Optimal Routing in Gossip Networks

Yuval Shavitt and Amir Shay
Tel Aviv University
{shavitt,amirshay}@eng.tau.ac.il

TECHNICAL REPORT: EES2003-4

30 June 2003

Abstract

In this paper we introduce the *Gossip Network* model where travelers can obtain information about the state of dynamic networks by gossiping with peer travelers using ad-hoc communication. Travelers then use the gossip information to recourse their path and find the shortest path to destination. We study optimal routing in stochastic, time independent networks, and demonstrate that an optimal routing policy may direct travelers to make detours to gather information. A dynamic programming equations that produce the optimal policy for routing in gossip networks is presented. In general the dynamic programming algorithm is intractable, however for two special cases a polynomial optimal solution is presented.

We show that usually gossiping helps travelers decrease their expected path cost. However in some scenarios, depending on the network parameters, gossiping could increase the expected path cost. The parameters that determine the effect of gossiping on the path costs are identified and their influence is analyzed. This dependency is fairly complex and was confirmed numerically on grid networks.

1 Introduction

Optimal routing in both deterministic and stochastic networks has been extensively studied in the past. While the solutions for the deterministic problem are well known [BG92] and based on the dynamic programming (Bellman-Ford) or label correcting (Dijkstra) algorithms, the solution to the stochastic problem depends profoundly on the problem modelling. One of the main characteristic of the stochastic problem model is how the information about the stochastic states of the network is obtained. The introduction of ad-hoc communication presents an opportunity for a new kind of network model – the *Gossip Networks*. In this paper we formulate, for the first time, the gossip networks model in which mobile agents obtain information about the state of a stochastic network by exchanging information with neighboring agents using peer to peer (P2P), ad-hoc communication. Mobile agents use the exchanged information to revel information about the network state and consequently optimize their routing.

There are variety of real life problems that can benefit from an optimal solution to the problem of routing in gossip networks. For example, airplanes or vessels finding the "best" route by exchanging information with their peers. This paper will focus on another example from the field of transportation. Road congestion is a known and acute urban menace with no signs of disappearing.

There are apparently many suggested approaches to tackle this problem, one of them is to supply vehicles and drivers with up-to-date information about road conditions.

There are two kinds of approaches to supply drivers with information that can aid them avoid congestion. One approach is based on fixed-structure communication networks, for example cellular networks or FM/AM radio [APG, TMC, TRM], the other approach is based on ad-hoc communication networks. Several innovative projects propose using ad-hoc networks as the communication infrastructure, for example FleetNet [ER01], and CarNet [MJK+00].

The advance in technology in recent years helps to bring into vehicles sophisticated onboard navigation systems at a reasonable price. Such a system contains a computing device with a detailed road map, GPS for locating the vehicle on the map, and communication means. One can use ad-hoc communication networks (such as Wi-Fi) to exchange information between neighboring vehicles. When two vehicles are at communication range they can exchange their information regarding road condition. The road condition information is thus propagated in the network without any need for external or central infrastructure. Each time new information is obtained by a vehicle, the onboard navigation systems recalculates the optimal route from its current location to the destination. For example, if the navigation system receives information that one of the street in its planned path is blocked it will plan a new path that avoids the blocked roads, the new path will be the shortest path from the vehicles current position to the destination taking into account the blockage.

Our gossip network model was built based on research done in "ad-hoc networks" and "stochastic shortest path routing". In this paper, mobile agents acquire and disseminate information about road conditions using wireless communication (ad-hoc networks) and use the information to minimize their traveling time (shortest path problem). There are two networks in our model, the "road network" on which the mobile agents roam and the "communication network" on which information flow. While there is an extensive literature about routing in each of the networks, to the best of our knowledge, this is the first attempt to formulate and solve the combine problem: shortest path routing of mobile agents in the context of gossip ad-hoc networks. [1]

There are currently several ongoing projects focusing on the idea of mobile agents (for example vehicles) exchanging information and forming communication networks without or with a little help from external infrastructure. Mobile Ad-hoc Networks (MANET) [MAN] is an IETF working group set to standardize these efforts. The FleetNet project [ER01] aims at the development and demonstration of a wireless ad-hoc network for inter-vehicle communications. FleetNet is a consortium of six companies and three universities looking into mostly the practical issues of providing drivers and passengers services over ad-hoc communication. Some of the proposed FleetNet services are: notifications about traffic jams and accident, and providing information about nearby available point of interest. Another project, CarNet [MJK+00] demonstrates the use of ad-hoc scalable routing protocol (Grid) to support IP connectivity as well as providing services similar to FleetNet. For a comprehensive overview of Inter-Vehicle ad-hoc communication see [Bri01].

FleetNet, CarNet, and similar projects aim at building communication infrastructure using ad-hoc communication and are researching for suitable routing protocols, medium access methods, radio modulation etc. In this paper we assume the existence of such ad-hoc network that enables mobile agents to exchange information. However, we don't implicity include here specification of the ad-hoc network such as routing or multi-access communication protocols, instead we abstract them into the *gossip probability*, the probability that a mobile agent will receive information about the status of some roads in the network from another mobile agents. The gossip probability is defined formally in Section 2.

[1] This paper focuses on the routing of mobile agents on the "roads networks" and not on the routing of data packets on the "information networks".

Classification	Description	Gossip Networks
Time	In time dependance problems the network's weights changes with respect to time	Time independent
Solution type	In deterministic problems the shortest path is determined a-priory while in dynamic problems the traveler is given a policy and adapt his path while traveling according to the obtained information (also known as the *recourse* problem)	Policy based
Information Source	The mechanism in which travelers obtain information about the weights of the network. Most models assume knowledge of emerged roads	Gossiping and emerged roads

Table 1: Shortest Path Problems Classification

The problem of *Shortest Path Routing* was investigated extensively in the literature, for a comprehensive summary of the various efforts in the field of transportation see [PS98]. Shortest path problems can be classified according to several criteria as outlined in Table 1.

In this paper we assume time independence, i.e., the network doesn't change during the course of the travel. Some of the road conditions are known to be alternating, however, a traveler may not know in advance the current condition of all these roads, termed stochastic roads. We assume that no waiting at roads or junctions is allowed and once a junction is reached the weights of all the roads that emerge from that junction become known. We investigate two different models of weight correlation. The first is the *Independent Weight Correlation* model (G-IWC) where there is no correlation between the states of different edges. The second is the *Dependent Weight Correlation* model (G-DWC) where the network can be in several different states, each state determines the weights of all stochastic edges [PT96]. Note that the G-IWC model is a generalization of the G-DWC model with substantially more states. The rational behind the G-DWC model is that in "real-life" transportation systems there is a correlation between roads weights, usually a traffic jam on one road effects the roads in its vicinity.

When the shortest path model is stochastic, like in this paper, the information about the actual state of the stochastic edges plays a crucial role in finding the optimal routing solution. Further more, due to the dynamic nature of the problem the solution is not a path but rather a policy that direct the traveler according to the information he obtains. In the literature there are several papers that discuss optimal routing policies in stochastic networks where the traveler can recourse his path according to information obtained during travel. However, the basic difference between these models and ours is that in gossip networks the information is obtained by gossiping with neighboring travelers thus a traveler can obtain data about the state of remote stochastic roads. In all the other models we survey the only way to obtain information about the state of a road is to visit the junction it emanates from. Andreatta and Romeo [AR88] assume that once a blockage is encounter a recourse path that consists of only deterministic roads is used. Orda, Rom, and Sidi [ORS93] investigated a model where link delay change according to Markov chains, they model several problems and showed that in general, the problems are intractable. Polychronopoulos and

Tsitsiklis [PT96] investigated a network where there is a correlation between the roads weights. In their model a traveler can deduce the stochastic state by visiting enough roads. Waller and Ziliaskopoulos [WZ02] solved a model with dependency between successor roads and a model with time dependency for the same road.

The introduction of information exchange in gossip networks leads to unique optimal routing policies. In this paper we will show that sometimes it is worth taking a detour to obtain more information about the state of the stochastic edges. The extra cost of the short detour can be compensated by the additional information gained, information that can improve the selection of the continuing path.

The rest of the paper in organized as follows. In the next section, the formal model of the gossip networks is introduced and an example that demonstrates the characteristics of the model is presented. An algorithm for optimal routing in gossip networks that is based on dynamic programming is developed in Section 3. In Section 4 we discuss the implications of traveling in gossip networks. Then, in Section 5, we use numerical analysis to demonstrate the influence of the various model parameters on the network behaviors. Finally in Section 6 we summarize and highlight our main findings.

