

Unstructured Networks and Network Size Estimation


Christian Grothoff

Technische Universität München


April 24, 2013

Unstructured Networks — Physical


Internet Map
city-to-city connections



Unstructured Networks — Wireless



Unstructured Networks — Social



Graph Theory

With graph theory, we can:

- ▶ create realistic topologies for experiments
- ▶ use idealized topologies for mathematical analysis
- ▶ deliberately change the network towards a particular topology

Terminology

- ▶ Graph has nodes and edges
- ▶ Graphs can be connected
- ▶ Graphs / edges can be directed or undirected
- ▶ Nodes / edges can be weighted or unweighted
- ▶ Directed graphs consist of strongly connected components

Common Topologies

- ▶ Ring (*)
- ▶ Star (*)
- ▶ Full Mesh / Clique (*)
- ▶ N-dimensional Grid (*)
- ▶ N-dimensional Torus (*)
- ▶ Erdős-Renyi / Random graph
- ▶ Small-world graph
- ▶ Scale-Free graph
- ▶ Expander graph (+)

Graph Properties

- ▶ Size (edges, nodes) (*)
- ▶ Diameter (*)
- ▶ Density
- ▶ Characteristic path length
- ▶ Clustering coefficient
- ▶ Degree distribution
- ▶ Spectral expansion (+)

Density

The density of an undirected graph is:

$$d = \frac{2|E|}{|V|(|V| - 1)} \in [0, 1] \quad (1)$$

Streinu & Theran defined (k, l) -sparse as having at most $k|V| - l$ edges. Examples:

- Forests are $(1, 1)$ -sparse

Characteristic path length

Let $d(i, j)$ be the distance between nodes i and j . The **diameter** of an undirected graph is

$$D := \max_{i, j \in V} d(i, j) \quad (2)$$

The **Characteristic path length** L is the average length of a shortest path in an undirected graph:

$$L := \text{avg}_{i, j \in V, i \neq j} d(i, j) \quad (3)$$

Erdős-Renyi / Uniform Random graphs

Construction:

- ▶ Nodes connected at random
 - ▶ Determined by $|V|$ and $|E|$
 - ▶ **uniform** random graphs as there is no bias for link selection
 - ▶ Erdős-Renyi construction: connect any two nodes with probability p
- $\Rightarrow |E| \approx p \cdot \frac{n^2}{2}$

Properties:

- ▶ Average distance is most likely close to optimal given n and m
- ▶ Node degree follows binomial distribution

Random graphs are generally too simple and uniform as a model of real networks.

Clustering coefficient

Graph $G = (V, E)$, then neighbourhood Γ_v of $v \in V$ is:

$$\Gamma_v := \{u \in V | (v, u) \in E\} \quad (4)$$

Given $U \subseteq V$, define:


$$E(U) := \{(u, v) \in E | u, v \in U\} \quad (5)$$

Clustering coefficient C_v of node v is then:

$$C_v := \frac{2}{\deg(v) \cdot (\deg(v) - 1)} \cdot |E(\Gamma_v)| \quad (6)$$

Clustering coefficient C of G is $C := \frac{1}{n} \sum_{v \in V} C_v$.

Example Calculation



Towards Real networks

- ▶ Characteristic path length is small (as in random graphs)
- ▶ Clustering coefficient is high (but is **low** in random graphs)

Small-world graph

A **Small-World graph** is a graph with a characteristic path length close to that of an equivalent random graph ($L \approx L_{random}$), but with a much larger clustering coefficient ($C \gg C_{random}$).

Example	$ V $	avg. deg.	L	L_{random}	C	C_{random}
Internet	260,000	3.39	11.4	10.1	0.023	0.000014
Gnutella	n/a	n/a	3.86	3.19	0.045	0.0068
Film collab.	225,000	61	3.65	2.99	0.79	0.00027
Power grid	4,900	2.67	18.7	12.4	0.080	0.005
Neural net.	282	14	2.65	2.25	0.28	0.05

(Sources: Watts & Strogatz 1999, Li et al 2004, Jin & Bestavros 2006)

Small-world graph construction

Watts-Strogatz construction:

- ▶ Create n -dimensional torus
- ▶ For each edge, with probability p , rewire one end of it

Small-world graph construction

Watts-Strogatz construction:

- ▶ Create n -dimensional torus
- ▶ For each edge, with probability p , rewire one end of it

Kleinberg construction:

- ▶ Create n -dimensional torus
- ▶ Additionally, connect nodes u, v at random with probability $p \sim d^{-n}(u, v)$.

