

The impact of dynamic events on the number of errors in networks

Christian Glacet^{a,1,*}, Nicolas Hanusse^b, David Ilcinkas^b

^aCNR, IEIIT, Turin, Italy

^bCNRS, LaBRI, UMR 5800, F-33400 Talence, France

Abstract

In order to achieve routing in a graph, nodes need to store routing information. In the case of shortest path routing, for a given destination, every node has to store an *advice* that is an outgoing link toward a neighbor. If this neighbor does not belong to a shortest path then the advice is considered as an error and the node giving this advice will be qualified a *liar*. This article focuses on the impact of graph dynamics on the advice set for a given destination. More precisely we show that, for a weighted graph G of diameter D with n nodes and m edges, the expected number of errors after \mathcal{M} edge deletions is bounded by $\mathcal{O}(n \cdot \mathcal{M} \cdot D/m)$. We also show that this bound is tight when $\mathcal{M} = o(n)$. Moreover, for \mathcal{M}' node deletions, the expected number of errors is $\mathcal{O}(\mathcal{M}' \cdot D)$. Finally we show that after a single edge addition the expected number of liars can be $\Theta(n)$ for some families of graphs.

Keywords: graph, network, dynamics, errors, routing information, distance changes

1. Introduction

Everyone has already faced the problem of reaching a destination in an uncertain network. This is typically the case whenever one is in an unknown city,

*Corresponding author

Email address: christian.glacet@gmail.com (Christian Glacet)

¹Partially supported by the ERCIM fellowship program.

without a map, and aims at reaching, let us say, the closest cash machine. The only thing one can do is ask people in the street for some information on the direction. Unfortunately, there is no evidence that all the information one gets is reliable.

Nowadays, in a communication network, a corresponding situation can occur. Let us consider the routing task. Due to its dynamic (change of topology, time required to update local information) and its large-scale size, current networks are not immune to faults and crashes. It is no more realistic to blindly trust the data stored locally at each node. For instance, the Border Gate Protocol (BGP) used in Internet to route messages between autonomous systems implicitly assumes that some paths are known to reach any target. Ideally, these paths are as short as possible. Unfortunately, many messages do not reach their destination because no paths are temporally known although some paths could exist.

In the following, for a given target t , we informally refer to a *liar* as a node containing bad information about the location of t . A series of papers [11, 10, 9] tackles the problem of locating a target (node, resource, data, ...) in presence of liars.

A first model was introduced by Kranakis and Krizanc [15]. They designed algorithms for searching in distributed networks having the ring or the torus topology, when a node has a constant probability of being a liar. A more realistic model was proposed by Hanusse *et al.* [11]: the number of liars is a parameter k and during a routing query, the information stored at every node is unchanged. The main performance measure is the number of edge traversals during a request. Several algorithms, either generic or dedicated to some topologies, and bounds are presented in [11, 10, 9] and are typically of the form $\mathcal{O}(d + k^{\mathcal{O}(1)})$ (for path, grids, expanders, ...) or $\mathcal{O}(k^3 \log^3 n)$ for bounded degree graphs, d being the distance between the source and the target.

In these papers, there is an implicit assumption: the number of liars is small. Our goal is to evaluate whether this is realistic or not. Starting from a network without any liar, we aim at estimating bounds on the number of liars obtained

after few changes of topology. It turns out that this problem is related to the problem of estimating the number of distance changes after few edge/node deletions or insertions. We will show that the expected number of liars created by random nodes/edges removals is proportional to the diameter of the graph and its average degree. Sections 2 and 4 will respectively cover the related works and the presentation of our results, while sections 5, 6 and 7 will show the details of the proofs. Section 5 contains general results, Sections 6 and 7 respectively concern deletions and additions.”

2. Related Works

The influence of topology changes in graphs has been studied in several works. In [4, 18], it has been proven that, in the worst case scenario, the deletion of \mathcal{M} edges within a graph of diameter D induces a multiplication of the diameter by a factor $O(\mathcal{M} + 1)$. A lower bound of $(\mathcal{M} + 1)D - 2\mathcal{M} + 2$ can be obtained, for odd D using the very simple example given in [4] and showed in Figure 1. In this example, the graph G has diameter $D(G) = 6$ and $n = 25$ nodes, the deletion of edges e_1, e_2, e_3 and e_4 induces a graph G' of diameter $D(G') = n - 1 = 24 = (\mathcal{M} + 1) \cdot D(G) - 2\mathcal{M} + 2$. In other words, every edge deletion adds roughly $D(G) = 6$ units to the diameter of the graph.

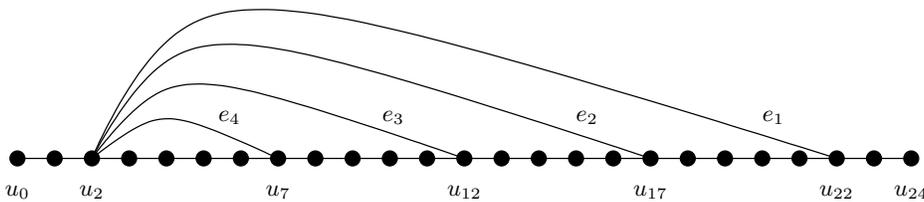


Figure 1: Lower bound example on the impact of dynamics on the diameter from [4].

Our work is also related to the computation of the most vital node of a shortest path [17], that is the node whose removal results in the largest increase of the distance for a given pair of source/target, and the Vickrey pricing of edges [12]. More precisely, the latter article shows a centralized algorithm that

allows to find the most important node for a user that would like to send data over a given graph. This algorithm runs in time $\mathcal{O}(m + n \log n)$ with a working memory of $\mathcal{O}(m)$.

Recently, some works on *dynamic* data structures for shortest paths/distance computation problems have been proposed. By dynamic, we mean that the data structures can tolerate some topology changes in a given network. A dynamic network model defines how the underlying graph changes/evolves over time. More precisely, the following types of models are usually considered:

- The most general model is the model of unconstrained *Evolving graphs* or *Time Varying Graphs*, introduced by Ferreira in [8]: An evolving graph \mathcal{G} is based on a static graph $G = (V, E)$ called *underlying graph*. The state of the graph \mathcal{G} at a given time t , $G_t = (V, E_t)$, is given by a presence function which determines for every edge of E if it belongs to E_t or not. Thus, for every t , the graph G_t is a sub-graph of G . Let remark that for any two given times t and t' the state of the graphs G_t and $G_{t'}$ can be very different, which means that the network can vary at any speed.