2 Model and Definitions

2.1 The formal model

The above discussion leads to the following formal model. The network[2] is represented by a directed graph $G = (V, E)$, where V is the set of vertices, and E is the set of edges, $|V| = n$ and $|E| = m$. An edge $e \in E$ is associated with a discrete random weight variable, w_e. Edges with degenerated weight function that has only one value are termed deterministic, and we denote the set of these edges by $D \subseteq E$. The number of edges in the network with stochastic weights (namely, non deterministic) is denoted by $\delta = |E \setminus D|$. We assume that under all weight distributions there are no negative cost cycles in the network and there is always a path between source and destination.

In the G-IWC model the weights, w_e, of the *stochastic edges* are random variables with discrete probability distribution that has β_e states. The expected cost of an edge is $\sum_{s=1}^{\beta_e} w_e^s q_e^s$, where q_e^s is the probability of an edge e to have the weight w_e^s. We denote by $\hat{w}_e$ the actual weight of the edge e. In the G-DWC model the network can be in only R realizations, each $r \in R$ realization determines the states of the network and thus the weights w_e^r of all the stochastic edges.

Traveling agents (TAs) are roaming the network. Each TA stores internally the weights of the stochastic edges in an *Information Vector* (IV), $I\{\cdot\}$. For example, an information vector of a traveler could look like this: $I = \{\hat{w}_1, X, \hat{w}_3, X, \ldots, X, X, \hat{w}_\delta\}$. For known edges, those that the traveler visited or received information about, the weights are written down explicitly, $\hat{w}_1, \hat{w}_3, \hat{w}_\delta$. Unknown edge weights are denoted by X. The number of possible states of the information vector in the G-IWC model, l_I is given by

$$l_I = \prod_{e \in E \setminus D} (\beta_e + 1) \tag{1}$$

and in the G-DWC model, the number of different information vector states is given by

[2] As mentioned above, there are two networks in our model, the "road network" and the "communication network". From this point on, when we say "network" we refer to the "road network". We assume the existence of communication network that enable mobile agent to exchange information but in this paper we don't include it in the formal model implicity, it is included in the gossip probability presented below.

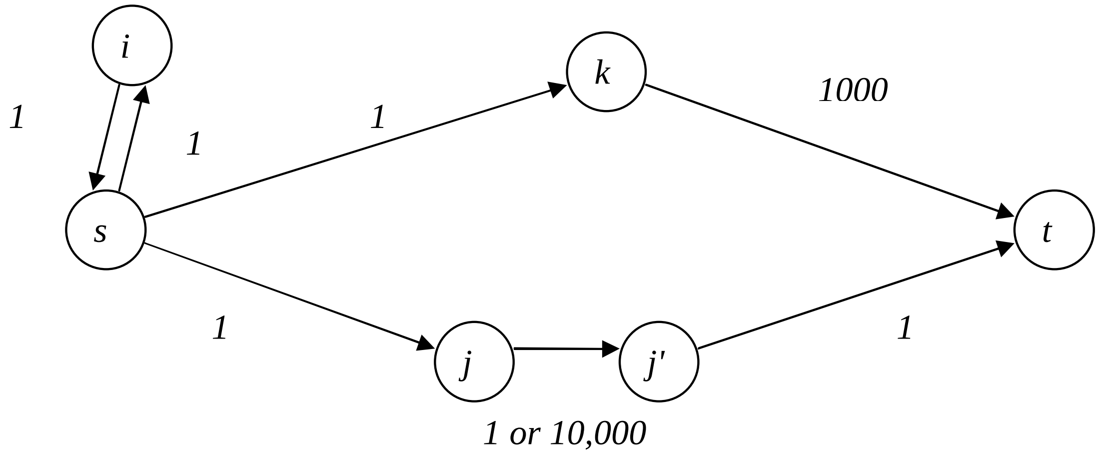

Figure 1: An example of the influence of gossiping on routing. We are looking for the optimal routing policy between the vertices s and t where the edge (j, j') is stochastic and on edge (i, s) the traveler can obtain information about the stochastic edge.

$$l_D = \sum_{i=1}^{R} \binom{R}{i} = 2^R - 1 \tag{2}$$

When two or more TAs are within communication range they can exchange their IV in order to gain missing data. The *gossip probability* is a function holding the probability that when a TA traverses an edge it will obtain new information that will update his IV. We denote the gossip probability by $P(I(i), I(j), T(i, j))$ where (i, j) is the edge traversed, $I(i)$ and $I(j)$ are the IV before and after (i, j) traversal, respectively, and $T(i, j)$ is the *topology probability.* The topology probability is determined by aspects like the number of TAs around the traveler, the other TAs previous paths, physical obstacles that interfere with the wireless communication, etc. It is a characteristic of the network structure and the flows of TAs in the network. $T(i, j)$ is a vector of probabilities, where each element corresponds to some stochastic network edge. For example, $T(i, j) = \{1, 0.5, \ldots 0\}$ means that when the TA slates edge (i, j) it will learn about stochastic edges 1,2, and δ with probability 1, 0.5, and zero, respectively. The gossip probability also depends on the IV before and after the edge traversal, for example, the probability to change an IV element from $\{\hat{w}\}$ to $\{X\}$ is zero – a known weight can not be changed into unknown.

In this paper we are looking for the optimal routing policy of a TA that start at the source vertex s with information vector $I(s)$ and traveler to a destination vertex t. We assume that the TA knows a-priory the network structure, weights distribution, and the topology probability. We are looking for an optimal routing policy, π^* with minimal expected cost, $C^*(s, t, I(s))$, of all possible routing policies $\pi^k \in \pi$.

$$\forall\ \pi^k \in\ \pi \quad C^*(s, t, I(s)) \leq C^k(s, t, I(s))$$

2.2 An Example

In the example network presented in Fig. 1, a traveler is located at vertex s and is looking for the optimal routing policy to vertex t. In this network there is one ($\delta = 1$) stochastic edge, (j, j'), that has two possible states. With probability $q^u_{jj'} = \xi_U$ the edge is in the "UP" state where $w^u_{jj'} = 1$, and with probability $q^d_{jj'} = (1-\xi_U)$ the edge is in the "DOWN" state where $w^d_{jj'} = 10000$. The traveler can obtain information about the state of the edge (j, j') only when traversing the edge (i, s), with the topology probability: $T(i, s) = \xi_T$. The gossip probability of this network is:

$$\begin{aligned} P(\{X\}, \{X\}, T(i,s)) &= 1-\xi_T \\ P(\{X\}, \{1\}, T(i,s)) &= \xi_T \\ P(\{X\}, \{10000\}, T(i,s)) &= \xi_T \\ P(\{1\}, \{1\}, T(i,s)) &= 1 \\ P(\{10000\}, \{10000\}, T(i,s)) &= 1 \\ \text{Else } \forall\ u, v \in V \quad P(I(u), I(v), T(u,v)) &= 0 \end{aligned}$$

The traveler has to chose between different travel options: *a*) The "safe" path through vertex k which guarantee a cost of 1001 or; *b*) The "risky"[3] path through vertex j with cost that depends on the state of edge (j, j'), either 10002 or 3 or; *c*) Travel to vertex i, obtain information about the status of edge (j, j') and then, according to the obtained information, chose whether to go through vertex k, j or return to vertex i.

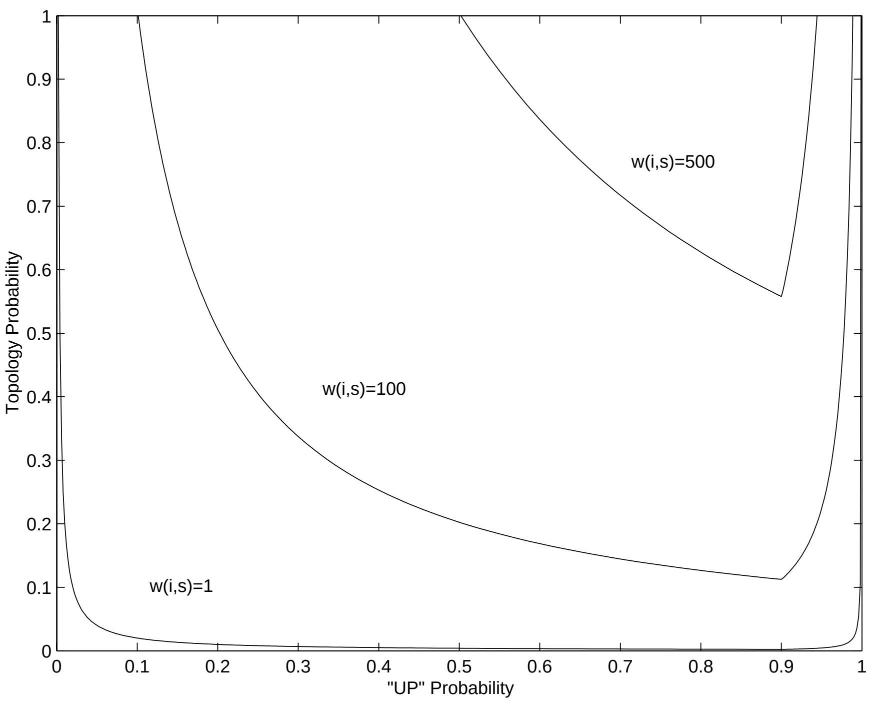

Figure 2: The relation between the "UP" and gossip probabilities for different w_{is} values. The area above the line is where $C^*(s,t,\{X\})_i < C^*(s,t,\{X\})_{kj}$ and the traveler will cycle for information