⇒ diameter is $\Theta(\log n)$

(Kleinberg 2001, Martel & Nguyen, 2004)

Real networks

- ▶ Characteristic path length is small (as in random graphs)
- ▶ Clustering coefficient is high (as in small-world graphs)
- ▶ Popularity / degree of nodes differs extremely

Power law distribution

In a **Power law distribution**, the frequency of the i -th most popular object is proportional to $i^{-\alpha}$.

Let $p(s_i) \geq p(s_{i-1})$ for all items $s_i \in S - \{s_0\}$ for some distribution $p : S \rightarrow \mathbb{R}$. If p is a power law distribution, then

$$p(s_i) \sim i^{-\alpha} \tag{7}$$

holds for some α , which is often called the Zipf coefficient.

Power law distribution

In a **Power law distribution**, the frequency of the i -th most popular object is proportional to $i^{-\alpha}$.

Let $p(s_i) \geq p(s_{i-1})$ for all items $s_i \in S - \{s_0\}$ for some distribution $p : S \rightarrow \mathbb{R}$. If p is a power law distribution, then

$$p(s_i) \sim i^{-\alpha} \tag{7}$$

holds for some α , which is often called the Zipf coefficient.

- ▶ Power-Law network \equiv Power Law distribution of node degree
- ▶ Also called **scale-free network** as degree distribution is independent of scale ($|V|$)

Construction of Scale-Free Networks

Barabasi-Albert construction (“The rich get richer”):

- ▶ Given n nodes
- ▶ Start with m_0 unconnected nodes, add a random link to each node ($\Rightarrow \deg \geq 1$) among those nodes)
- \Rightarrow The result is called the “seed network”
- ▶ For $i \in [1, t]$, add node v' , connecting it to m nodes selected by linear preferential attachment:

$$p(v) := \frac{\deg v}{\sum_{u \in V} \deg u} \quad (8)$$

Construction of Scale-Free Networks

Barabasi-Albert construction (“The rich get richer”):

- ▶ Given n nodes
- ▶ Start with m_0 unconnected nodes, add a random link to each node ($\Rightarrow \deg \geq 1$) among those nodes)
- \Rightarrow The result is called the “seed network”
- ▶ For $i \in [1, t]$, add node v' , connecting it to m nodes selected by linear preferential attachment:

$$p(v) := \frac{\deg v}{\sum_{u \in V} \deg u} \quad (8)$$


This construction is also insightful as to **why** real networks are often scale-free.

Early P2P Networks

- ▶ Napster
- ▶ Gnutella
- ▶ Gnutella2
- ▶ BitTorrent


Napster

- ▶ Star topology for search (central Napster Server)
- ▶ “P2P” because data transfers are done directly between clients
- ▶ Napster Inc.’s central server shutdown in 2001
- ▶ You can **still** run your own Napster Server today




Gnutella (0.4)

- ▶ “random” mesh topology (ping/pong) for search (query/hit)
 - ▶ Limited flooding as routing principle
 - ▶ Push requests for NAT traversal
 - ▶ Transfers directly between peers
- ⇒ Unstructured search = large overhead
- ▶ queries without addresses
- ⇒ weak anonymity



Gnutella2

- ▶ Peers and super peers
- ▶ Super peers index hundreds of normal peers
- ⇒ structured search = better scalability
- ▶ Relatively few super peers
- ⇒ censorship, malware, operator abuse, ...