This type of unconstrained evolving graphs makes however most of the static graph problems unsolvable. This is why weaker models are most commonly considered. For example, in an unconstrained evolving graph, it is impossible to broadcast a piece of information since the graph is not necessary connected over time, i.e. the union of the G_t is not a connected graph. Different variants of this model can be found in the survey made by Santoro *et al.* in [2].

- However, in many studies the dynamic model studied is even more constrained. A very common model in the routing community is the *failure model*. Intuitively, it is considered in this model that, starting from an initial graph, after \mathcal{M} additions/deletions of nodes/edges the graph stays stable for a sufficiently long period of time. This model is widely used in self-stabilization and in routing or distance oracle problems (also referred as *forbidden-set* problems).

In the case of the failure model, the most naive solution to allow routing is to recompute every routing tables after some dynamic events occur. Two dynamic and centralized algorithms that perform a global re-computation are presented by Demetrescu and Italiano [6] and Thorup [19]. These two algorithms are the fastest known and allow to update every routing tables within amortized time of $\mathcal{O}(n^2 \text{polylog}(n))$ per atomic topological modification.

However, in the failure model, it is not always necessary to recompute every shortest path in order to guarantee that routing is possible. In fact, several other works deal with this problem, and propose for example distributed structures that allow to retrieve either shortest paths or distances within a network where one or several edges/nodes have failed. This kind of distributed structures are of two kinds, they either give exact values (shortest paths or exact distances)[7, 1], or they give constant approximation [13, 3]. The factor of approximation is: (i) a stretch factor of the induced route in the case of routing approximations ; (ii) the ratio between the proposed distance over actual distance within the graph in the case of distance oracles.

The real challenge is to propose an algorithm that can tolerate an arbitrary number of failures. The most general known result is the \mathcal{M} -sensitivity distance oracle proposed by Chechik *et al.* in [3]. This oracle uses a data structure of size $\mathcal{O}(\mathcal{M} \cdot s \cdot n^{1+1/s} \log n)$ and allows to approximate distances for every pair of nodes up to a multiplicative factor $\mathcal{O}(s \cdot \mathcal{M})$, with a query time in $\mathcal{O}(\mathcal{M} \log^2 n)$ (s is a parameter of their oracle).

Throughout [3], all analyzes are made in the worst case, the one of a *strong adversary* that decides the worst forbidden set of nodes/edges, more generally, the worst possible dynamic scenario. This is sometimes a pessimistic vision of dynamics. For example if dynamics in a given network is resulting from material failures, then it is more likely that the adversary is making random choices. Thus, as far as no attacks need to be modeled, it is interesting to have tools to analyze this kind of algorithm in the average case. This is why we study here a model of *random failures*: any sequence of \mathcal{M} topological modifications has the same probability of occurrence. An estimation of the expected number of

distance changes after \mathcal{M} random modifications within a given network could be used for example to give update times of algorithms such as the one presented in [14, section 2.1]. In this work, King shows an algorithm that allows to maintain a shortest path tree up to date with a convergence time of $\mathcal{O}(D \cdot \kappa)$ with κ the number of distances to the root that have changed. The analysis we present here would allow to give the update time of King's algorithm within a random fault model.

3. Model

We consider a weighted and connected graph $G = (V, E)$. We will use two different notions of distance, the classical one: $d_G(u, v)$, and the *hop-distance* which is equal to the minimum number of edges traversal that are required to travel from u to v via a shortest path. The hop-distance is denoted $hd_G(u, v)$, finally, the *hop-diameter* is equal to the maximal hop-distance over every pair of nodes.

3.1. Advice-Liar model

A particular node t of V is called *destination*, and every other node $u \in V \setminus \{t\}$ has an *advice* $\text{Adv}(u)$ that points toward a neighbor node:

$$\text{Adv}(u) \in (\Gamma(u) \cup \{u\})$$

with $\Gamma(u)$ the set of neighbors of node u . If the node $\text{Adv}(u)$ is on a shortest path from u to t then node u is said to be *truthful*, otherwise it is called a *liar*. The set of all pieces of advice (a set of arcs) \mathbb{A} of cardinality $n - 1$, called a *configuration*, induces an oriented sub-graph of G , noted $G_{\mathbb{A}} = (V, \mathbb{A})$. If the configuration of advice does not contain any liar, then it is said to be a *truthful configuration*. For any two nodes u and v in V , there exists an arc (u, v) in $G_{\mathbb{A}}$ ($(u, v) \in \mathbb{A}$) if and only if $\text{Adv}(u) = v$. Whenever \mathbb{A} is truthful, the graph $G_{\mathbb{A}}$ is a shortest path tree rooted at node t that covers V .

3.2. Dynamic Model

In this study, we will consider separately the case of deletions and additions. The number of edges, respectively nodes, added or deleted from G is \mathcal{M} (respectively \mathcal{M}'). We will consider that the number of deletions/additions is dominated by the number of elements in G , $\mathcal{M} = o(m)$ and $\mathcal{M}' = o(n)$. Moreover, since the reasoning for nodes and edges cases are very similar we will only present the problem for edges in detail.

For a given graph $G = (V, E)$, the set of graphs $\mathcal{G}_{G, \mathcal{M}}^-$ is defined as the set of graphs that can be obtained after \mathcal{M} edge deletions in graph G :

$$\forall G^- = (V, E^-) : \left[G^- \in \mathcal{G}_{G, \mathcal{M}}^- \right] \iff [(E^- \subset E) \wedge (|E^-| = |E| - \mathcal{M})]$$

Similarly we define $\mathcal{G}_{G, \mathcal{M}}^+$ to be the set of graphs that can be obtained after \mathcal{M} edge additions:

$$\forall G^+ = (V, E^+) : \left[G^+ \in \mathcal{G}_{G, \mathcal{M}}^+ \right] \iff [(E^+ \supset E) \wedge (|E^+| = |E| + \mathcal{M})]$$

Given the graphs G , and $G' \in (\mathcal{G}_{G, \mathcal{M}}^- \cup \mathcal{G}_{G, \mathcal{M}}^+)$, together with a truthful configuration \mathbb{A} , we define the following sets:

- $\mathcal{S}(G, G')$ as the set of nodes whose distance to t is different in G and G' . This will be denoted \mathcal{S} if there is no ambiguity;
- $\mathcal{K}(G', \mathbb{A})$ as the set of liars in G' . This will be denoted \mathcal{K} if there is no ambiguity. The size of this set will often be denoted k .

