

Four degrees of separation (and how we did it)



Sebastiano Vigna

*Dipartimento di Informatica Laboratory for
Web Algorithmics Università degli Studi di Milano*

Amphi 34B

4, place Jussieu
75005 Paris
Metro Jussieu

**6 Mai 2019
à 18h00**

Since the 60's, sociologists have been interested in how many friendship links you must traverse in average to get from any person in the world to any other one. Milgram's celebrated experiment, involving a few hundred people, concluded that there were six "degrees of separation". More generally, sociologists were interested in the distance distribution of friendship: how many pairs of people are separated by k degrees? We will discuss some new, high-performance diffusion-based approximate algorithms that made it possible to conclude that on Facebook there are 3.74 degrees of separation using commodity hardware, and to analyze how this value decreased in time. The same algorithms can be used to compute on very large graphs centrality measures based on distances, such as closeness and harmonic centrality, which provide interesting, high-quality rankings.

Sebastiano Vigna's research focuses on the interaction between theory and practice. He has worked on highly theoretical topics such as computability on the reals, distributed computability, self stabilization, minimal perfect hashing, succinct data structures, query recommendation, algorithms for large graphs, pseudorandom number generation, theoretical/experimental analysis of spectral rankings such as PageRank, and axiomatization of centrality measures, but he is also (co)author of several widely used software tools. In 2011 he collaborated to the first computation of the distance distribution of the whole Facebook graph, from which it was possible to evince that on Facebook there are just 3.74 degrees of separation. Recently, he participated to the analysis of the largest available public web crawl (Common Crawl 2012), which led to the publication of the first open ranking of web sites (<http://wwwranking.webdatacommons.org/>). His work on Elias-Fano coding and quasi-succinct indices is at the basis of the code of Facebook's "folly" library. His pseudorandom number generator xorshift128+ is the current stock generator of Google's V8 JavaScript engine, used by Chrome, Safari, Firefox, Edge, and Node.js; it is also the stock generator of the Erlang language.

contact : colloquium@lip6.fr
<http://www.lip6.fr/colloquium/>
Vidéo disponible sur le site