BitTorrent (Cohen, 2003)

Network Size Estimation: Purpose


Purpose of Network Size Estimation

- ▶ Human curiosity
- ▶ Detection of unusual events
- ▶ Value of the botnet
- ▶ Tuning parameter

Today: Unstructured Methods

- ▶ Sample and Collide
- ▶ Hop Sampling
- ▶ Gossip-based aggregation
- ▶ Gossipico

Sample and Collide [3]



Sample and Collide


Advantages

- ▶ $O(\sqrt{n})$ messages (message size: $O(\sqrt{n})$ or worse?)

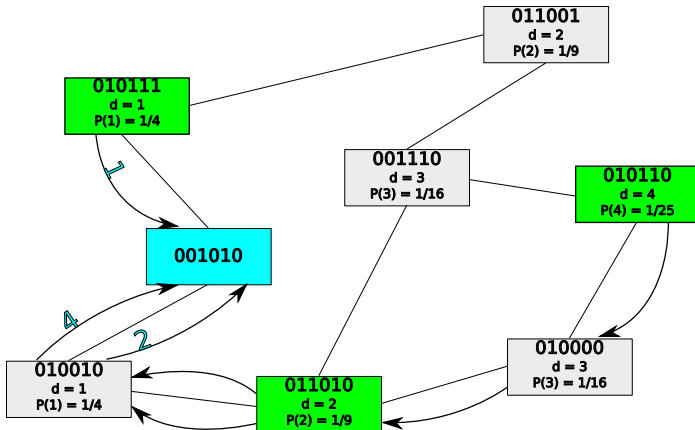
Disadvantages

- ▶ Message size could be up to $O(\sqrt{n})$ depending on the topology
- ▶ Fails in some topologies (ring, highly clustered networks) depending on sampling method
- ▶ Centralized or each node has to do the estimate, complexity in that case is $O(n^2)$ per round
- ▶ Design fails to address denial-of-service potential
- ▶ m malicious participants can force size estimate of m^2

Hop Sampling [2]



Hop Sampling



Hop Sampling

Advantages

- ▶ Works with all topologies
- ▶ $O(|E|)$ messages (message size: $O(1)$)

Disadvantages

- ▶ High load on neighbors of initiator
- ▶ If each node needs an estimate: $O(|N| \cdot |E|)$ per round
- ▶ If not centralized, each node has to keep $O(|N|)$ state to track distance to origin
- ▶ Design fails to address denial-of-service potential
- ▶ Tiny fraction of malicious participants can always create significant size overestimates

Gossip-based aggregation [1]

- ▶ One node starts with a value of $v_i = 1$, all others with $v_o = 0$
 - ▶ Select a random edge (A, B) and update values to $\frac{v_A + v_B}{2}$
 - ▶ If node A leaves, set $v_B := v_B + v_A$ for some neighbour B of A
- ⇒ Value globally converges to $\frac{1}{|N|}$

Gossip-based aggregation

Advantages

- ▶ Work with all topologies
- ▶ Network wide agreement
- ▶ $O(|N|)$ messages in total (message size: $O(1)$)

Disadvantages

- ▶ Start point agreement problem
- ▶ Slow convergence
- ▶ Very vulnerable to denial-of-service
- ▶ Very vulnerable to result manipulation
- ▶ One malicious participant affect the whole network

Gossipico [4]

- ▶ New gossip-based network counting algorithm (2012)
- ▶ Significantly higher precision than previous algorithms
- ▶ Significantly better performance than previous algorithms
- ▶ Uses “count” to accumulate number of peers in network
- ▶ Uses “beacons” to structure counting message propagation

Count

Count does not Scale


Spreading the Result

Beacons


Towards Shortest Paths

Topology Changes

Gossipico Precision



Gossipico Performance



Gossipico aggregation

Advantages

- ▶ Quite efficient
- ▶ Very good precision
- ▶ No designated node to start the process
- ▶ Supports churn

Disadvantages

- ▶ Not secure against denial-of-service attacks
- ▶ Malicious participants can change result to any value

Network Size Estimation in GUNet ¹

Functional Goals

- ▶ Supports churn
- ▶ Fully decentralized
- ▶ Efficient
- ▶ All peers obtain the network size estimate
- ▶ Operates in unstructured topologies

¹Evans, Polot, Grothoff: “Efficient and Secure Decentralized Network Size Estimation”, Networking 2012

