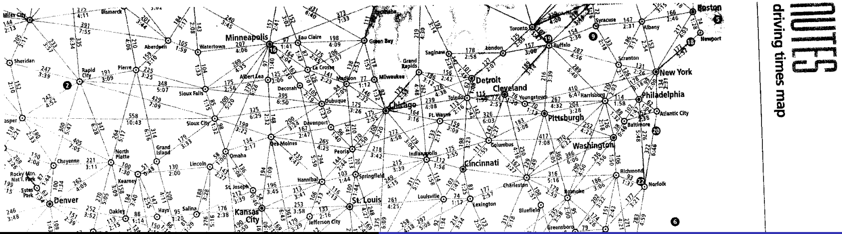

Second, a “street map” showing every road within a town.

OpenStreetMap http://www.openstreetmap.org/index.html

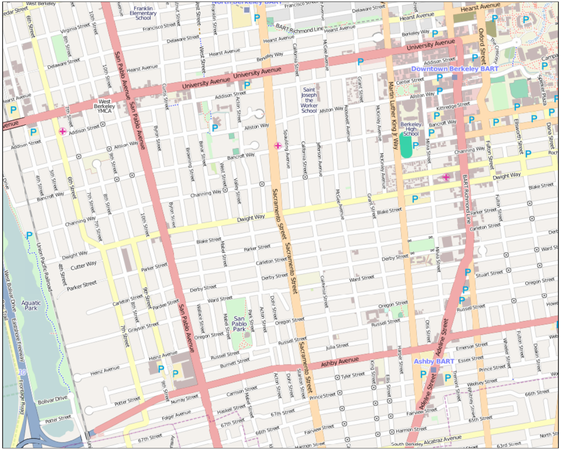

(fun aside): every road in the U.S.

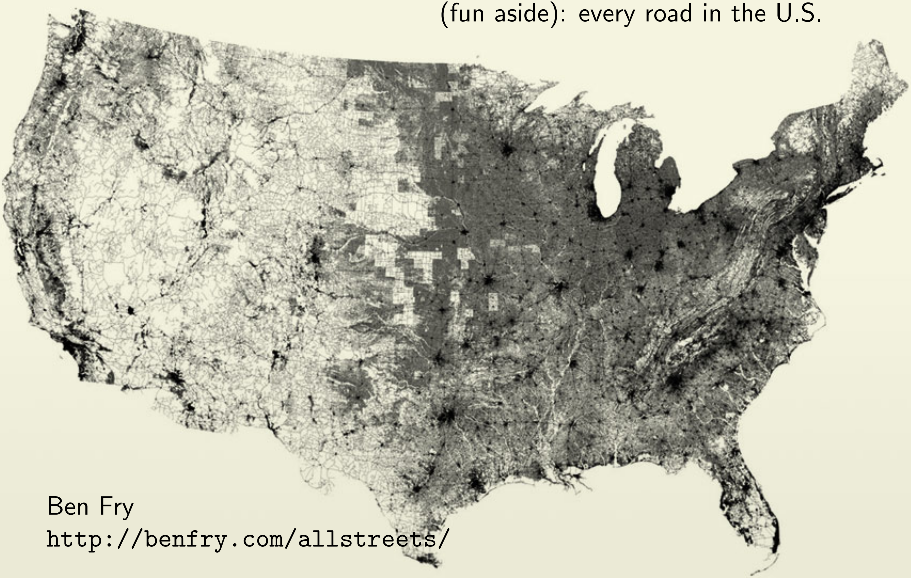

Ben Fry
http://benfry.com/allstreets/

Conceptual point: Online road maps/routes differ from paper maps in 2 [obvious] ways that will motivate our modeling.

- Can zoom in – see greater detail in window covering less area.
- Can get routes between any two specified addresses.

Print - Maps http://www.bing.com/maps/print.aspx?mkt=en-us&z=6&s=r&c...

