

➔ Overlays and small-worlds

Applicative perspectives

Gwendal Simon¹

november 2006

¹gwendal.simon@enst-bretagne.fr

➔ Milgram experiment

1967 : a social psychologist and 60 letters

- sources : various people from Nebraska
- destination : a stockbroker in Massachusetts
- constraint : only pass the letter to personal acquaintances

Main result : it works !

⇒ 6 hops in average : “ *six degrees of separation* ”

→ often distant friends : “ *the strength of weak ties* ”

→ similarities and inspiration

- people = internet
- network of acquaintances = overlay
- pass to personal acquaintances = distributed routing

→ similarities and inspiration

- people = internet
- network of acquaintances = overlay
- pass to personal acquaintances = distributed routing


Wow!!!

Billions of humans... and only 6 hops for a guaranteed routing!!

→ studying the (so efficient) social network

Clustering coefficient :


- friends of my friends are friend
- measure triangle-like neighborhood



→ studying the (so efficient) social network

Clustering coefficient :

- friends of my friends are friend
- measure triangle-like neighborhood
- here $c = \frac{9}{15}$





→ studying the (so efficient) social network

Clustering coefficient :

- friends of my friends are friend
- measure triangle-like neighborhood
- here $c = \frac{9}{15}$

Power law degree distribution

- a vast majority of nodes have few friends
- few nodes have a huge number of friends
- ... or $P(k) = k^{-\gamma}$ with $2 < \gamma < 3$



⊖ far from what we want

Many models construct network with small diameter :

- random graph
- preferential attachment
- augmented ring
- ...

→ how to find the shortest path without global knowledge ?

⊙ toward an useful definition of small-world

A small-world **navigable** overlay :

→ a *decentralized routing algorithm* finds paths of $\theta((\log^k n))$ hops

➔ toward an useful definition of small-world

A small-world **navigable** overlay :

→ a *decentralized routing algorithm* finds paths of $\theta((\log^k n))$ hops

Decentralized routing algorithm : each node only knows

- the target
- its neighbors
- *the distance between the target and each neighbor*

➔ toward an useful definition of small-world

A small-world **navigable** overlay :

→ a *decentralized routing algorithm* finds paths of $\theta((\log^k n))$ hops

Decentralized routing algorithm : each node only knows

- the target
- its neighbors
- *the distance between the target and each neighbor*

In Milgram experiment, people *knows* that :

- an accountant is closer to a stockbroker than a dentist
- New York is closer to Massachusetts than Los Angeles


→ Kleinberg model

An **augmented** d -dimensional grid

- $pr \propto \frac{1}{\|i,j\|^d}$
- many “close” *long links*

Grid + *Manhattan dist.* is a **metric** :

- $d(i,j) \geq 0$
- $d(i,j) = d(j,i)$
- $d(i,j) \leq d(i,k) + d(k,j)$



Greedy routing succeeds $\rightarrow \theta(\log^2(n))$ paths

 further

How to make a small-world ?

- 1 your overlay : **any d -bounded doubling dimension metric**
- 2 add one additional link per node
- 3 implement the simplest routing algorithm you have ever implemented

➔ further

How to make a small-world ?

- 1 your overlay : **any d -bounded doubling dimension metric**
- 2 add one additional link per node
- 3 implement the simplest routing algorithm you have ever implemented

A ball : $B(p, r) = \{q : \|p, q\| < r\}$

A d -bounded doubling dimension metric :

- any ball of radius $2 * r$ can be filled by 2^d balls of radius r

➔ further

How to make a small-world ?

- 1 your overlay : **any d -bounded doubling dimension metric**
- 2 add one additional link per node
- 3 implement the simplest routing algorithm you have ever implemented

A ball : $B(p, r) = \{q : \|p, q\| < r\}$

A d -bounded doubling dimension metric :

- any ball of radius $2 * r$ can be filled by 2^d balls of radius r

Do you know any bounded doubling dimension graph metric ?

→ some bounded doubling dimension graphs

Internet :

- the router-level topology : $d \simeq 14$
- latencies between 2500 sites over Internet : $d \simeq 9$

→ some bounded doubling dimension graphs

Internet :

- the router-level topology : $d \simeq 14$
- latencies between 2 500 sites over Internet : $d \simeq 9$

Social graphs :

- author citations in scientific papers : $d \simeq 12$
- actors having played in same movies : $d \simeq 14$

→ some bounded doubling dimension graphs

Internet :

- the router-level topology : $d \simeq 14$
- latencies between 2500 sites over Internet : $d \simeq 9$

Social graphs :

- author citations in scientific papers : $d \simeq 12$
- actors having played in same movies : $d \simeq 14$

The Kleinberg torus : $d = 2$

⊙ typical applications

Nearest neighbor graph over the planet

- *overlay* : linking device-equipped people with their neighborhood
- *service* : joining the closest guy to a specified position

Affinity-based graph on a lexicon of d categories :

- *overlay* : Delaunay triangulation of people
- *service* : discovering the guy being the closest of a “character”

➔ typical applications

Nearest neighbor graph over the planet

- *overlay* : linking device-equipped people with their neighborhood
- *service* : joining the closest guy to a specified position

Affinity-based graph on a lexicon of d categories :

- *overlay* : Delaunay triangulation of people
- *service* : discovering the guy being the closest of a “character”

Another idea : choose long links suitable for another metric

- *overlay* : Solipsis peer-to-peer virtual world
- *long links* : choosing preferentially peers :
 - far in the virtual world
 - close in the network

→ ultra-efficient teleporting service !!!

➔ conclusion

The topic has inflamed the French peer-to-peer scientific community

Some key questions when constructing overlays :

- bounds of the doubling dimension, distance function. . .

New perspectives for routing in overlays. . .

➔ conclusion

The topic has inflamed the French peer-to-peer scientific community

Some key questions when constructing overlays :

- bounds of the doubling dimension, distance function. . .

New perspectives for routing in overlays. . .

. . . but routing may be a not-so-big issue