Next we will calculate the expected cost of the different routing policies. The cost of the path through vertex k is deterministic and does not depend on the a-priory knowledge of the state of the edge (j, j')

[3]The risky policy is taken by a traveler that must reach the destination at some specific time (for example to catch a plane that leaves in 10 time units). If not there by that time the traveler care less about the path cost (anyway he needs to reschedule).

$$C(s,t,\{\cdot\})_k = 1001 \tag{3}$$

The cost of the path through vertex j without any a-priory knowledge about the state of the edge (j, j')

$$C(s,t,\{X\})_j = 10002(1-\xi_U) + 3\xi_U \tag{4}$$

If the traveler needs to choose between traveling through k or j (without first traveling to vertex i) then his optimal routing policy depends on the value of his information vector:

$$\begin{aligned} C^*(s,t,\{X\})_{kj} &= \min(1001, (1-\xi_U)10002 + 3\xi_U) \\ C^*(s,t,\{1\})_{kj} &= 3 \\ C^*(s,t,\{10000\})_{kj} &= 1001 \end{aligned}$$

If the traveler knows that the stochastic edge is in the "DOWN" state he will travel to vertex k; in the case he knows that the edge is in the "UP" state he will travel to vertex j; and in the case the traveler doesn't know the state of the stochastic edge he will decide according to the value of ξ_U.

When the traveler moves to vertex i without any a-priory knowledge about the state of the edge (j, j') the expected cost of his routing policy assuming one trial to obtain information is:

$$\begin{aligned} C(s,t,\{X\})_i^{(1)} &= 2 + \xi_T[\xi_U C^*(s,t,\{1\})_{kj} + \\ &(1-\xi_U)C^*(s,t,\{10000\})_{kj}] + (1-\xi_T)C^*(s,t,\{X\})_{kj} \\ &= 2 + \xi_T[3\xi_U + 1001(1-\xi_U)] + (1-\xi_T)C^*(s,t,\{X\})_{kj} \end{aligned} \tag{5}$$

When the traveler routing policy is to cycle between vertices s and i until it obtains information, the expected number of cycles he will need is $1/\xi_T$. Therefore

$$C(s,t,\{X\})_i = 2(1/\xi_T) + 3\xi_U + 1001(1-\xi_U)$$

For the above example there is a threshold probability, ξ_0, such that for $\xi_T \geq \xi_0$

$$C^*(s,t,\{X\})_i < C^*(s,t,\{X\})_{kj} \tag{6}$$

Meaning that for $\xi_T \geq \xi_0$ the traveler's optimal routing policy when he has no information about the state of the stochastic edge is the one that makes a detour through node i until it obtains information about the state of the stochastic edge. Fig. 2 illustrates this by plotting the equilibrium line of Eq. (6) for different values of $\hat{w_{is}}$. The area above the line is where the inequality holds and the traveler is making a detour to gather information. The minimum of the plots in Fig. 2 is when Eq. 4 and Eq. 3 are equal, at $\xi_U = 0.90028$ in this example.

The optimal routing policy for a traveler that starts on vertex s is outlined in the EXAMPLE_POLICY below. And the corresponding routing table for source vertex s is outlined in Table 2.

Topological Probability(ξ_T)	$I(s)$	next hop
$\geq \xi_0$	$\{X\}$	i
$\geq \xi_0$	$\{1\}$	j
$\geq \xi_0$	$\{10000\}$	k
$< \xi_0$	$\{X\}$	α
$< \xi_0$	$\{1\}$	j
$< \xi_0$	$\{10000\}$	k

Table 2: Routing table of the source vertex s. The value of α is k or j according to the value of ξ_U.

EXAMPLE_POLICY

```
IF ξ_T ≥ ξ_0
    WHILE I = {X} cycle in the path {s, i, s}
IF I = {1}
  Then take the path L_U = {s, j, j', t}
ELSE IF I = {10000}
  Then take the path L_D = {s, k, t}
ELSE IF I = {X}
  Then take the path min(L_U, L_D)
END
```

3 The Routing Algorithm

3.1 Solution approach

The optimal routing policy in the gossip networks is the one with the minimum expected cost from source to destination for a given information vector. In this subsection we will show how one can calculate the expected cost of a routing policy in the network, in the next subsection we will introduce an algorithm that uses these calculations to find the optimal routing policy to destination.

A traveler starts his journey on vertex s with information vector $I(s)$ and wants to reach vertex t. During his journey, there is a probability that he will learn, through gossiping, about the states of the stochastic edges and accordingly update his information vector $I(\cdot)$. At every vertex $r \in V$ he reaches, the traveler makes a routing decision, based on his updated information vector. The expected cost of a routing policy between a source vertex, s, and a destination vertex, t through a neighbor vertex r is:

$$C(s,t,I(s))_r = \hat{w_{sr}} + \sum_{I(r)\in B(I(s),(s,r))} P(I(s),I(r),T(s,r)) \cdot Q(I(r)) \cdot C(r,t,I(r)) \tag{7}$$

The weight of edge (s,r) is known and its value is $\hat{w_{sr}}$. $B(I(s),(s,r))$ is the set of all the possible information vectors $I(r)$ of the traveler when reaching vertex r, assuming that at vertex

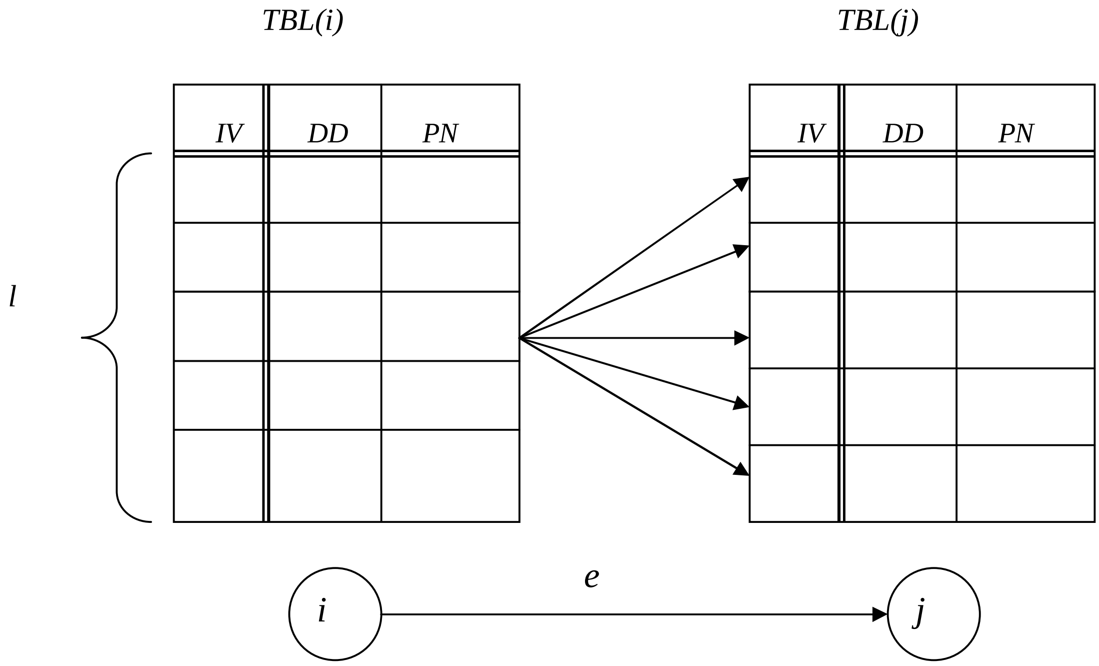

Figure 3: The relaxation process for one state of one edge.

s it has the information vector $I(s)$. $P(I(s), I(r), T(r,s))$ is the probability that the information vector will change from $I(s)$ into $I(r)$ on the edge (s,r). And $Q(I(r))$ is the a-priory probability that the network G is in a state corresponding to the information in $I(r)$.

3.2 Dynamic Programming Algorithm

The optimal routing policy from vertex s to vertex t in the gossip networks, $C^*(s,t,I(s))$, is the one that minimizes the expression in Eq. 7. Namely, the one that selects the policy with the smallest expected cost. Thus, we can write the following dynamic program:

$$C^*(s,t,I(s)) = \min_{r \in \mathcal{N}_s} \{ \hat{w_{sr}} + \sum\nolimits_{I(r) \in B(I(s),(s,r))} P(I(s), I(r), T(r,s)) \cdot Q(I(r)) \cdot C^*(r,t,I(r)) \} \tag{8}$$

where $\mathcal{N}_i$ is the group of neighbors of vertex i.