The purpose of this paper is to analyze these sets, and in particular their size.

Remark. It is possible that some pieces of advice from \mathbb{A} had to be changed after edge deletions. Indeed, for two nodes u and v , if $(u, v) \in \mathbb{A}$ and edge $\{u, v\}$ is deleted, then within G^- node u has a piece of advice that points toward a non-neighbor node. Therefore, since we consider that every node has to point toward a neighbor in its current graph, then u will have to pick a new advice among its neighbors. This new choice will depend on the model of adversary

we will consider (adversary models will be discussed next). Also, if its neighborhood is empty then its advice will be a self loop $\text{Adv}(u) = u$ whatever the model is.

It is also possible that graph G^- is disconnected. In this case, every node that does not belong to the same component as t is considered a liar.

Adversary model. We will study the number of liars and distance changes under two models of adversary:

- **Strong adversary model.** This adversary models the worst case scenario, which can be seen as an attack on the network by an omniscient entity. For a given graph $G = (V, E)$, the strong adversary coordinates the choice of a truthful set of advice \mathbb{A} among all possible sets of advice, together with the choice of the \mathcal{M} edges to remove from/add to G with the objective of maximizing the number of liars $k(G^-, \mathbb{A})$ (respectively $k(G^+, \mathbb{A})$). Moreover, for any node $u \in V$, if $\text{Adv}(u) = v$ and $\{u, v\} \notin E^-$, then this adversary can redirect the advice of node u to any node in the neighborhood of u in E^- .
- **Random adversary.** This model can be used to analyze the expected number of errors created by a set of random faults that happened in the network. For a given graph $G = (V, E)$ and a given target node t , the random adversary chooses uniformly at random a set of advice \mathbb{A} among all possible truthful sets of advice. It also chooses uniformly at random a set of \mathcal{M} distinct edges within E for deletions or within $V^2 \setminus E$ for additions. In other words it chooses a graph G' uniformly at random within $\mathcal{G}_{G, \mathcal{M}}^-$ for deletions (within $\mathcal{G}_{G, \mathcal{M}}^+$ for additions). Moreover, for any node $u \in V$, if $\text{Adv}(u) = v$ and $\{u, v\} \notin E^-$, then this adversary redirects the advice of node u to a node chose uniformly at random within the neighborhood of u in graph the G^- .

4. Contributions

4.1. Under the strong adversary

Under the strong adversary model the number of liars created is very high even for one single edge or node deletion and without graph disconnection:

Theorem 4.1. *Under the strong adversary, for one node or edge deletion (or addition), the numbers of liars and distance changes can be at least $n - D - \mathcal{O}(1)$ even if the graph remains connected.*

Since the proof of this theorem is quite simple and short we will use it as a warm-up. This can be useful to better understand the model of dynamics and what is measured.

Proof. The proof is based on Figures 2 and 3 that show lower bounds for the deletion and addition cases respectively. The proof is split in two parts:

Deletion case:. Take any unweighted graph $G = (V, E)$, of diameter D , and connect every node of this graph to the extremities a and b of two paths that both end in a new node t , one path A of length D and the other, B , of length $D - 1$. In this graph H , every node $v \in V$ is at distance D from t and has its shortest path to t going through path B . Therefore, every such node v has its advice pointing to the first node of path B , node b . Let us now consider the graph H^- which is graph H minus the closest edge from t on path B , edge $\{x, t\}$. This graph has diameter $D + 1$, the advice of every node $v \in V$ is necessary a lie since it goes to a dead-end. Moreover, every node from path B , except x , turns into a liar. Therefore the number of liars in this graph is $n - |A| - |\{x\}| = n - D - 1$.

Addition case:. Consider the graph H^- described in the deletion case. In this graph H^- , every node $v \in V$ is at distance $D + 1$ from t and has its shortest path to t going through path A . Therefore, every such node v has its advice pointing to the first node of path A , node a and every node from B has its advice pointing toward G (the sub-graph used to construct H^- earlier). Let us

now consider the graph H which is graph H^- plus one edge connecting path B back to t , adding back the edge $\{x, t\}$ (this graph is the same H as the one from the deletion case). It has diameter D , the advice of every node $v \in V$ is a lie since the path from v to t via path A is one unit longer than the one going through B . Moreover, every node from path B turned to a liar for the same reason. In other words every node except nodes from A turn into a liar. Therefore the number of liars in this graph can be up to $n - |A| = n - D$.

□

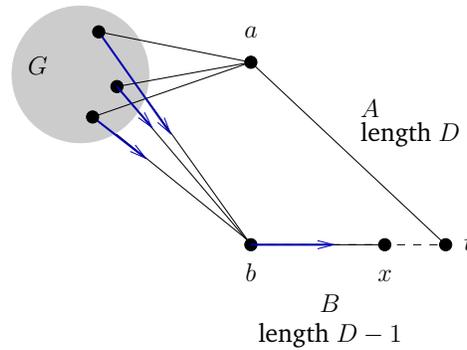


Figure 2: An example of an edge deletion (edge $\{x, t\}$) that creates $n - (D + 1)$ liars and distance changes.

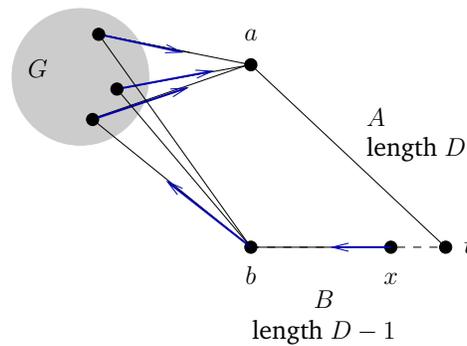


Figure 3: Example of a single edge addition (edge $\{x, t\}$) that creates $n - D$ liars and distance changes.

This means that in the general case, many liars can appear after an attacker removes (adds) a single communication link. That is why the rest of the paper focuses on random deletions (additions), that can be seen as a model of failure, where the probability for any node or link to fail is the same.

4.2. Under the random adversary

We will first present an upper bound on the expected number of liars after \mathcal{M} edge deletions, then we will show that this upper bound is tight for some family of graphs, even if the graph stays connected after the deletions.

