

Optimal Routing in Gossip Networks

Yuval Shavitt Amir Shay
 School of Electrical Engineering
 Tel Aviv University, Israel
 Email: {shavitt, amirshay}@eng.tau.ac.il

March 1, 2004

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 gossip networks, and demonstrate that an optimal routing policy may direct travelers to make detours to gather information. A dynamic programming equation that produces 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 ordinarily 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 [1] 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 then use the exchanged information to reveal information about the network state and consequently optimize their routing.

There are varieties of real life problems that can benefit from an optimal solution to the problem of routing in gossip networks. This paper will focus on an 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 main 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 [2–4], the other approach is based on ad-hoc communication networks, as proposed by innovative projects such as FleetNet [5], and CarNet [6].

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 recalculate the optimal route from its current location to the destination. For example, if the navigation system receives information that one of the streets in its planned path is blocked it will plan a new path that avoids the blocked roads.

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. FleetNet [5], CarNet [6], 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 implicitly 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.

The problem of *Shortest Path Routing* was investigated extensively in the literature. 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 [7]. Note that the G-IWC model is a generalization of the G-DWC model with substantially more states. The rationale 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 network 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

¹This paper focuses on the routing of mobile agents on the “roads networks” and not on the routing of data packets on the “information networks”.

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 [8] assume that once a blockage is encounter a recourse path that consists of only deterministic roads is used. Orda, Rom, and Sidi [9] 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 [7] 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 [10] solved a model with dependency between successor roads and a model with time dependency for the same road.

The primary contribution of this paper is in the introduction and analysis of the gossip model and the new directions it opens for building P2P mobile systems. The introduction of information exchange 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. Further more, we were able to quantify an optimal policy that balance between information gathering costs and path costs. Other main contribution is the regime state diagram we produced. Using the diagram one can determines the influence of gossiping on the traveling costs in different networks characteristics.

The rest of the paper is 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.

2 Model and Definitions

2.1 The formal model

The network² 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 $\bar{w}_e = \sum_{s=1}^{\beta_e} w_e^s q_e^s$, 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 at the time of travel. 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*, $I\{\cdot\}$. For example, an information vector

²As 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 implicitly, it is included in the gossip probability presented below.

of a traveler could look like this: $I = \{\hat{w}_1, X, \hat{w}_3, X, \dots, X, X, \hat{w}_\delta\}$. 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

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

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

$$l_D = \sum_{i=1}^R \binom{R}{i} = 2^R - 1 \quad (2)$$

When two or more TAs are within communication range they can exchange their information vectors in order to gain missing data. The *gossip probability* is the probability that when a TA traverses an edge it will update his information vector.

$$P(s, s', T(i, j)) = \mathcal{P}\{I(j) = s' | I(i) = s, T(i, j)\} \quad (3)$$

where $s, s' \in I$ are the information vector before and after the edge (i, j) traversal, respectively, $I(i)$ is the information vector at vertex $i \in V$, and $T(i, j)$ is the *topology probability*. The topology probability is the probability that a TA will receive information from other TAs during the traversal on an edge. 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. Assuming that there are “enough” mobile agents in the network $T(i, j)$ is a vector of probabilities, where each element corresponds to some stochastic network edge. For example, $T(i, j) = \{1, 0.5, \dots, 0\}$ means that on average when the TA traverses edge (i, j) it will learn about stochastic edges 1, 2, and δ with probability 1, 0.5, and zero, respectively. The gossip probability depends on the topology probability and on the information vector before and after the edge traversal. For example, the probability to change an information vector element from $\{\dots, \hat{w}, \dots\}$ to $\{\dots, X, \dots\}$ is zero. Regardless of the topology probability, a known weight can not be changed into unknown.

In this paper we are looking for the optimal routing policy of a TA that starts at the source vertex s with information vector $I(s)$ and travels to a destination vertex t . We assume that the TA knows a priori 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 Assumptions and Reality

In this section we will analyze the formal assumptions in our model and relate them to real life scenarios in transportation networks. The first assumption is that the network is time independent. In many situations, a driver can assume that during his commute (30 to 60 minutes) the traffic patterns in his area doesn't change significantly. Thus, in many cases, an optimal routing policy calculated at the beginning of the journey will yield satisfying results throughout the journey.

Another assumption is that the agent knows a priori the network structure, edge weight distribution and topology probability. While network structure can be obtained from any GIS, the edge weight and topology probability are calculated from historical information gathered over time. Currently there are several commercial and academic projects that use historical data to predict future traffic patterns, for example the MIT's DynaMIT project [11]. While the edge




Figure 1: An example of the influence of gossiping on routing.




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

weight distribution can be computed directly from the historical traffic data, in order to compute the topology probability one needs information about the agents movement in the network. Given that information, we can calculate and record fairly easily the edge weight distribution and topological probability for a given time. For example, we will have one distribution for morning commute, second for evening commute, third for holidays etc. Then, each time the agent will compute his optimal routing policy using gossip networks time independent algorithm he will use the appropriate distributions.

2.3 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 $q_{jj'}^u = \xi_U$ the edge is in the “UP” state where $w_{jj'}^u = 1$, and with probability $q_{jj'}^d = (1 - \xi_U)$ the edge is in the “DOWN” state where $w_{jj'}^d = 10000$. 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 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 .