The proposed dynamic programming algorithm is similar to the Bellman-Ford algorithm [BG92]. In the classical Bellman-Ford equations one finds for each vertex the shortest path to a destination. In our gossip networks, we need to find for each vertex the shortest path for each possible state of the vertex's information vector $I(\cdot)$.

Specifically, for each vertex $u \in V$ we keep a table, $TBL(u)$, (see Fig. 3) that has l rows (l is defined in Eq. 1 or Eq. 2 according of the model in use. Each row holds the information vector state, $s_k \in I$, the distance to destination, DD, and a pointer to next vertex, PN.

The relaxation processes for each edge (u,v) and for each information vector state s_k is:

$$DD(u,s_k) = \hat{w_{uv}} + \sum_{m=1}^{l} P(s_k, s_m, T(u,v)) Q(s_m) DD(v, s_m)$$

For each source vertex state, s_k, the algorithm checks what is the probability that during the travel on the edge (u,v) the state s_k will change into s_m, $(m = 1 \ldots l)$, each of the possible information vector states of the destination vertex v. Each gossip probability $P(s_k, s_m, T(u,v))$ is multiplied by the destination vertex distance $DD(v, s_m)$ and the probability $Q(s_m)$ that the network will be in state s_m.

The complete algorithm GOSSIP_DP is presented below.

GOSSIP_DP(G,w,T,s,t)

```
For k = 1 To l
    DD(t, s_k) ← 0; PN(t, s_k) ← t
Cont ← TRUE
For Each u ∈ V \ t
    For k = 1 To l
        DD(u, s_k) ← ∞; PN(u, s_k) ← NIL
For count = 1 to |V| − 1 && While Cont = TRUE
    Cont ← FALSE
    For Each e ∈ E
        if GOSSIP_RELAX(e) then Cont ← TRUE
End

Function GOSSIP_RELAX(e)
    u ← Source(e); v ← Destination(e)
    Relax ← FALSE
    For k = 1 To l
        tempDD ← ŵ_e +
        Σ_{m=1}^{l} TRANS_PROBE(s_k, s_m) · DD(v, s_m)
        If tempDD < DD(u, s_k) Then
            DD(u, s_k) ← tempDD
            PN(u, s_k) ← v
            Relax ← TRUE
    Next k
    Return(Relax)
End Function

Function TRANS_PROBE(s_k, s_m)
    P ← probability to move from s_k to s_m on e
    Q ← probability of the network to be in s_m
    Return(P · Q)
End Function
```

Before the traveler starts his journey he builds his optimal routing policy by calculating TBL for all the vertices of the network using the algorithm GOSSIP_DP. During his journey the traveler updates his information vector and navigates on the network using the information in TBL. Every time the traveler reaches a new vertex $u \in V$ with information vector state $s_k = I(u)$ he looks for the next vertex in $PN(u, s_k)$.

3.3 Complexity of G-IWC and G-DWC

Theorem 3.1 *The complexity of the GOSSIP_DP algorithm under the G-IWC model is* $O(nm\delta(2\beta+1)^{\delta})$.

Proof: When there is no correlation between the edges weights we must examine all the edges ($O(|E|)$); for each edge we must examine all the source vertex stochastic states ($O(l_I)$); and for each source vertex stochastic state we examine all the destination vertex's stochastic state ($O(l_I)$), here we assume that the number of stochastic states is bounded by β. Notice however that not all state transfers are possible and actually the number of possible state transfer we need to examine is only $(2\beta+1)^{\delta}$. The first $\beta+1$ states are for the transfer from state {X} to all the available states, the second β states are for staying in the same state when the weight of the stochastic edge is known. In each state transfer we need to calculate the transfer probability P and the a-priory probability Q, for that we need to examine all stochastic edges $O(\delta)$. In the worst case a vertex has $O(|V|)$ neighbors and the algorithm terminates either after repeating for each of the neighbors or when there is no difference between successive iterations. ■

Theorem 3.2 *The complexity of the GOSSIP_DP algorithm under the G-DWC model is* $O(nm\delta 2^{2R})$.

Proof: The complexity of GOSSIP_DP algorithm under the G-DWC model is similar to the complexity of the algorithm under the G-IWC model. The only different is that we need to examine $O(l_D)$ transfer states instead of $O(l_I)$ states. According to Eq. 2, $O(l_I) = 2^R$.
■

Although the optimal solution to the gossip networks problem is intractable in general, we presented above two special cases where the optimal solution is polynomial in respect to the network size. In the first case a polynomial solution is obtained when the number of stochastic edges δ is small. The second case is when the number of realizations in the network is relatively small.

4 Discussion

4.1 Gossiping and Learning

In this subsection we will illustrate the importance of gossiping by comparing the learning rates of the gossip and non-gossip travelers. We assume the G-DWC model with R possible realizations. When the traveler starts his journey he doesn't know what is the current network realization $r \in R$. Each time he gathers information about some edge weights he could eliminate zero or more network realizations in which those weights are inconsistent. Depending on the network weights distribution, the traveller will be able to determine the current realization of the network after obtaining information about the state of enough edges. Since each time the traveler visits a vertex he gathers information about the state of all the emerging roads we define information vertices as the set of vertices the traveler needs to visit in order to find the current network realization, and senote it by k. In the following subsection we assume that the traveler doesn't visit a vertex more than once and that the information vertices are distributed uniformly at random in the network.

We first analyze the non-gossip traveler which we call a *Step-By-Step* (SBS) traveler, he receives information about a vertex only when he visits it. The probability that after i steps in the network (visiting i vertices) the SBS traveler already visited j out of the k information vertices is given by the hypergeometric distribution.

$$Pr(n,k,i;x=j) = \frac{\binom{k}{j}\binom{n-k}{i-j}}{\binom{n}{i}} \quad where \quad j \leq k; \; j \leq i \leq n$$

The probability that after visiting i vertices the SBS traveler already visited all k information vertices and thus found the current network realization is

$$Pr(n,k,i;x=k) = \frac{\binom{k}{k}\binom{n-k}{i-k}}{\binom{n}{i}} \quad where \quad k \leq i \leq n$$

The expected number of steps the SBS traveler needs to take to find all k information vertices is

$$\sum_{i=k}^{n} iPr(n,k,i;x=k) = \sum_{i=k}^{n} i\frac{\binom{n-k}{i-k}}{\binom{n}{i}}$$

After normalizing the above expression we get

$$\frac{\sum_{i=k}^{n} i\frac{\binom{n-k}{i-k}}{\binom{n}{i}}}{\sum_{i=k}^{n} \frac{\binom{n-k}{i-k}}{\binom{n}{i}}} = \frac{\sum_{i=k}^{n} i\frac{i!}{(i-k)!}}{\sum_{i=k}^{n} \frac{i!}{(i-k)!}} = \frac{(n+1)k+n}{2+k} \tag{9}$$

Eq. 9 indicates that the number of steps the SBS traveler needs to take in order to find the current network realization is proportional to the network size, n.

Unlike the SBS traveler that can only gather information about one new vertex in each step, the gossip traveler has an additional probability to receive information about all the network's remaining unknown vertices. In his first step the gossip traveler receives information about $\xi_T n$ vertices and in the i step about $\xi_T(1-\xi_T)^i n$ vertices. In each step the gossip traveler has information about all the vertices he learned about in his previous steps. Therefore, in the i step the gossip traveller has information about $g(i)$ vertices:

$$g(i) = \sum_{j=0}^{i-1} \xi_T(1-\xi_T)^j n = (1-\bar{\xi_T}^i)n$$

where $\bar{\xi_T} = 1 - \xi_T$.

Obviously, when the traveler gathers information about all n network's vertices he has information about all k information vertices and knows the network current realization. Thus, an upper bound on the expected number of steps the gossip traveler needs to take is the number of steps needed to gather information about all the network vertices. Since the number of vertices is discrete we are looking for the step number, r, such as

$$g(r+1) - g(r) = n(\bar{\xi_T}^r - \bar{\xi_T}^{r+1}) < 1$$

Solving the above equation yield

$$r < -\frac{\ln(n\xi_T)}{\ln(1-\xi_T)} \tag{10}$$

In the gossip model presented in this paper r could be even smaller since the gossip traveler gather information by both gossiping and visiting vertices. In the above analysis we took into account only gossiping. Thus, Eq. 10, is an upper bound on the expected number of steps the gossip traveler needs to take in order to find the current network realization. Comparing Eq. 10, to the expected number of steps the SBS traveler needs to take, Eq. 9, we conclude that the outcome of gossiping is higher learning rate. While the SBS traveler needs on average to visit $\mathrm{O}(n)$ vertices of the network to learn its state, the gossip traveler needs to visit only $O(\log(n))$ of them. In most cases higher learning rate in stochastic networks will result in shortest path to destination. Once the traveler knows the network edge's states he can reduce his path cost, for example by avoiding blocks roads.