Print

Route: **849.6 mi, 12 hr 22 min**

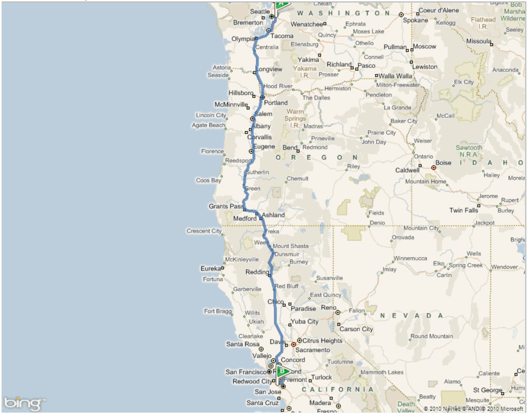

This was your map view in the browser window.

Idea behind our set-up: start with **routes** instead of **roads**.

We abstract *Google maps* as an "oracle" that for any start/desination pair (z_1, z_2) in the plane gives us a route $\mathrm{r}(z_1, z_2)$.

Analogous to ergodic theory regarding the *Hamlet* text as one realization from a stationary source, we regard *Bing maps* as containing one realization of a "continuum random spatial network". We will axiomatize this object by axiomatizing properties of random routes $\mathcal{R}(z_1, z_2)$.

The key assumption is **scale-invariance**, described intuitively as follows.

7 points in a window.

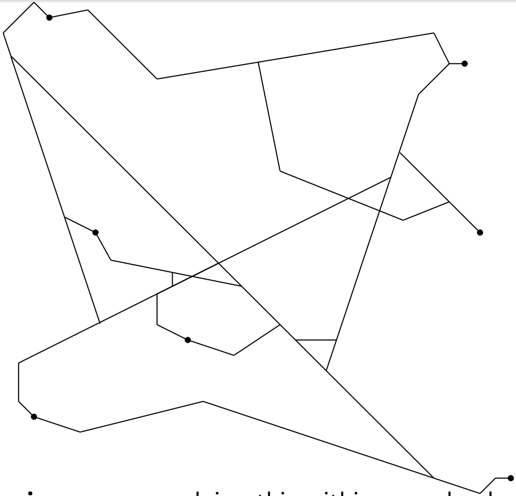

Scale-invariance means: doing this within a randomly positioned window, the statistics of the subnetwork observed don't depend on the scale, i.e. don't depend on whether the side length is 5 miles or 100 miles.

Wiener_process_animated.gif (GIF Image, 500x100 pixels) http://upload.wikimedia.org/wikipedia/commons/2/2a/Wiener_p...

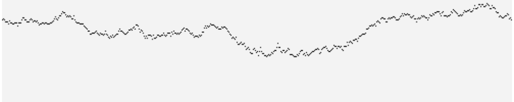

Comments re scale-invariance:

- Wikipedia has nice "zoom in" animation for Brownian scale-invariance.
- For any real-world phenomena one might model by BM, there is some "bottom level", some microscopic scale at which the BM model breaks down.
- Visualize animation of zooming into a particular model in our S-I class
- To have a network model which is exactly scale-invariant, we need to work in the continuum.
- Naive Euclidean scaling, not "scaling exponent".

Thoughts from a well-known mathematician 120 years ago.

"What do you consider the largest map that would be really useful?"

"About six inches to the mile."

"Only six inches!" exclaimed Mein Herr. "We very soon got to six yards to the mile. Then we tried a hundred yards to the mile. And then came the grandest idea of all! We actually made a map of the country, on the scale of a mile to the mile!"

"Have you used it much?" I enquired.

"It has never been spread out, yet," said Mein Herr: "The farmers objected: they said it would cover the whole country, and shut out the sunlight! So we now use the country itself, as its own map, and I assure you it does nearly as well."

Sylvie and Bruno Concluded. Lewis Carroll (1889)

Axiomatic setup: 1

Details are pretty technical, but ......

Process is presented as FDDs of random routes $\mathcal{R}(z_1, z_2)$; in other words we are given a distribution for the random subnetwork spanning each finite set $\{z_1, \ldots, z_k\}$, Kolmogorov-consistent.

Assume

- Translation and rotation-invariant
- Scale-invariant

So route-length D_r between points at (Euclidean) distance r apart must scale as $D_r \overset{d}{=} rD_1$.

Assume $ED_1 < \infty$

Axiomatic setup: 2

Envisage the route $\mathcal{R}(z_1, z_2)$ as the path that optimizes *something* but do not formalize that idea; instead

Assume a route-compatability property.

Convenient to study the process via the subnetwork $\mathcal{S}(\lambda)$ spanning a Poisson point process (rate λ per unit area).