Theorem 4.2. *For any weighted graph $G = (V, E)$ of hop-diameter D with n nodes and m edges, for a randomly located target, after $\mathcal{M} \leq \frac{m}{2}$ random edge deletions, the expected number of liars is $\mathbb{E}(k) \leq \mathcal{M} \cdot \frac{nD}{m}$. Moreover, the expected number of distance changes is also upper-bounded by $\mathcal{M} \cdot \frac{nD}{m}$. Finally the number of liars and distance changes after \mathcal{M}' node deletions is at most $\mathcal{M}'D$.*

As we said before, this kind of upper bound allows one to know, for example, the average update time of King’s algorithm, which would be $\mathcal{O}(D^2\mathcal{M} \cdot \frac{n}{m})$ after \mathcal{M} random edge deletions and $\mathcal{O}(D^2\mathcal{M}')$ after \mathcal{M}' random node deletions. More generally we can compute the number of errors for certain graphs or family of graphs. If we look at the internet autonomous system network in 2007, which had, according to DIMES maps [21], 17144 nodes, 46621 links and diameter 8, the expected number of liars created for every falling link is about 2.95 ($\simeq 8 \times \frac{17}{47}$). More generally, many natural and artificial complex networks have scale-free properties. These networks are often modeled by random graphs that have constant average degree and logarithmic diameter[16] while some of them even have a smaller diameter[5]. Thus, in most common scale-free networks[20] (transportation, world-wide web, AS relationships, Math. co-authorship, ...) every random edge deletion creates $\Theta(\log n)$ liars are created on average.

Theorem 4.3. *For every unweighted \mathcal{M} -edge-connected graph $G = (V, E)$ of diameter D and such that $\mathcal{M} \leq \frac{n}{2}$, there exists a unweighted \mathcal{M} -edge-connected H*

of diameter $D + \Theta(1)$ with $n = \Theta(|V|)$ nodes, $m = \Theta(|E|)$ edges and having G as a sub-graph, in which the expected number of liars after \mathcal{M} deletions under the random adversary model is $\Omega(\mathcal{M} \cdot (n - \mathcal{M}) \frac{D}{m})$.

Remark that the bounds from [Theorem 4.2](#) and [Theorem 4.3](#) are tight for $\mathcal{M} \leq \frac{n}{2}$.

Secondly, we have considered the case of random edge additions and actually found that considering a single edge addition can already have a huge impact on distances and information stored at nodes.

Theorem 4.4. *There exists a family of n -node unweighted graphs in which the expected number of liars and distance change created by a single edge addition is $\Theta(n)$.*

This looks as bad news, but liars created by edge addition(s) do not create any cycle of advice, but simply miss new shortcuts. This means that routing would still be guaranteed whereas in the case of edge or node deletions routing can be compromised.

5. First Observations

5.1. Useful tool

The estimation of the number of distance changes and the number of liars is correlated with the probability of disconnecting a given path within a graph as we will see later on. The following Lemma and Corollaries concern path disconnections and will be used to prove both the lower and the upper bounds given in [Section 4.2](#).

Lemma 5.1 (Probability of disconnecting a given path). *For any path P composed of ℓ edges, within a graph $G = (V, E)$ such that $|V| = n$ and $|E| = m$, the*

probability of disconnecting P by deleting \mathcal{M} random edges from E is equal to²:

$$\sum_{i=0}^{\mathcal{M}-1} \left(\frac{\ell}{m-i} \prod_{j=0}^{i-1} \frac{m-j-\ell}{m-j} \right) \quad (1)$$

Proof. In the following, the event E_i is "the i -th edge deletion is the first deletion that disconnects the path P ". Also, the event C_i is "the i -th deleted edge belongs to P ". The probability p of disconnecting path P is equal to:

$$\begin{aligned} p &= \sum_{i=1}^{\mathcal{M}} \Pr(E_i) \\ &= \sum_{i=1}^{\mathcal{M}} \Pr(\overline{C_1} \wedge \overline{C_2} \wedge \dots \wedge \overline{C_{i-1}} \wedge C_i) \\ &= \sum_{i=0}^{\mathcal{M}-1} \left(\frac{\ell}{m-i} \prod_{j=0}^{i-1} \frac{m-j-\ell}{m-j} \right) \end{aligned}$$

□

Corollary 5.1 (General upper bound). *For any path P composed of ℓ edges, within a graph $G = (V, E)$ such that $|V| = n$ and $|E| = m$, the probability of disconnecting P by deleting $\mathcal{M} \leq \frac{m}{2}$ random edges from E is upper-bounded by $\frac{2\ell}{m} \cdot \mathcal{M}$.*

Proof. From [Lemma 5.1](#) this probability is equal to equation (1):

$$\sum_{i=0}^{\mathcal{M}-1} \left(\frac{\ell}{m-i} \prod_{j=0}^{i-1} \frac{m-j-\ell}{m-j} \right).$$

Since $\mathcal{M} \leq \frac{m}{2}$, we have $\frac{\ell}{m-i} \leq \frac{2\ell}{m}$ for any $i \in [0, \mathcal{M} - 1]$. The product from equation (1) is always smaller than 1, therefore we immediately obtain the upper bound $\frac{2\ell}{m} \cdot \mathcal{M}$. □

Corollary 5.2 (Tightness of the bound). *For any path P composed of $\ell \leq \frac{m}{4}$ edges, within a graph $G = (V, E)$ such that $|V| = n$ and $|E| = m$, the probability of disconnecting P by deleting $\mathcal{M} \leq \frac{m}{2}$ random edges from E is in $\Theta\left(\frac{\ell}{m} \cdot \mathcal{M}\right)$ when $\frac{\ell}{m} \cdot \mathcal{M} \in \mathcal{O}(1)$.*

²Notice that when $i=0$ the value of the product is equal to 1.