4.2 Characteristics of traveling in Gossip Networks

In this section we will discuss the characteristics of optimal routing in gossip networks under the proposed GOSSIP_DP algorithm. The discussion is under the following assumptions: The network is in the G-IWC model with one stochastic edge, and the traveler needs to traverse over the stochastic edge on his way from source to destination. The stochastic edge can be either in the "UP" or "DOWN" states. In the "UP" state the stochastic edge weight is similar to the weight of the deterministic edges, in the "DOWN" state its weight is higher than the weights of the deterministic edges. Once we analyze the parameters that influence routing under those assumptions expanding the model to the case of several stochastic edges with several stochastic states is straightforward.

A traveler in the gossip networks that is navigating using our optimal routing policy can be viewed as operating in three different regimes: "WIN", "LOSE", and "NEUTRAL". In the *"WIN"* regime the traveler reduces his travel cost by gossiping. In the *"NEUTRAL"* regime obtaining information doesn't increase or decrease the gossip traveler's path cost. In the *"LOSE"* regime obtaining information actually increases the traveler path cost. The operating regime is a result of the following parameters: the magnitude of the difference between the values of the actual weight of the stochastic edges ($\hat{w}_e$) and their expected weights (ω_{SE}), the values of the topology probability (ξ_T), and the magnitude of the difference between the values of the stochastic edges actual state (ξ_A) and a-priory probability to be in the "UP" state (ξ_U) (see Table 4 for notation summary). Next we will explain the influence of each parameter.

The magnitude of the difference between the traveler's a-priory knowledge (ω_{SE}) and the actual weight of the stochastic edges ($\hat{w}_e$), denoted by $\Delta_\omega = |\omega_{SE} - \hat{w}_e|$, determines the influence of obtaining information on the traveler's path cost. When ω_{SE} and $\hat{w}_e$ are similar, a gossip traveler will not have an advantage over a non-gossip traveler, they both know a-priory the "correct" stochastic state. However, above some critical difference, $\Delta_\omega > \Delta_C$ obtaining information will decrease the traveler's path cost. For example, when ω_{SE} "tells" the travelers that a stochastic edges is in the "UP" state and the actual state is "DOWN" a non-gossip traveler may include this edge in its path while a gossip traveler will reduce his path cost by bypassing it in advance. Similarly, when ω_{SE} "tells" the traveler that a stochastic edge is in the "DOWN" state and the actual state is "UP" a non-gossip traveler will not include the stochastic edge in its planned route trying to bypassing it and therefore will increase his path cost compared to the gossip traveler. The value of Δ_C is determined by the difference that will cause the non-gossip traveler to take the wrong path, meaning that he will bypass the stochastic edge when it's "UP" or traveler through it when it's "DOWN".

Fig. 4 illustrates the different possible types of paths a traveler can have for different values of topology probability (ξ_T). When there is no gossiping (a) the probability to receive information is zero thus the optimal policy is determined a-priory before the start of the journey and has no

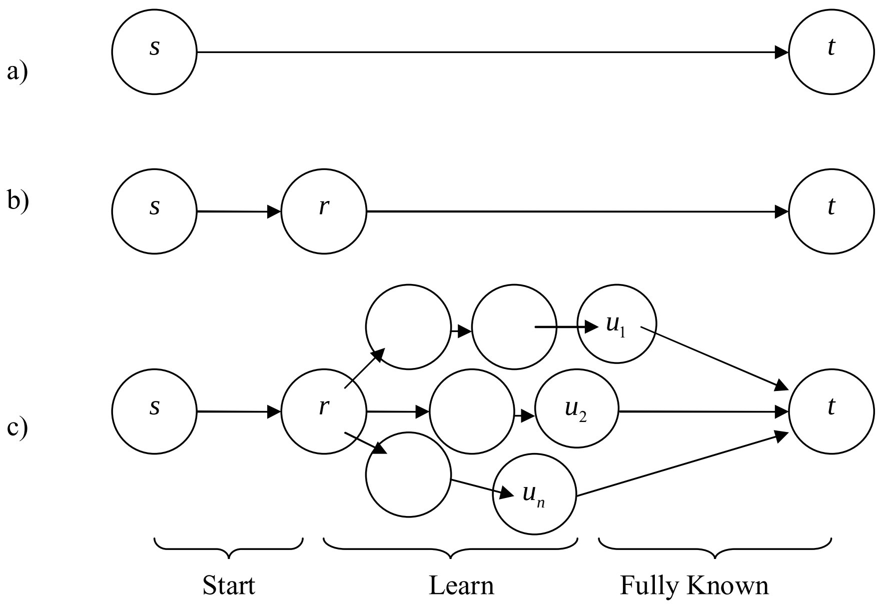

Figure 4: The different possible paths a traveler can have for different topology probabilities. (a) No gossiping, (b) Maximal gossiping, and (c) In between.

recourse. In this case the optimal policy is the one that minimize the expected weights. When ξ_T is maximal (b) the traveler learns about the state of all the stochastic edges on the traversal of the first edge, (s, r), and then travels to the destination t with full knowledge about $\hat{w}_e$ and therefore without changing his course. When ξ_T is in between (c) the traveler's path is composed of three phases, the initial phase is until the traveler obtains any information about the state of the stochastic edges. Then, in the learning phase, the traveler may recalculate and recourse his path according to the updated information vector – his optimal policy is a collection of different branches. When the traveler has full information about $\hat{w}_e$, at some vertex u, he travels to the destination without changing his course. The higher ξ_T the quicker the gossip traveler will learn about the state of the network and therefore minimize the learning phase in his travel which leads to decrease in the policy cost.

According to the optimal policy, stated in Eq. 8, one of the parameters that determines the relative weight of each branch in the path is the a-priory probability of the network to be in certain stochastic state, denoted here by ξ_U. The closer ξ_U is to ξ_A (small $\Delta_\xi = |\xi_U - \xi_A|$) the more efficient the learning phase will be. Efficient learning means that the traveler is directed toward the "right" direction by giving higher relative weight to the "right" branch. When there is a relatively large difference between ξ_U and ξ_A the branches in the learning phase will direct the traveler to the "wrong" direction and as a result the cost of his policy will increase. For example, when the a-priory probability of the stochastic edge to be in the "UP" state is small ($\xi_U \approx 0$) the optimal policy will direct the gossip traveler to branches that detour the stochastic edge. When the stochastic edge is actually in the "DOWN" state this direction is justify, however when the actual state of the stochastic edge is "UP" the direction will maximize the gossip traveler learning phase and his total traveling costs.

The parameters that influence the traveler optimal policy cost in the gossip networks are summarized in Table 3.

The operating regime that the traveler experiences is determined by the combined values of the parameters, Δ_ω,ξ_T, and Δ_ξ. Fig. 5 illustrates the influence of the parameters on the network

Parameter	Description
$\Delta_\omega = \|\hat{w}_e - \omega_{SE}\|$	Increase in Δ_ω directs the traveler to the "wrong" direction
ξ_T	Increase in the ξ_T decreases the learning phase
$\Delta_\xi = \|\xi_A - \xi_U\|$	Increase in Δ_ξ increases the learning phase

Table 3: The parameters that determine the operating regimes of a traveler in the gossip network.