Intuitive Idea

- ▶ Set of elements distributed in a space
- ▶ Pick a random spot
- ▶ Measure distance to nearest element
- ▶ More elements \Rightarrow smaller distance, more *overlapping*

Intuitive Idea



Intuitive Idea



Intuitive Idea



Intuitive Idea



Intuitive Idea - Applied to networks


- ▶ Space: all possible IDs
- ▶ Population: randomly distributed peer IDs
- ▶ Overlap: number of leading bits in common with a random ID

Theorem

Let \bar{p} be the expected maximum number of leading overlapping bits between all n random node identifiers in the network and a random key. Then the network size n is approximately

$$2^{\bar{p}-0.332747}$$

Empirical Measurement of the 0.33... correction



Proof (1/3)

Let X be the random variable for all n identifiers and let X_i be the number of overlapping bits for an individual random node identifier i .

The probability that a single random node identifier i overlaps with at least α bits with a random key is

$$P(X_i \geq \alpha) = 2^{-\alpha}. \quad (9)$$

Then, the probability that a single random node identifier overlaps with less than α bits with a random key is

$$P(X_i < \alpha) = 1 - 2^{-\alpha}. \quad (10)$$

The probability that the maximum number of leading overlapping bits for all n random nodes is strictly less than α is

$$P_n(X < \alpha) := P\left(\bigwedge_i X_i < \alpha\right) = (P(X_i < \alpha))^n = (1 - 2^{-\alpha})^n.$$

Proof (2/3)

Then $E_n(X)$, the expected maximum number of leading overlapping bits between n random node identifiers in the network is:

$$\begin{aligned} E_n(X) &:= \sum_{\alpha=0}^{\infty} \alpha \cdot P_n(X = \alpha) = \sum_{\alpha=1}^{\infty} P_n(X \geq \alpha) \\ &= \sum_{\alpha=1}^{\infty} (1 - P_n(X < \alpha)) = \sum_{\alpha=1}^{\infty} (1 - (1 - 2^{-\alpha})^n) \\ &= \sum_{\alpha=1}^{\log_2 n} (1 - (1 - 2^{-\alpha})^n) + \sum_{\alpha=\log_2 n+1}^{\infty} (1 - (1 - 2^{-\alpha})^n) \end{aligned}$$

Suppose n is sufficiently large such that we can use

$$\lim_{n \rightarrow \infty} (1 - \frac{x}{n})^n = e^{-x}.$$

Proof (3/3)


By substituting $\beta := \alpha - \log_2 n$ and $\gamma := \log_2 n - \alpha$ we then get:

$$\begin{aligned} E_n(X) &= \log_2 n - \sum_{\gamma=0}^{\log_2 n-1} \left(1 - 2^{\gamma-\log_2 n}\right)^n + \sum_{\beta=1}^{\infty} \left(1 - \left(1 - 2^{-(\beta+\log_2 n)}\right)\right)^n \\ &= \log_2 n - \sum_{\gamma=0}^{\log_2 n-1} \left(1 - \frac{2^\gamma}{n}\right)^n + \sum_{\beta=1}^{\infty} \left(1 - \left(1 - \frac{2^{-\beta}}{n}\right)\right)^n \\ &\approx \log_2 n - \sum_{\gamma=0}^{\log_2 n-1} e^{-2^\gamma} + \sum_{\beta=1}^{\infty} \left(1 - e^{2^{-\beta}}\right) \\ &\approx \log_2 n - 0.521865 + 0.854613 = \log_2 n + 0.332747 \end{aligned}$$