Proof. From [Lemma 5.1](#) and the inequality $\mathcal{M} \leq \frac{m}{2}$ we have:

$$\begin{aligned} p &= \sum_{i=0}^{\mathcal{M}-1} \left(\frac{\ell}{m-i} \cdot \prod_{j=0}^{i-1} 1 - \frac{\ell}{m-j} \right) \\ &\geq \sum_{i=0}^{\mathcal{M}-1} \left(\frac{\ell}{m} \cdot \prod_{j=0}^{i-1} 1 - \frac{2\ell}{m} \right) \\ &\geq \frac{\ell}{m} \cdot \sum_{i=0}^{\mathcal{M}-1} \left(1 - \frac{2\ell}{m} \right)^i \end{aligned}$$

since by hypothesis $\mathcal{M} \leq \frac{m}{\ell} \cdot c$ for some constant $c > 0$, we have:

$$p \geq \frac{\ell}{m} \cdot \sum_{i=0}^{\mathcal{M}-1} \left(1 - \frac{2\ell}{m} \right)^{\frac{m}{2\ell} \cdot c}$$

The function $\left(1 - \frac{1}{x}\right)^x$ is non-decreasing on the interval $[2, +\infty)$ and has a value of $1/4$ for $x = 2$. Therefore, in this interval, for $c' = 1/4$, we have $\left(1 - \frac{1}{x}\right)^x \geq c'$. In our case x actually belongs to the given interval since $\ell \leq \frac{m}{4}$. As a consequence, we have:

$$p \geq \frac{\ell}{m} \cdot \sum_{i=0}^{\mathcal{M}-1} c'^c$$

This implies that p is in $\Omega\left(\frac{\ell}{m} \cdot \mathcal{M}\right)$. Combining this with [Corollary 5.1](#), we obtain the claimed result. \square

Before using these corollaries to analyze the number of errors we will characterize what makes truthful nodes turn into liars.

5.2. How truthful nodes turn into liars

First, notice that edge deletions do not necessary imply the creation of liars, even if these deletions have an impact on the distances in the graph. For example, in an unweighted complete graph with $n > 2$, no liar can be created by a single edge deletion (u, v) :

- if $v = t$, then $\text{Adv}(u)$ in G is v and the distance from u to v in G^- is 2, thus every neighbor is closer to v than u and is therefore a truthful advice for u in G^- ;

- if $v \neq t$, then $\text{Adv}(u) = t$ in both graphs and u is not a liar.

Nevertheless there exists a relationship between distance changes and the creation of liars. To define this relationship, let us consider two graphs $G = (V, E)$ and $G^- = (V, E^-) \in \mathcal{G}_{G, \mathcal{M}}^-$ together with the truthful set of advice on G , \mathbb{A} .

Lemma 5.2. *If a node $u \in V$ is a liar in $G^- \in \mathcal{G}_{G, \mathcal{M}}^-$ then either its advice points toward a node v whose distance to the destination has changed or its advice has been changed due to a deletion. More formally:*

$$\forall (u, v) \in \mathbb{A} : u \text{ is a liar within } G^- \implies (v \in \mathcal{S}(G, G^-)) \vee (\{u, v\} \notin E^-)$$

Proof. Consider two nodes u and v such that $v \notin \mathcal{S}(G, G^-)$ and $\{u, v\} \in E^-$. From both hypotheses, we know that the distance from u to t is the same in both G and G^- . Since the set \mathbb{A} of advice is truthful in G , and because a distance cannot decrease from G to G^- , the advice of u is still correct. □

We are now ready to show the impact of random edge deletions on the number of liars.

6. Edge Deletions Under the Random Adversary

As [Lemma 5.2](#) shows liars come either directly from the deletion of the edge leading to their advice or from distance changes. We will therefore start by analyzing the number of distance change created by \mathcal{M} random edge deletions on a graph G with advice set \mathbb{A}_G . We will denote \mathbb{A} the modified advice set on the resulting graph G^- (recall that some illegal pieces of advice needed to be randomly redirected in this model).

6.1. Impact on the number of distance changes

Under the random adversary, for a given node, the probability of changing its distance after \mathcal{M} edge deletions can be bounded using the upper bound on the probability of disconnecting a given path ([Corollary 5.1](#)):

Lemma 6.1. For any graph $G = (V, E)$, and any graph $G^- = (V, E^-)$ chosen uniformly at random in $\mathcal{G}_{G, \mathcal{M}}^-$, with $\mathcal{M} \leq \frac{m}{2}$, the probability that a node $u \in V$ has a different distance in G and G^- is bounded as follows:

$$\Pr[u \in \mathcal{S}] \leq \mathcal{M} \cdot \frac{hd_G(u, t)}{m}$$

Proof. Distances $d_G(u, t)$ and $d_{G^-}(u, t)$ are different if and only if for every shortest path $P = \{u_0 = u, u_1, \dots, u_\ell = t\}$ at least one edge $\{u_i, u_{i+1}\}$ does not belong to E^- . The probability of this happening is upper-bounded by the probability of deleting at least one edge in any given shortest path from u to t . In particular, it is upper-bounded by the probability of deleting at least one edge in the shortest path with the minimum number of hops, which is, from [Corollary 5.1](#), less than or equal to $\frac{hd_G(u, t)}{m} \cdot \mathcal{M}$. Therefore we have:

$$\Pr[d_G(u, t) \neq d_{G^-}(u, t)] \leq \mathcal{M} \cdot \frac{hd_G(u, t)}{m}$$

□

The upper bound on the number of distance changes given by [Theorem 4.2](#) can be derived from this Lemma, by bounding for every node u its distance to t in G by the hop-diameter D of G .

Proof of [Theorem 4.2](#), the part on $|\mathcal{S}|$. Let us denote by X_u the random variable that has value 1 if $u \in \mathcal{S}$ and 0 otherwise. Since expectation is linear, the size of the set \mathcal{S} satisfies:

$$\begin{aligned} \mathbb{E}(|\mathcal{S}|) &= \sum_{u \in V} \mathbb{E}(X_u) \\ &= \sum_{u \in V} \Pr(u \in \mathcal{S}) \end{aligned}$$

Thus the expectation of the number of distance changes after \mathcal{M} edge deletions is:

$$\begin{aligned} \mathbb{E}(|\mathcal{S}|) &\leq \sum_{u \in V} \mathcal{M} \cdot \frac{hd_G(u, t)}{m} \\ &\leq \mathcal{M} \cdot \frac{nD}{m} \end{aligned}$$

□

6.2. Impact on the number of liars

The following Lemma can be deduced from Lemmas 6.1 and 5.2:

Lemma 6.2. Any node $u \in V$ is a liar in graph G^- with a probability at most:

$$\Pr[u \in K(G^-, \mathbb{A})] \leqslant hd_G(u, t) \cdot \frac{\mathcal{M}}{m}$$

Proof. From Lemmas 6.1 and 5.2 it can be deduced that, for a given node $u \in V$ such that $(u, v) \in \mathbb{A}$, the probability for a node u of being a liar can be bounded as follows:

$$\begin{aligned} \Pr[u \in K(G^-, \mathbb{A})] &\leqslant \Pr[v \in \mathcal{S}] + \Pr[\{u, v\} \notin E^-] \\ &\leqslant hd_G(v, t) \cdot \frac{\mathcal{M}}{m} + \frac{\mathcal{M}}{m} \\ &\leqslant (hd_G(u, t) - 1) \cdot \frac{\mathcal{M}}{m} + \frac{\mathcal{M}}{m} \end{aligned}$$

Finally:

$$\Pr[u \in K(G^-, \mathbb{A})] \leqslant hd_G(u, t) \cdot \frac{\mathcal{M}}{m}$$

□

From this Lemma we can show the upper bound on the total number of liars, given by [Theorem 4.2](#).