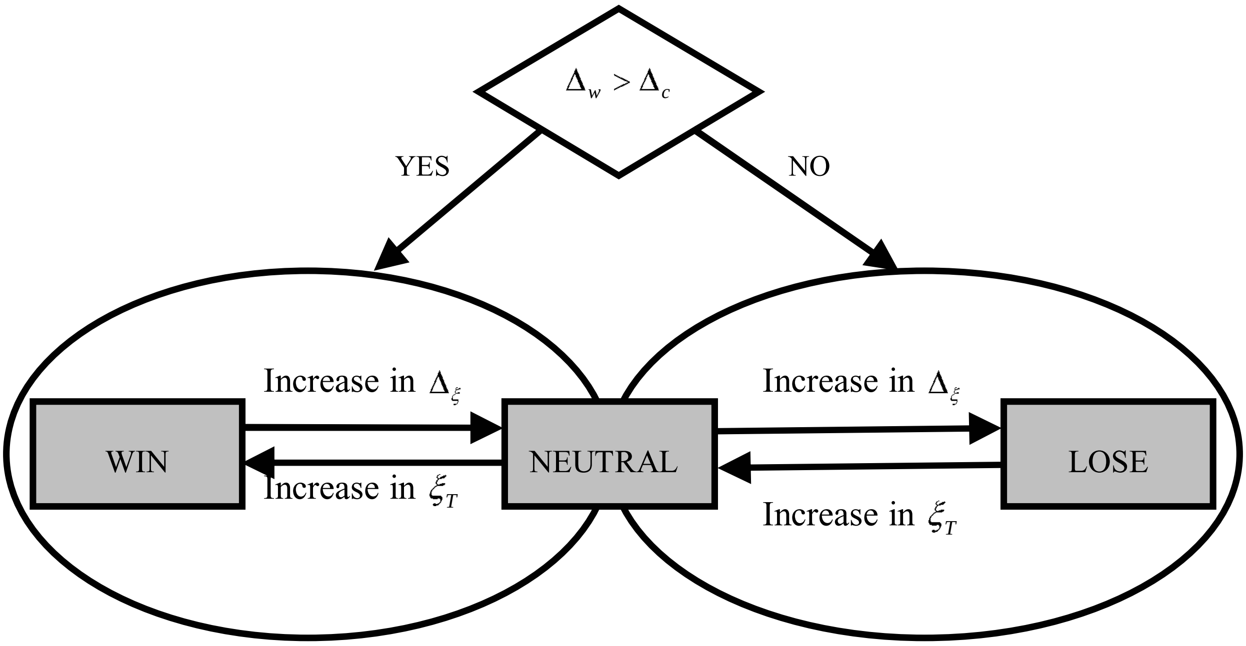

Figure 5: The influence of the gossip network parameters on the network regime

regime. When Δ_ω is below some threshold, Δ_C, the a-priory knowledge of the network state is close enough to the true value, and thus increasing the path length to obtain information can not benefit the gossip traveler. As a results, in this case, the network can be either in the "NEUTRAL" or "LOSE" regimes. The "LOSE" regime is obtained when the learning phase is relatively large (increase in Δ_ξ), however a larger topology probability shortens the learning phase and pushes the network into the "NEUTRAL" regime. The ultimate network regime is determined by the relation between those two parameters ξ_T, and Δ_ξ. Similarly, when Δ_ω is above the threshold, Δ_C, gossiping helps the gossip traveler to reduce his policy costs. The network can be either in the "WIN" or "NEUTRAL" regimes according to the relation between ξ_T, and Δ_ξ. In the next section, we will demonstrate the above discussion using the simulation results.

5 Numerical Analysis

The main purpose of the simulations was to investigate the influence of gossiping on the traveler's optimal policy cost under the different parameters used in the gossip networks. The performance and behavior of the proposed algorithm on the gossip networks are examined through numerical

Notation	Description
ω_D :	Weights of the deterministic edges
ω_{SE}:	Expected weights of the stochastic edges
$\hat{w_e}$:	Actual weight of the stochastic edges
ω_{SD} :	Weights of the stochastic edges in "DOWN" state
ξ_A :	Stochastic edges actual state
ξ_T :	Topology probability
ξ_U :	A-priory probability of the stochastic edges to be in the "UP" state
θ_E:	Expected cost of the optimal policy
θ_R:	Relative expected cost (θ_E) at some topology probability; $\theta_E(\xi_T)/\theta_E(0)$
θ_A:	The average of relative expected (θ_R) over the whole range of ξ_T
Configuration:	A set of values for the above parameters
Operation Regime:	Determined by the network configuration. Can be either "WIN", "NEUTRAL" or "LOSE"

Table 4: Notation Summary

experiments on various grid networks configurations with random generated weights under the G-IWC model. In each network configuration the simulation derived results comparing the traveler's expected optimal policy cost for different topology probabilities.

First, for each random generated network configuration the optimal routing policy tables are calculated using the GOSSIP_DP algorithm. Then, using the calculated routing tables the simulation computes the expected optimal policy cost from each vertex to the destination. For notation of the parameters we use see Table 4.

5.1 Simulation design

The simulation was conducted on fully connected grid networks representing, for example, road structure in many urban areas. Fig. 6 shows such a network for a 4 × 4 grid. The weights of the different deterministic edges were selected uniformly randomly. Three specific edges in the grid were chosen to be stochastic. The stochastic edges could be in two states, with probability ξ_U in the "UP" state, then the edge weights are randomly selected exactly like the deterministic edges. When the stochastic edges are in the "DOWN" state their weights are set to different values as explain further below. The stochastic edges were selected such that they will have a significant influence on the optimal policy to the destination vertex t. For the same reason, the weight of the deterministic edge that is adjacent to t was set to be higher than the other deterministic edges.

The following list details the range of values we used in the simulation:

Deterministic weight (ω_D) : Uniform random in [1,100].

Stochastic "DOWN" weight (ω_{SD}) : In each configuration all stochastic edges had the same weight which was selected uniformly at random in [0,800].

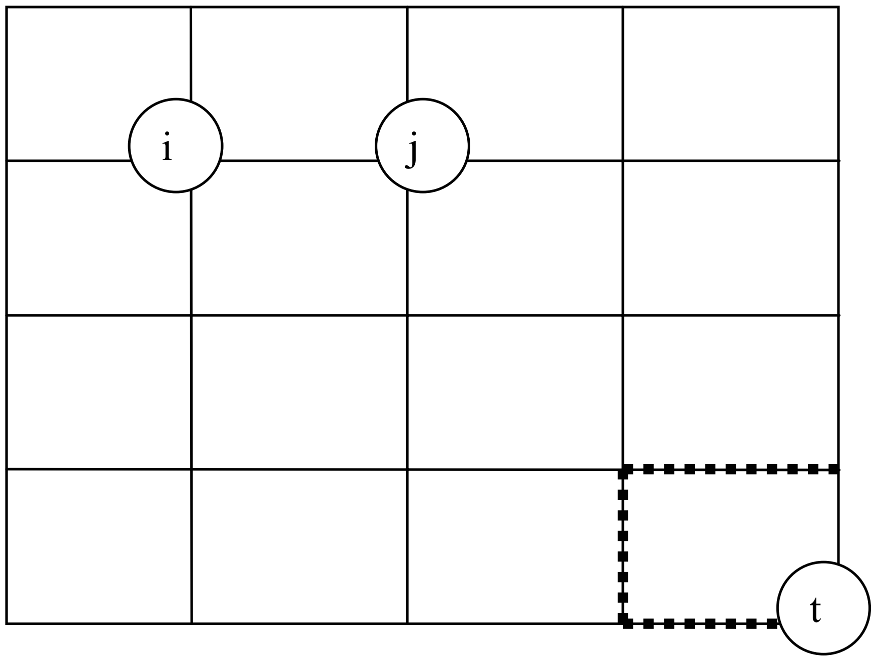

Figure 6: A 4×4 grid network used in the simulations. The dashed lines are stochastic edges with probability ξ_U to be in the "UP" state. Larger grids had the same structure.

Topology probability (ξ_T) : In each configuration the same value of ξ_T was set to all the edges in the network. The range of tested values was in [0,1].

A-priory probability (ξ_U) : Different values in the range [0,1] were used to test the influence of ξ_U. In each configuration all stochastic edges were set to the same value.

Stochastic actual state (ξ_A) : The actual state of all three stochastic edges was set equally to either "UP" or "DOWN".

Network structure (Grid Size) : Two different network grids were used with sizes of 4 × 4 and 8 × 8.

Totally we tested $21(\xi_T) \cdot 9(\omega_{SD}) \cdot 11(\xi_U) \cdot 2(\xi_A) = 4158$ different configuration for each grid size.

In order to remove the influence of specific random network weights the same set of experiments were repeated with the same network configuration for ten different random seeds. The analyzed results are averaged over the ten different runs.