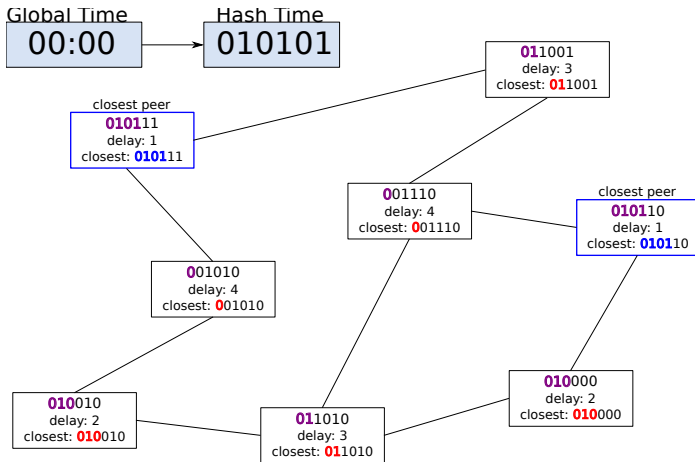
Thus, for sufficiently large values of n ,

$$E_n(X) \approx \log_2 n + 0.332747. \tag{11}$$

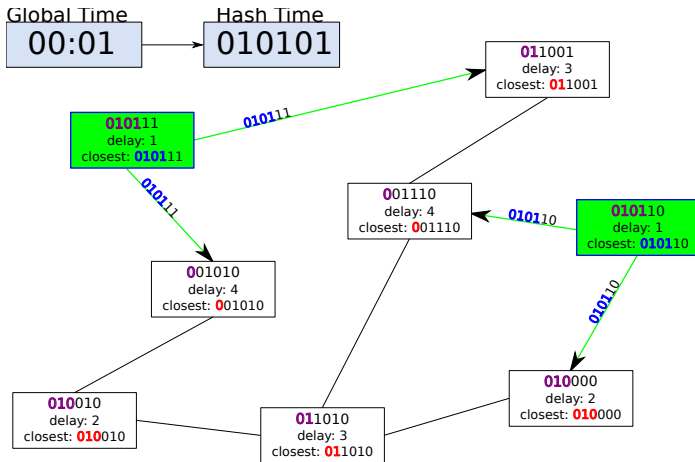
Empirical Measurement of the 0.33... correction