Define a statistic

$$\ell = \text{length-per-unit-area of } \mathcal{S}(1).$$

Assume $\ell < \infty$.

There's one final assumption, which we first describe qualitatively.

Sample points of a rate-λ PPP; draw only routes between points in A and points in B.

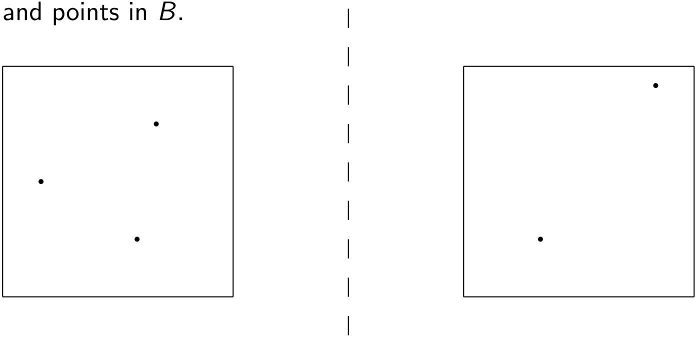

A real-world road network would have the property: as $\lambda \to \infty$ the number of places where one of these routes crosses a intervening line stays finite. We'll rephrase this in a tidier way, below.

Sample points of a rate-λ PPP; draw only routes between points in A and points in B.

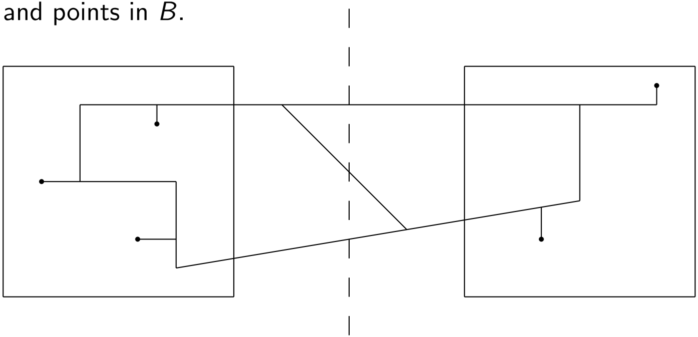

A real-world road network would have the property: as $\lambda \to \infty$ the number of places where one of these routes crosses a intervening line stays finite. We'll rephrase this in a tidier way, below.

Sample points of a rate-λ PPP; draw only routes between points in A and points in B.

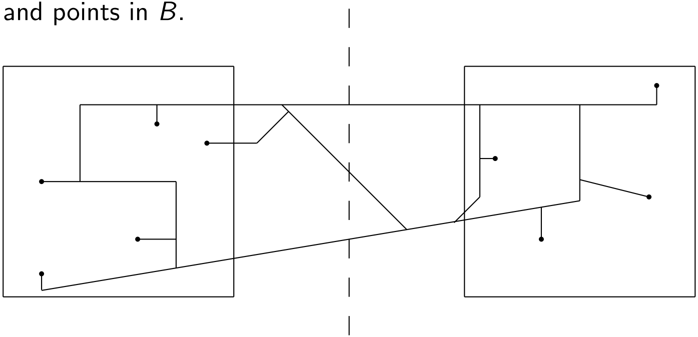

A real-world road network would have the property: as $\lambda \to \infty$ the number of places where one of these routes crosses a intervening line stays finite. We'll rephrase this in a tidier way, below.

Technical point. Recall Kolmogorov - Doob - Skorokhod theory of continuous-time processes (X_t). We nowadays typically pass from FDDs to "the process as a whole" by using a version with sample path regularity. Not obvious how to do this in our model of $\mathcal{R}(z_1, z_2)$. A realization can't have unique paths for every (z_1, z_2).

Conceptual point. There are several real-world measures of "size" of a road segment, quantifying the minor road – major road spectrum

- number of lanes
- highway numbering system
- traffic volume.

Both paper and online maps designed to indicate size of road.
Online maps designed to make minor roads become visible as we zoom in.

In our setup we are just given routes without an explicit notion of "size of road" forming the route. But, as a first "interesting consequence" of our setup, a notion of "size" automatically emerges.