5.2 Performance Measurement

After the routing tables were build for a given network configuration the *Expected Cost* (θ_E) from each vertex to the destination was calculated. θ_E is calculated by following all the possible paths from source to destination assuming that the traveler starts his travel with no information $IV = \{X, X, X\}$. The paths were weighted according to their probability to occur. The results are presented using the value of *Relative Expected Cost* (θ_R), where

$$\theta_R(\xi_T) = \frac{\theta_E(\xi_T)}{\theta_E(\xi_T = 0)}$$

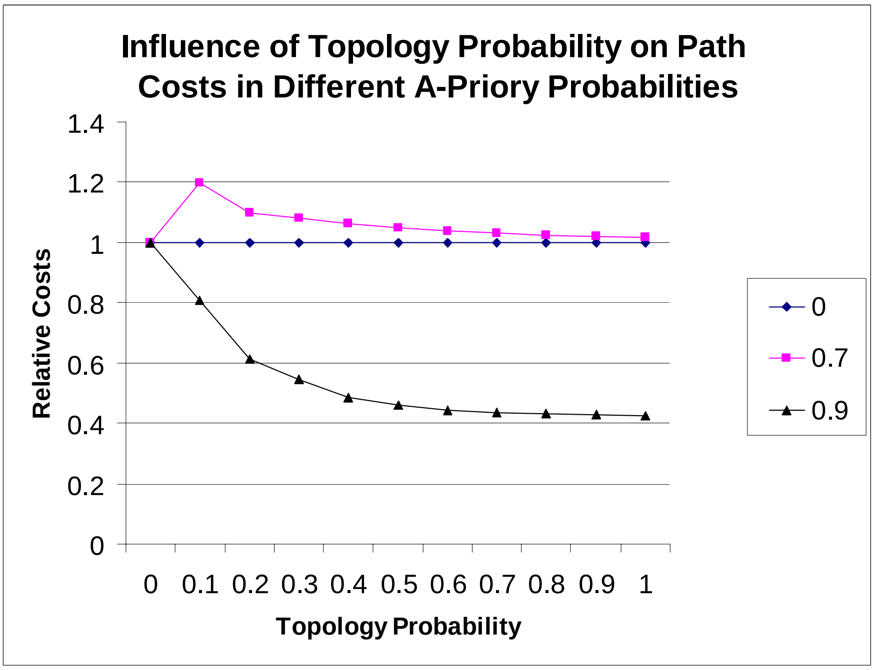

Figure 7: θ_R (Y axe) as function of ξ_T (X axis). This simulation was done with the following parameters: grid size = 4×4 ; ξ_A = "DOWN"; $\omega_{SD} = 700$; ξ_U = 0,0.7 and 0.9.

When $\theta_R = 1$ gossiping doesn't change the gossip traveler's θ_E and we are in the "NEUTRAL" regime. For $\theta_R < 1$ obtaining information leads to decrease in θ_E – the "WIN" regime. In the case of $\theta_R > 1$ obtaining information leads to increase in θ_E, contradicting the desirable outcome – the "LOSE" regime. We are interested in the value of θ_R and less in the value of θ_E since we are mainly interested in the influence of obtaining information on the performance of a given network configuration.

Some of the produced results are presented using the values of θ_A which is the *Average* of θ_R over all the different measured gossip probabilities for a given network configuration.

5.3 Results Discussion

The results presented in Fig. 7 demonstrates the role of obtaining information in different network configurations. In this example ξ_A ="DOWN", thus when ξ_U=0 obtaining information doesn't change the traveler's optimal policy cost. When $\Delta_\omega = \Delta_\xi = 0$ obtaining information will not help the gossip traveler, both travelers are directed in the "right" direction and the gossip traveler has a minimal learning phase, as a result the network operates in the "NEUTRAL" regime. When ξ_U=0.7 obtaining information increases the traveler's optimal policy cost – the network is in the "LOSE" regime. In this case ω_{SE} is such that the non-gossip traveler bypass the stochastic edges, which is justify since ξ_A="DOWN". Therefore, the non-gossip traveler knows the "right" direction. Obtaining information only puzzles the gossip traveler due to Δ_ξ that implies that the learning phase will be relatively large, as a result the gossip traveler will increase his optimal policy cost. Increase in the ξ_T leads to shorter learning phase which leads to smaller θ_R. When ξ_U=0.9 the network is in the "WIN" regime. In this case $\Delta_\omega > \Delta_C$, thus the non-gossip traveler roam to the "wrong" direction. Increase in ξ_T leads to reduce in θ_R since the gossip traveler finishes his learning phase quicker. Fig. 7 also illustrates that the magnitude of the "WIN" effect is substantial larger

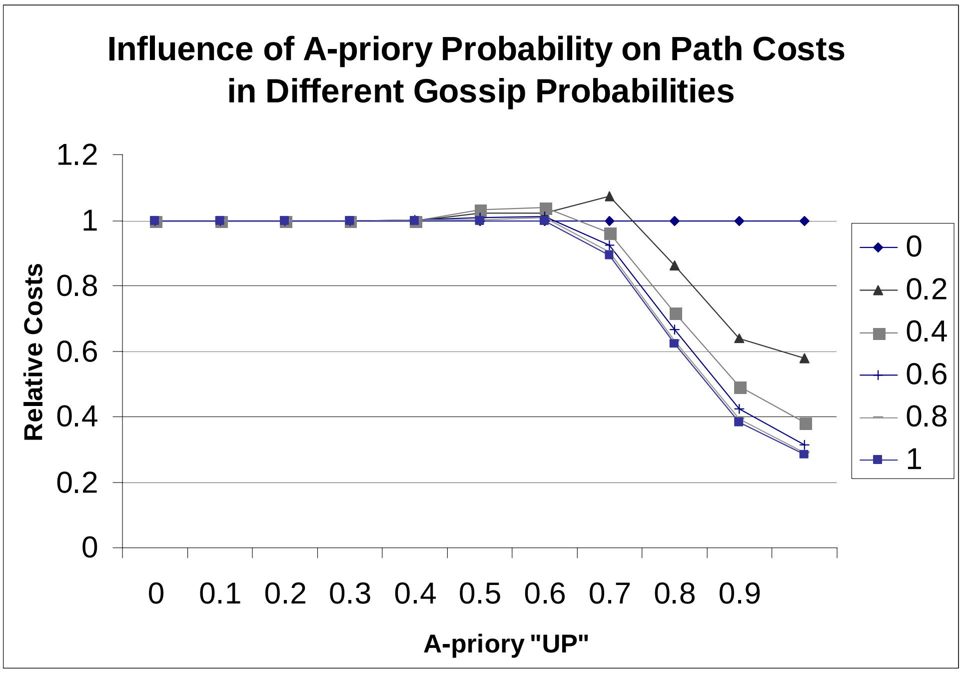

Figure 8: θ_R (Y axis) as function of ξ_U (X axis) for different ξ_T. The different graphs are drawn for $\xi_T = 0, 0.2, 0.4, 0.6, 0.8,$ and 1. This simulation was done with the following parameters: grid size = 4 ; ξ_A = "DOWN"; $\omega_{SD} = 600$.

than the "LOSE" effect.

Fig. 8 depicts the relation between ξ_U and θ_R for different ξ_T values. The curves move between three regimes. When ξ_U is below a threshold value,an increase in ξ_U doesn't change θ_R – the network is in the "NEUTRAL" regime. Then, increase in ξ_U leads to increase of θ_R and the network is in the "LOSE" regime. Further increase of ξ_U moves the network into the "WIN" regime. Comparing the graphs for different ξ_T reveals that in the "NEUTRAL" regime the behavior of all the graphs is almost identical. In the "LOSE" regime, the θ_R peak is reached at $\xi_T = 0.2$. In the "WIN" regime increase in ξ_T leads to decrease in θ_R.

In this graph the network is in the "NEUTRAL" regime when ω_{SE} and $\hat{w}_e$ are similar and the difference between ξ_A and ξ_U is small. In the "LOSE" regime the increase in Δ_ξ leads to a longer learning phase and as a result an increase in θ_R. In the "WIN" regime the increase in Δ_ξ increases the learning phase while an increase in ξ_T decrease it, however the non-gossip traveler moves towards the stochastic edge which increases his θ_E significantly compare to the θ_E of the gossip traveler. As a result, taking both parameters into account, the relative optimal policy cost of the gossip traveler, θ_R, is reduced.

Fig. 9 illustrates the relation between ξ_U and ω_{SD} for averaged ξ_T when the grid size is 4 X 4. Here are several observations from the results:

1) When ξ_U is zero, ω_{SE} is equal to ω_{SD}, in this case the traveler knows a-priory $\hat{w}_e$ and there is no benefit in obtaining information – the network is in the "NEUTRAL" regime.

2) At lower ω_{SD} $(0-200)$ increase of ξ_U leads the network into the "WIN" regime. In this case the stochastic edges weights is similar to the weights of the deterministic edges, therefore information helps the gossip traveler to find the optimal path in the network and decrease his θ_A only moderately.