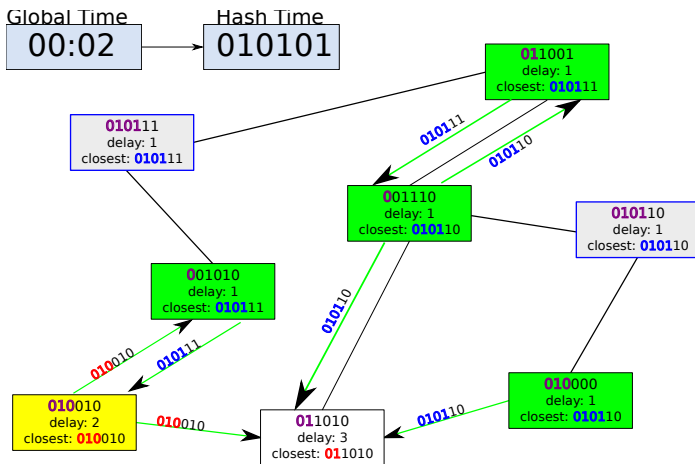
Time: 0



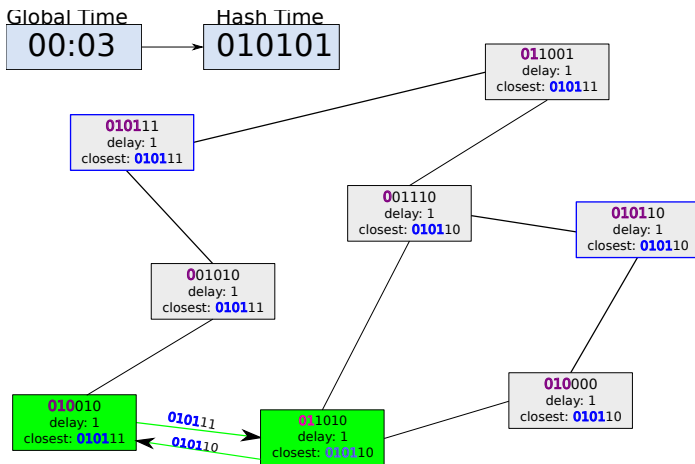
Time: 1



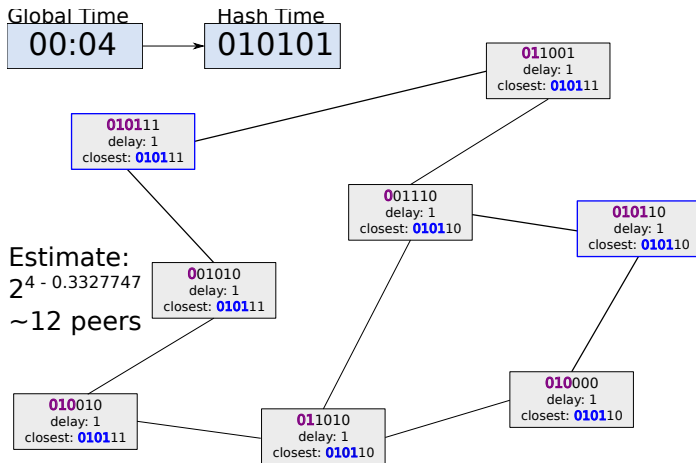
Time: 2



Time: 3



Time: 4



Our Approach: Key Points

- ▶ Use the current time to generate a random number
- ▶ More overlapping bits \Rightarrow gossip earlier
- ▶ Also delay gossip randomly to avoid traffic spikes
- ▶ Proof-of-Work to make Sybil attacks harder

Message Format

Offset	Contents
0	Message header magic code
4	Hop-Count (updated at each peer)
8	Signed data header magic code
16	Time S of the round
24	Proximity p in bits
28	Public key (2048 bit RSA)
288	Proof-of-work
296	Signature (signing bytes 8–295)

Security

Security Properties

- ▶ No trusted third parties
- ▶ Reliable
- ▶ Resistant to malicious participants


Attacker Model

- ▶ Freely participate
- ▶ Multiple identities
- ▶ May alter, drop, send/receive data
- ▶ Same resources as “normal” peers


Processing results

- ▶ Final agreed value fluctuates around the actual size
- ▶ Average and std dev over last i protocol rounds is provided


Precision vs. Rounds of Measurement




Network Size Estimate (64 rounds) under Churn




Small-World Topology




Random Graph Topology




Without Clock Skew



With Clock Skew



Agreement between peers



Compare


	S. & Coll.	Gossip	H. Sampling	GNUnet
MEM	$O(\sqrt{N})$	$O(1)$	$O(N)$	$O(1)$
CPU	$O(\sqrt{N})$	$O(N)$	$O(E)$	$O(E / N)$
NET	$O(N \sqrt{N})$	$O(N ^2)$	$O(N \cdot E)$	$O(E)$
SEC	DoS, BE	DoS, BE	DoS, BE	Pr.-of-Work
IMP	Simulation	Simulation	Simulation	Yes

(BE = Bad Estimates)

Conclusion

- ▶ Mathematical foundation applicable broadly for group size estimates
- ▶ Secure & Efficient Network Size Estimation Protocol
- ▶ Arbitrary Topologies, Clock Skew harmless, DoS resistant
- ▶ Simple to implement, free software implementation in GNUUnet
- ▶ Trade-off between precision (Gossipico) and security (GNUUnet)

Questions?



References



Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu.

Gossip-based aggregation in large dynamic networks.

ACM Trans. Comput. Syst., 23:219–252, August 2005.



Dionysios Kostoulas, Dimitrios Psaltoulis, Indranil Gupta, Ken Birman, and Al Demers.

Decentralized schemes for size estimation in large and dynamic groups.

In *Proceedings of the Fourth IEEE International Symposium on Network Computing and Applications*, pages 41–48, Washington, DC, USA, 2005. IEEE Computer Society.



Laurent Massoulié, Erwan Le Merrer, Anne-Marie Kermarrec, and Ayalvadi Ganesh.

Peer counting and sampling in overlay networks: random walk methods.

In *Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, PODC '06, pages 123–132, New York, NY, USA, 2006. ACM.



Ruud van de Bovenkamp, Fernando Kuipers, and Piet Van Mieghem.

Gossip-based counting in dynamic networks.

In *IFIP International Conferences on Networking (Networking 2012)*, pages 404–419, Prague, CZ, 05/2012 2012. Springer Verlag, Springer Verlag.