Proof of Theorem 4.2, the part on K . Let us set X'_s to be the random variable that has value 1 if s is a liar in G^- and 0 otherwise. From [Lemma 6.2](#) we can deduce that the expected number of liars in G^- is:

$$\begin{aligned} \mathbb{E}(|K(G^-, \mathbb{A})|) &= \mathbb{E}\left(\sum_{s \in V \setminus \{t\}} X'_s\right) \\ &= \sum_{s \in V \setminus \{t\}} \mathbb{E}(X'_s) \\ &= \sum_{s \in V \setminus \{t\}} \Pr[u \in K(G^-, \mathbb{A})] \\ &\leqslant \sum_{s \in V \setminus \{t\}} hd_G(s, t) \cdot \frac{\mathcal{M}}{m} \\ &\leqslant \sum_{s \in V \setminus \{t\}} D \cdot \frac{\mathcal{M}}{m} \end{aligned}$$

This gives the following upper bound on the number of liars:

$$\mathbb{E}(|K(G^-, \mathbb{A})|) \leq \mathcal{M} \cdot \frac{Dn}{m}$$

and concludes the proof of the second part of [Theorem 4.2](#). □

Remark on the case of node deletions. The expected number of liars created by \mathcal{M}' node deletions can be analyzed in the exact same way. First it can be shown that the probability of disconnecting every shortest path from a node u to a node t is upper bounded by $\mathcal{M}' \cdot \frac{hd_G(u,t)}{n}$. Then, from that, the bound on the number of distance changes $\mathcal{M}'D$ can be obtained, which finally gives that the number of liars is also upper bounded by $\mathcal{M}'D$.

6.3. Lower Bound

This section handles the proof of [Theorem 4.3](#), which states that the lower bound on the number of liars after \mathcal{M} random edge deletions can be up to $\Omega(\mathcal{M} \cdot \frac{nD}{m})$ for some family of small-diameter unweighted graphs that we will define right away and show on [Figure 4](#).

Graph settings. It is desirable to have a lower bound that is not based on graph disconnections. In order to guarantee that, in the random model, it is needed to use a \mathcal{M} -connected graph, so that \mathcal{M} deletions would never disconnect the graph. For that purpose we will use a gadget-graph parametrized by its diameter, δ , and its connectivity, κ . We denote this graph by $P_{\delta,\kappa}$. It is composed of κ paths of length δ where, for every $i \in [1, \delta]$, the i -th node of every path is connected to the $(i + 1)$ -th node of every other paths. For every $i \in [1, \delta + 1]$ and every $j \in [1, \kappa]$, we will denote the i -th node of the j -th path by $u_{i,j}$. The construction of the graph, H , used for the lower bound is based on two copies, $G = (V, E)$ and $G' = (V', E')$, of any \mathcal{M} -edge-connected graph. Every node of graph G (respectively G') is connected to two nodes a and b (respectively a' and b'). Nodes a and a' are linked by a path $X = \{x_0 = a, x_1, \dots, x_{D-1}, x_D = a'\}$ of length D . Also, node b and b' are linked by a graph $P = P_{D,\mathcal{M}}$ as follows: for

every $j \in [1, \mathcal{M}]$, an edge is added between nodes b and $u_{1,j}$ and between nodes b' and $u_{D+1,j}$. Finally every node x_i from path X is connected to \mathcal{M} nodes, and each of these node is connected, for every $j \in [1, \mathcal{M}]$, to node $u_{i+1,j}$. If D denotes the diameter of G , then graph H has:

- diameter $D + \Theta(1)$;
- $n = 2|V| + \Theta(D\mathcal{M})$ nodes;
- and $m = 2|E| + 4|V| + \Theta(D\mathcal{M}^2)$ edges.

However notice that, since G is \mathcal{M} -connected, $m \geq n\mathcal{M}$ and $n \geq (D - 1)\mathcal{M}$, thus $m \geq (D - 1)\mathcal{M}^2$. Therefore, we have $n = \Theta(|V|)$ and $m = \Theta(|E|)$.

Target location and advice set.. Since the graph H is symmetric, we can consider, without loss of generality, that the target node t is located within graph G' . Since the original configuration has to be truthful, every node u that belongs to G has its advice pointing toward node a .

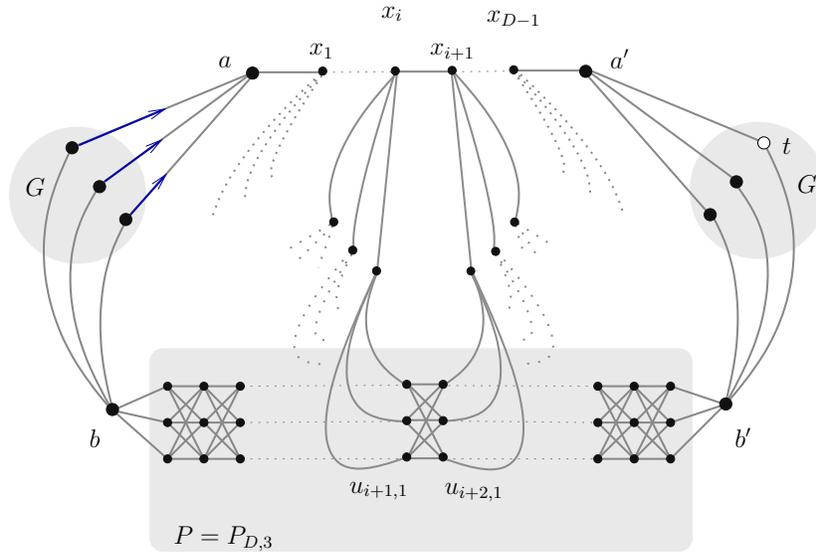


Figure 4: Graph used in the lower bound in the case of edge deletions.