3) At higher ω_{SD} $(300-)$ increase of ξ_U leads the network from the "NEUTRAL" to the "LOSE" and then to the "WIN" regime. In the "NEUTRAL" and "LOSE" regimes the non-gossip traveler

Averaged Path Costs at Different Network Configurations for 4 x 4 Grid

A-priory "UP" \ "DOWN" Weight	0	100	200	300	400	500	600	700
1	**0.973**	**0.973**	**0.851**	**0.728**	**0.640**	**0.569**	**0.516**	**0.476**
0.9	**0.973**	**0.973**	**0.897**	**0.790**	**0.728**	**0.678**	**0.709**	**0.653**
0.8	**0.973**	**0.972**	**0.914**	**0.847**	**0.814**	**0.739**	**0.774**	**0.868**
0.7	**0.973**	**0.976**	**0.927**	**0.883**	**0.816**	**0.888**	**0.970**	**1.071**
0.6	**0.973**	**0.994**	**0.975**	**0.863**	**0.964**	**0.958**	**1.049**	**1.044**
0.5	**0.988**	**0.985**	**0.972**	**0.935**	**0.966**	**1.030**	**1.029**	**1.029**
0.4	**0.988**	**0.993**	**0.966**	**0.989**	**1.012**	**1.012**	**1.012**	**1.012**
0.3	**1.000**	**0.993**	**0.991**	**1.002**	**1.002**	**1.002**	**1.002**	**1.002**
0.2	**1.000**	**1.000**	**0.998**	**1.001**	**1.001**	**1.001**	**1.001**	**1.001**
0.1	**1.000**	**1.000**	**1.000**	**1.001**	**1.001**	**1.001**	**1.001**	**1.001**
0	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**	**1.000**

Figure 9: θ_A for different values of ω_{SD} (X axis) and ξ_U (Y axis). White cells represents the "WIN" regime, gray the "NEUTRAL" regime and darker gray the "LOSE" regime. This simulation was done with the following parameters: grid size = 4; ξ_A = "DOWN"; θ_A was averaged over $\xi_T = 0$ to 1.

bypass the stochastic edges, therefore in this case obtaining information doesn't help the gossip traveler. When $\xi_U > 0$ obtaining information actually increases the learning phase due to relatively large Δ_ξ and thus there is an increase in the θ_A. Then, with the increase in ξ_U the non-gossip traveler tries to travel through the stochastic edges which leads to increase in his path cost and decrease in θ_A of the gossip traveler that bypass the stochastic edge. The move from the "LOSE" to "WIN" regime is not due to the fact that the gossip traveler decreases his path cost, he actually increases it. However the non-gossip traveler increases his path cost even more due to the fact that now he doesn't bypass the stochastic edges.

4) At higher ω_{SD} (300−), with the increase in ω_{SD} there is increase in the size of the "LOSE" regime. The "LOSE" regime ends when the non-gossip traveler decides to travel through the stochastic edges. This is happening when his ω_{SE} reaches ≈ 200 which is the cost of bypassing the stochastic edges in this example.

5) At higher ω_{SD} (300−), in the "LOSE" regime, the value of θ_A increases with the increase in ξ_U and doesn't change with the increase in ω_{SD}. This phenomena is due to the parameter Δ_ξ, at higher ξ_U there is more probability to paths that leads to the "wrong" direction.

6) In the "WIN" regime, increase in ω_{SD} leads to decrease in θ_A. In higher ω_{SD} the non-gossip traveler travel through the stochastic edges that have increased weights, therefore the gossip traveler can reduce his path cost relatively more.

7) In the "WIN" regime, increase in ξ_U leads to decrease in θ_A. The change here is more moderate and is the result of two parameters. On the one hand with the increase in ξ_U the difference between ω_{SE} and $\hat{w}_e$ is increased which leads to increase in the non-gossip traveler path cost and decrease in θ_A. On the other hand increase of ξ_U leads to increase in the learning phase which leads to the opposite result of increase in θ_A. The outcome of the two parameters is total decrease in θ_A.

Averaged Path Costs at Different Network Configurations for 8 x 8 Grid

A-priory "UP" \ "DOWN" Weight	0	100	200	300	400	500	600	700
1	0.953	0.994	0.939	0.839	0.739	0.663	0.603	0.555
0.9	0.972	0.994	0.938	0.838	0.756	0.678	0.659	0.607
0.8	0.972	0.994	0.936	0.852	0.794	0.754	0.736	0.743
0.7	0.972	0.994	0.947	0.891	0.849	0.824	0.807	0.892
0.6	0.983	0.994	0.977	0.930	0.876	0.901	0.941	1.006
0.5	0.983	0.994	0.976	0.948	0.894	1.004	1.003	1.003
0.4	0.983	0.994	0.976	0.959	0.969	1.002	1.002	1.002
0.3	0.987	0.994	0.991	0.958	1.001	1.001	1.001	1.001
0.2	1.000	1.000	0.999	0.989	1.000	1.000	1.000	1.000
0.1	1.000	1.000	0.999	1.000	1.000	1.000	1.000	1.000
0	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

Figure 10: θ_A for different values of ω_{SD} (X axis) and ξ_U (Y axis). White cells represents the “WIN” regime, gray the “NEUTRAL” regime and darker gray the “LOSE” regime. This simulation was done with the following parameters: grid size = 8; ξ_A = “DOWN”; θ_A was averaged over $\xi_T = 0$ to 1.

Fig. 9 illustrates that for this network configuration gossiping helps in more than half of the cases. In addition, the gain from gossiping is far greater, as much as 50% reduction of the expected path cost, compare to the possible loss which is only up to 7%. However, the fact that one can lose from trying to obtain information dictates the need to understand gossip networks behaviors.

Fig. 10 illustrates that the “LOSE” regime is less significant in larger grid sizes. The reason is that in a small grid the number of steps to destination is small therefore even one wrong step can lead to significant increase in the path cost. In larger networks, where the number of steps is relatively large, the influence of wrong moves is smaller. In real life traffic applications the smaller grid size behavior is more likely due to the small number of options the traveler have especially when the network is in the “DOWN” state - high traffic.

6 Conclusions

This paper presented and investigated a new model for information gathering in stochastic networks, the gossip networks. Gossiping could lead to some unusual phenomena, in our example network the optimal routing policy directs travelers to make a detour in order to gather information and minimize their expected path cost. The optimal traveling policy in gossip network is given by a dynamic programming equations. Although the algorithm that solves the equations, GOSSIP_DP, is intractable in general, we presented two special scenarios where the optimal solution is polynomial in respect to the network size. The first scenario is when the number of stochastic edges is small (G-DWC), the second is when there are small number of realizations in the network (G-DWC). We showed that under the network realizations model (G-DWC) non-gossip travelers on average will need to visit almost all of the network to identify its realization while gossip travelers will identify the network realization after visiting only logarithmic number of the vertices. The influence of gossiping on the optimal policy cost is determined by three parameters: the difference between the stochastic edges actual and expected costs, the topology probability and the difference between the a-priority and actual state of the stochastic edge. We analyzed the relation between these parameters and concluded that it leads to three regimes of operation. In each regime gossiping has different effect on the traveler optimal path cost, “WIN” (reduce), “NEUTRAL” (doesn't change) and “LOSE” (increase). Numerical studies on gossip grid networks confirmed the regime analysis. The numerical studies illustrate that in the networks we study the ”WIN” regime is larger both in size and magnitude and that the “LOSE” regime is more common in small networks.

Future research in gossip networks is needed to understand the influence of gossiping on different network model such as time dependent network for example. Another possible future direction involves developing general approximation algorithms.

References

[APG] Traffic information by monitoring cellular networks. http://www.appliedgenerics.com.

[AR88] G. Andreatta and L. Romeo. Stochastic shortest paths with recourse. *Networks*, 18:193–204, 1988.

[BG92] Dimitri Bertsekas and Robert Gallager. *Data networks*. Prentice-Hall, Inc., second edition, 1992.

[Bri01] Linda Briesemeister. *Group Membership and Communication in Highly Mobile Ad Hoc Networks*. PhD thesis, School of Electrical Engineering and Computer Science, Technical University of Berlin, Germany, November 2001.

[ER01] A. Ebner and H. Rohling. A self-organized radio network for automotive applications. In *Conference Proceedings ITS 2001, 8th World Congress on Intelligent Transportation Systems*, Sydney, Australia, October 2001.

[MAN] Official charter of Mobile Ad-hoc Networks maintained by the IETF. http://www.ietf.org/html.charters/manet-charter.html.

[MJK+00] Robert Morris, John Jannotti, Frans Kaashoek, Jinyang Li, and Douglas S. J. De Couto. CarNet: A scalable ad hoc wireless network system. In *Proceedings of the 9th ACM SIGOPS European workshop: Beyond the PC: New Challenges for the Operating System*, Kolding, Denmark, September 2000.

[ORS93] A. Orda, R. Rom, and M. Sidi. Minimum delay routing in stochastic networks. *IEEE/ACM Transactions on Networking*, 1(2):187–198, 1993.

[PS98] Stefano Pallottino and Maria Grazia Scutellà. Shortest path algorithms in transportation models: classical and innovative aspects. In P. Marcotte and S. Nguyen, editors, *Equilibrium and Advanced Transportation Modelling*, pages 245–281. Kluwer, 1998.

[PT96] G.H. Polychronopoulos and J.T. Tsitsiklis. Stochastic shortest path problems with recourse. *Networks*, 27:133–143, 1996.

[TMC] Traffic information via fm radio. http://www.tmcforum.com.

[TRM] Traffic information to telematic systems. http://www.trafficmaster.net.

[WZ02] S.T. Waller and A.K. Ziliaskopoulos. On the online shortest path problem. *Networks*, 40(4):216–227, 2002.