Axiomatic setup: 3

Define $\mathcal{E}(\lambda, r)$ to be the road sections which are in some route $\mathcal{R}(\xi, \xi')$ in $\mathcal{S}(\lambda)$ at distance $\geq r$ from each of the endpoints ξ, ξ'.

Then let $\lambda \uparrow \infty$ to define $\mathcal{E}(\infty, r)$. Up to null sets (technical) $\mathcal{E}(\infty, r)$ is a random line process, independent of the sampling PPP, with some "mean length per unit area" $:= p(r)$.

Scale-invariance implies $p(r) = p(1)/r$.

Final assumption: $p(1) < \infty$.

We have defined a class of processes we'll call

SIRSN: Scale-invariant random spatial networks.

Technical point: we used one of 3 possible ways to start axiomatics

(1) Start with routes $\mathcal{R}(z, z')$.
(2) Start with sub-networks $\mathcal{E}(\varepsilon)$ of major roads.
(3) Start with a random metric $d(z, z')$ and define routes as geodesics.

(2) or (3) fine for constructing a **particular** model but don't seem to work to define a **class** of processes – there is no simple way to guarantee unique routes $\mathcal{R}(z, z')$ – need to add assumption of uniqueness.

- $p(1) < \infty$ is the promised tidier form of the property above, natural for a real-world road network.
- $p(1) < \infty$ serves as a technical "regularity assumption" that makes $\mathcal{E}(\infty, r)$ a tractable object and enables us to talk about the process $(\mathcal{R}(z_1, z_2);\ z_1, z_2 \in \mathbb{R}^2)$ as a whole.
- $p(1)$ is an interesting numerical statistic of a particular process in this class.
- For a road segment we can define its "size" (importance) as the largest r such that the segment is in $\mathcal{E}(\infty, r)$, that is such that it is in the route between some two points are distance $\geq r$ from itself.

Repeat page 1

Actual content is math theory: inventing and studying a (slightly) new class of random process, an abstraction of road networks.

- Axiomatic setup
- Math examples – particular processes within the class.
- Consequences – properties of such processes.
- Realism?

Framing the topic:

- How does your car GPS device find routes?
- Online road maps/routes differ from paper maps in 2 ways that influence the way we set up the model.
- (end of talk) reflections upon data vs specific models vs. classes of process.

Properties of SIRSNs.

[xxx discuss on board]

- Unique semi-infinite geodesics?
- Continuity of $(z_1, z_2) \to \mathcal{R}(z_1, z_2)$?

Repeat earlier page – GPS navigation

Could regard this key fact (10,000 transit nodes such that . . . . . .) as merely an empirical property of real network. And in some qualitative sense it's obvious – there's a hierarchy of roads from freeways to dirt tracks, and "transit nodes" are intersections of major roads.

Is there some Theory? How do transit nodes arise in a math model? Why 10,000 instead of 1,000 or 100,000?

Abraham - Fiat - Goldberg - Werneck (SODA 2010) define **highway dimension** as the smallest integer h such that for every r and every ball of radius $4r$, there exists a set of h vertices such that every shortest route of length $> r$ within the ball passes through some vertex in the set.

They analyse algorithms exploiting transit nodes and other structure, giving performance bounds involving h and number of vertices and network diameter.

Back-of-envelope calculation – model real road network as a SIRSN, truncated below.

A: area of country
η: ave number road segments per unit area.
$p(r)$: "length per unit area" of subnetwork $\mathcal{E}(\infty, r)$
Assume scale-invariant (over distances $r \gg 1$ mile), translation-invariant.

Choose any r we like; then can find a set of transit nodes (depend on r) such that
(i) Number of local (distance $< r$) transit nodes is $O(p(1))$.
(ii) regarding time-cost of single Dijkstra search as $O($ number edges$)$, the time-cost of local search is $O\left((\eta r^2) p(1)\right)$
(iii) space-cost of a $k \times k$ inter-transit-node matrix is $O(k^2)$; so this space-cost is $O\left((p(1)A/r^2)^2\right)$.

After combining costs and optimizing over r, the total cost scales as $(A\eta)^{2/3}\ p^{4/3}(1) = M^{2/3}\ p^{4/3}(1)$ for
$M =$ number of road segments in country (say 20 million).

A particular model based on "binary lattice hierarchy".