Impact of \mathcal{M} random edge deletions. Let us now focus on the case where an edge is deleted in the path X that links node a and a' , this case is depicted in [Figure 5](#). If one of these edges is deleted then a shortest path A , from a to t , has the form $a \rightsquigarrow x_i \rightsquigarrow u_{i+1,\bullet} \rightsquigarrow b'$ with $\{x_i, x_{i+1}\}$ being the first edge that has been deleted from X . Whereas a shortest path B , from b to t , has the form $b \rightsquigarrow u_{i+1,\bullet} \rightsquigarrow b'$. Notice that such paths A and B always exist since the graph H is \mathcal{M} -connected. Since $d(b, u_{i+1,\bullet}) = i + 1$ and $d(a, x_i) + d(x_i, u_{i+1,\bullet}) = i + 2$, the path A is longer than the path B , therefore any node from G that is still connected to b is a liar since its advice point toward node a . If such an event occur, then the number of liars created is at least $|V| - (\mathcal{M} - 1)$, that is in $\Omega(n - \mathcal{M})$ when $\mathcal{M} \leq \frac{|V|}{2}$. Now, the probability that such an event occurs is equal to the probability of deleting at least one edge in the path X , which is, from [Corollary 5.2](#), in $\Omega(\frac{D}{m}\mathcal{M})$. Therefore the expected number of liars is in $\Omega(\mathcal{M} \cdot (n - \mathcal{M}) \cdot \frac{D}{m})$, which concludes the proof of [Theorem 4.3](#).

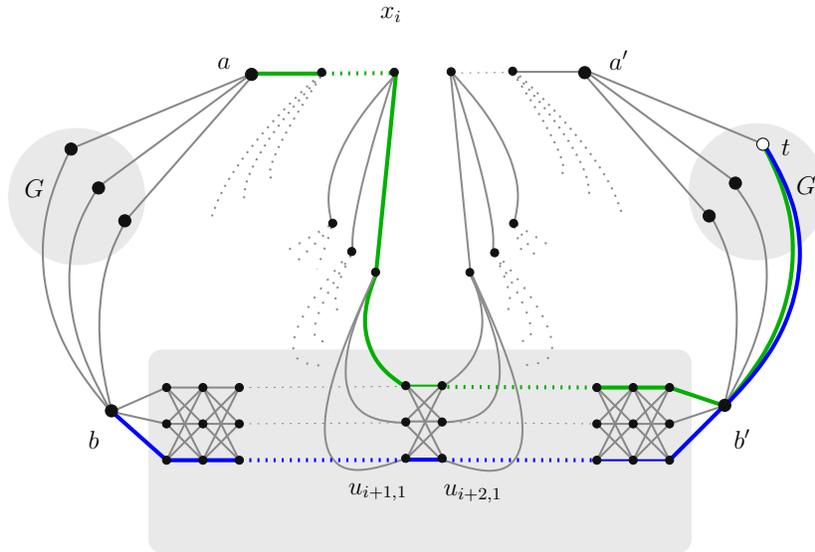


Figure 5: The case where an edge of the upper path has been removed.

Remark. If we do not want to restrain the analysis to \mathcal{M} -edge-connected graphs,

then the graph that connects b to b' can be changed to a line of length $D+2$, and connections between the path from a to a' and this new path become useless.

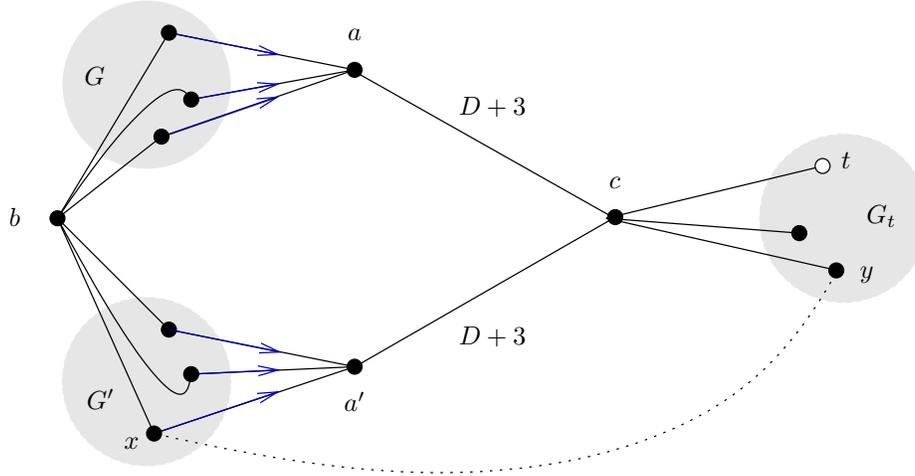


Figure 6: An $\Omega(n)$ lower bound for the edge addition.

7. Edge Addition Under the Random Adversary

We will show in this section that with the random adversary model, the expected number of liars after one single edge addition is $\Omega(n)$ for some family of graphs.

In order to prove the lower bound, we consider a graph H with $\mathcal{O}(n)$ nodes and diameter $\mathcal{O}(D)$ whose construction is based on three copies of any given n -node graph of diameter D , namely G, G' and G_t . Every node from graph G (respectively G') is connected to two new nodes a and b (respectively a' and b). Also, every node that belongs to graph G_t is connected to a fourth node, named c . The graph H is then built by connecting node a (respectively a') to node c by a path of length $D+3$. Figure 6 shows the graph H with some indications on node positions that will be used and detailed right after.

Let us now observe the configuration of advices in H when the destination node t belongs to graph G_t . Any node u that belongs to G has two possible ways

for exiting graph G in order to get closer to t , either through node a or node b .

The distances $d_H(a, t)$ and $d_H(b, t)$ have the following values in graph H :

$$d_H(a, t) = d_H(a, c) + 1 = D + 4$$

$$d_H(b, t) = 2 + d_H(a', t) = 2 + d_H(a, t) = D + 6$$

Since $d_H(a, t) < d_H(b, t)$ the shortest path from any node of graph G to t in H goes through node a . Therefore every node from G will have its advice pointing toward node a . Let us now add an edge $\{x, y\}$ to graph H and call this new graph H^+ . The edge $\{x, y\}$ has one of its extremities in G' with constant probability and the other extremity, y , belongs to G_t also with constant probability. Some distances have been changed by this edge addition, in particular there is a shortcut $b \rightsquigarrow x \rightsquigarrow y \rightsquigarrow c \rightsquigarrow t$ which implies that:

$$d_{H^+}(b, t) = \min(4, d_H(b, t)) = \min(4, D + 6) = 4$$

Thus $d_{H^+}(b, t) < d_{H^+}(a, t)$ and therefore, for any node u that belongs to G , the shortest path to node t in the graph H^+ goes through node b . Every node from G is then turned into a liar by adding the edge $\{x, y\}$. These nodes represent a constant proportion of nodes in graph H . Therefore the expected number of liars for a single random edge addition is in $\Theta(n)$, which concludes the proof of [Theorem 4.4](#).

8. Conclusion

This study shows that, in the context of routing, links and routers failures usually creates very few errors. More precisely it shows that the number of errors grows linearly (at most) with the diameter. It is promising in the case of dynamic routing since most networks where routing protocols are used have very small diameter. This result could be used to design dynamic protocols that decide to update their routing tables based on the tradeoff between the cost of update and the cost of dealing with existing errors. Such a protocol could make these decision using an estimated number of errors giving (local) observations on the probability of failure. In this context it would also be interesting

to have a better understanding of the appearance of liars. For example, it can be important to know that some deletions create no liar at all. This actually happens with a probability $p_0 \geq 1 - n/m$, which is the probability of deleting no advice (and therefore preserving all distances). Nevertheless it is not trivial to give a tight bound on this probability, since it is possible that no liars are created when deleting an advice. This question would allow to know the average number of liars created conditioned on the event of creating at least one liar, and this second expectation $\mathbb{E}(k \mid \text{at least 1 liar})$ could be much larger than the overall expected number of liars for some graphs. We can, using p_0 , state that $\mathbb{E}(k \mid \text{at least 1 liar}) \geq \mathbb{E}(k) \cdot \frac{m}{n}$. It could be interesting to give a non trivial upper bound on this last expectation but it is not clear if it would be possible. More generally, it could be interesting to have a distributed algorithm estimating the amount of liars to decide if an update mechanism should be executed.

Bibliography

- [1] Bernstein, A., Karger, D. R., 2009. A nearly optimal oracle for avoiding failed vertices and edges. In: Mitzenmacher, M. (Ed.), STOC. ACM, pp. 101–110.
- [2] Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N., 2012. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems* 27 (5), 387–408.
- [3] Chechik, S., Langberg, M., Peleg, D., Roditty, L., 2010. f -sensitivity distance oracles and routing schemes. In: Proceedings of the 18th annual European conference on Algorithms: Part I. ESA'10. Springer-Verlag, Berlin, Heidelberg, pp. 84–96.
URL <http://portal.acm.org/citation.cfm?id=1888935.1888946>
- [4] Chung, F. R. K., Garey, M. R., 1984. Diameter bounds for altered graphs. *Journal of Graph Theory* 8 (4), 511–534.
URL <http://dx.doi.org/10.1002/jgt.3190080408>

- [5] Cohen, R., Havlin, S., Feb 2003. Scale-free networks are ultrasmall. *Phys. Rev. Lett.* 90, 058701.
URL <http://link.aps.org/doi/10.1103/PhysRevLett.90.058701>
- [6] Demetrescu, C., Italiano, G. F., 2004. A new approach to dynamic all pairs shortest paths. *J. ACM* 51 (6), 968–992.
- [7] Demetrescu, C., Thorup, M., Chowdhury, R. A., Ramachandran, V., January 2008. Oracles for distances avoiding a failed node or link. *SIAM J. Comput.* 37, 1299–1318.
URL <http://portal.acm.org/citation.cfm?id=1350684.1350685>
- [8] Ferreira, A., Jarry, A., 2004. Complexity of minimum spanning tree in evolving graphs and the minimum-energy broadcast routing problem. In: *Proceedings of WiOpt*. Vol. 4.
- [9] Hanusse, N., Ilcinkas, D., Kosowski, A., Nisse, N., 2010. Locating a Target with an Agent Guided by Unreliable Local Advice. In: *Proceedings of the 29th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing PODC 2010*. ACM New York, NY, USA, Zurich Suisse, pp. 355–364.
URL <http://hal.archives-ouvertes.fr/hal-00516695/en/>
- [10] Hanusse, N., Kavvadias, D. J., Kranakis, E., Krizanc, D., 2008. Memoryless search algorithms in a network with faulty advice. *Theor. Comput. Sci.* 402 (2-3), 190–198.
- [11] Hanusse, N., Kranakis, E., Krizanc, D., 2004. Searching with mobile agents in networks with liars. *Discrete Applied Mathematics* 137, 69–85.
- [12] Hershberger, J., Suri, S., 2001. Vickrey prices and shortest paths: What is an edge worth? In: *FOCS*. pp. 252–259.
- [13] Khanna, N., Baswana, S., 2010. Approximate shortest paths avoiding a failed vertex: Optimal size data structures for unweighted graphs. In: *Mar-*

- ion, J.-Y., Schwentick, T. (Eds.), STACS. Vol. 5 of LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 513–524.
- [14] King, V., 1999. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In: FOCS. pp. 81–91.
- [15] Kranakis, E., Krizanc, D., 1999. Searching with uncertainty. Proc. 6th Int. Colloq. on Structural Information and Communication Complexity (SIROCCO'99), 194–203.
- [16] Lu, L., 2001. The diameter of random massive graphs. In: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms. Society for Industrial and Applied Mathematics, pp. 912–921.
- [17] Nardelli, E., Proietti, G., Widmayer, P., March 2003. Finding the most vital node of a shortest path. Theor. Comput. Sci. 296, 167–177.
URL <http://portal.acm.org/citation.cfm?id=763698.763710>
- [18] Schoone, A., Bodlaender, H., van Leeuwen, J., 1987. Improved diameter bounds for altered graphs. In: Tinhofer, G., Schmidt, G. (Eds.), Graph-Theoretic Concepts in Computer Science. Vol. 246 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 227–236.
- [19] Thorup, M., 2004. Fully-dynamic all-pairs shortest paths: Faster and allowing negative cycles. In: Algorithm Theory - SWAT 2004, 9th Scandinavian Workshop on Algorithm Theory. pp. 384–396.
- [20] Wang, X. F., Chen, G., 2003. Complex networks: small-world, scale-free and beyond. Circuits and Systems Magazine, IEEE 3 (1), 6–20.
- [21] Yuval Shavitt, Noam Koenigstein, U. W. N. Z., 2009. DIMES maps - as relationships. <http://www.netdimes.org/new/>, last visit : 2013-07-18.