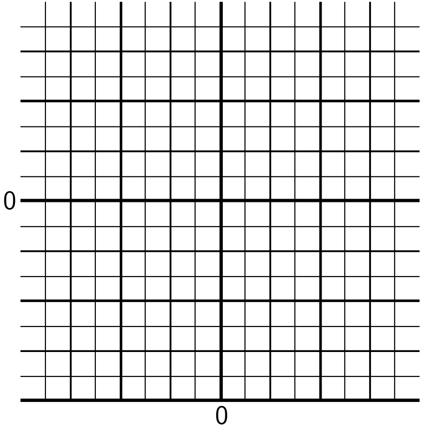

Start with square grid of roads, but impose "binary hierarchy of speeds": a road meeting an axis at $(2i+1)2^s$ has speed limit γ^s for a parameter $1 < \gamma < 2$. Define a <u>route</u> to be a shortest-time path.

(weird – axes have infinite speed limits!)

Key point of construction: given a minimum cost path from $z_1 = (x_1, y_1)$ to $z_2 = (x_2, y_2)$, scaling by 2 gives a minimum cost path from $2z_1$ to $2z_2$. (Each possible path within $2\mathbb{Z}^2$ has same relative cost as in $\mathbb{Z}^2$; never optimal to use odd-numbered streets).

Aside from the (technically hard) issue of uniqueness of routes, "soft" arguments extend this construction to a scale-invariant network on the plane.

- Consistent under binary refinement of lattice, so defines routes between points in $\mathbb{R}^2$.
- Force translation invariance by large-spread random translation.
- Force rotation invariance by randomization.
- Invariant under scaling by 2; scaling randomization gives full scaling invariance.

Need calculations (bounds) to show finiteness of the parameters.

Unexpectedly hard to prove a.a. uniqueness of routes.

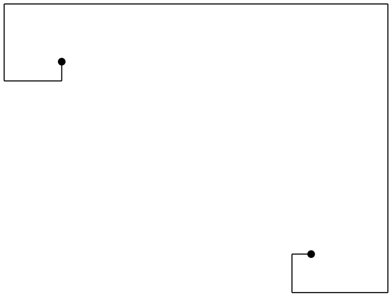

Routes like this are possible. Indeed they must happen on some scales because the process has the "unique semi-infinite geodesics" property.

Other particular models of a SIRSN: where we haven't worked through the details.

2: Use a Poisson line process, the lines having with varying speeds. (Nicer than the binary lattice model because it automatically has the desired invariance properties; harder to analyze paths to prove uniqueness of minimum-time paths).

3: "Dynamic" variants of the "static" proximity graphs like

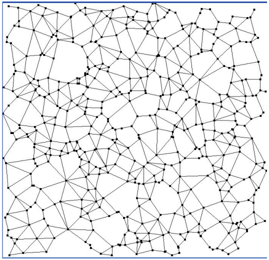

Throw points one-by-one; for a new point ξ, put an edge to an existing point ξ' iff the disc with diameter (ξ, ξ') contains no third existing point.

What about real-world road networks? Scale-invariance cannot be exactly true, but is it roughly true over some scales? What are some testable predictions?

1. Scale-invariance implies, for instance, that
(average route-length between addresses at distance $r)/r$ is constant in r which is empirically roughly correct.

2. Kalapala - Sanwalani - Clauset - Moore (2006) give data on average proportion of total route-length in the five largest segments

0 - 750 mi	0.46	0.21	0.12	0.07	0.04
750 - 1250 mi	0.40	0.21	0.13	0.08	0.05
1250 + mi	0.38	0.20	0.13	0.08	0.05

3. Could look (future undergrad project?) at topologies of the subnetworks on 4 points forming a square; do the frequencies of the different topologies vary much with scale?

Conceptual point: distinguish between a specific probability model (few parameters) a "class of processes".

Discussing a "class of processes" risks degenerating into "general abstract nonsense" theory. To stay concrete, I want to have a numerical statistic for each process in the class; the statistic should be non-obvious but say something about behavior of that process.

Class	statistic of particular model
stationary process	entropy
finite-state Markov chain	mixing time
random fractal	fractal dimension
social network models	??
S-I road network models	$p(1)$ **(this talk)**

Discussion point: in applied probability we tend to invent/study some specific model even when it's not a good fit to data; maybe more fruitful to view data as coming from some unspecified model within a specified class of models.