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 priori knowledge of the state of the edge (j, j')

$$C(s, t, \{\cdot\})_k = 1001 \quad (4)$$

³The 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).

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

$$C(s, t, \{X\})_j = 10002(1 - \xi_U) + 3\xi_U \quad (5)$$

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:

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

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

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

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 topological probability, ξ_0 , such that for $\xi_T \geq \xi_0$

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

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 curves in Fig. 2 is when Eq. 5 and Eq. 4 are equal; in this example at $\xi_U = 0.90028$.

3 The Routing Algorithm

3.1 Solution approach

The problem of finding the optimal routing in gossip networks belong to the class of online decisions problems. In these problems an agent is faced with the opportunity of influencing the behaviors of a probabilistic system as it evolve. At each step the agent receives information about the system state and performs an action accordingly. His goal is to chose a sequence of actions which causes the system to perform optimally with respect to some predetermine criteria. In the literature such problems can be found under the topics of *Markov decision processes*, *stochastic programming*, and *optimal control*. Similar to other online decisions problems, we solve the problem of optimal routing in gossip networks using dynamic programming and, in general, share the same "curse of dimensionality", which leads to intractable solutions. It is well known in online decisions problems that information pays off, in our algorithm we were able to quantify the importance of information.

A traveler starts his journey from 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:

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

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 s it has the information vector $I(s)$. $P(I(s), I(r), T(s, r))$ is the gossip 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 priori 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:

$$C^*(s, t, I(s)) = \min_{r \in \mathcal{N}_s} \{ \hat{w}_{sr} + \sum_{I(r) \in B(I(s), (s, r))} P(I(s), I(r), T(s, r)) \cdot Q(I(r)) \cdot C^*(r, t, I(r)) \} \quad (8)$$

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

In Bellman-Ford's dynamic programming algorithm for deterministic shortest path [1] one finds for each vertex the shortest path to a destination. In 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, $TBL(u)$, 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:

$$DD(u, s_k) = \hat{w}_{uv} + \sum_{m=1}^l P(s_k, s_m, T(u, v)) Q(s_m) DD(v, s_m) \quad (9)$$

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 , ($m = 1 \dots l$). 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 in [12]. The algorithm can be used to produce the optimal routing policy in gossip networks by the following steps: 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 $PN(u, s_k)$.

The proof that the algorithm GOSSIP_DP provides the optimal solution for routing in gossip network is a direct extension of a general dynamic programming optimality proof [13]. The proof is omitted due to space limitations.

Due to space limitations we state here the following theorem without proof, which appears in [12].

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

Although the optimal solution to the gossip networks problem is intractable in general, we presented above two special cases where the optimal solution is tractable regardless of 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 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. For the simplicity of the discussion we use the following assumptions: The network is in the G-IWC model with one stochastic edge, and the traveler's expected shortest path from source to destination contains a stochastic edge. 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 as we demonstrate in the numerical analysis in the next section.

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 change 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 priori probability to be in the "UP" state (ξ_U) Next we will explain the influence of each parameter.

The magnitude of the difference between the traveler's a priori 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 priori 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. 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. 3 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 priori before the start of the journey and has no 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

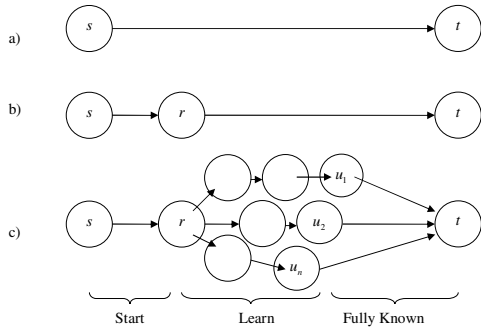


Figure 3: The different possible paths a traveler can have for different topology probabilities. (a) No gossiping, (b) Maximal gossiping, and (c) In between.




Figure 4: The regime state diagram determines the influence of gossiping on routing in different networks characteristics

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 priori 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 priori 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 operating regime that the traveler experiences is determined by the combined values of the parameters, Δ_ω, ξ_T , and Δ_ξ . Fig. 4 is a state diagram that illustrates the influence of the parameters on the network regime. When Δ_ω is below some threshold, Δ_C , the a priori 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 “NETRAL” 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 “NETRAL” 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 “NETRAL” 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

We preformed an extensive simulation study on grid networks under the G-IWC model. The results obtained in the simulations confirmed the discussion in Section 4.1 regarding the char-

acteristics of optimal routing in gossip networks. Due to space limitations the reader is referred to [12] for details.

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 equation. Although the algorithm that solves the equations, GOSSIP_DP, is intractable in general, we presented two tractable special scenarios. We are currently working on heuristics for solving the general case.

We showed in [12] that the influence of gossiping on the optimal policy cost is determined by several parameters. Surprisingly, in some parameter combinations gossiping can lead to slightly longer expected cost for the traveler. However, in most cases, a traveler can benefit, and in some cases significantly, from gossiping.

References

- [1] D. Bertsekas and R. Gallager, *Data networks*, 2nd ed. Prentice-Hall, 1992.
- [2] "Traffic information by monitoring cellular networks," <http://www.appliedgenerics.com>.
- [3] "Traffic information via FM radio," <http://www.tmcforum.com>.
- [4] "Traffic information to telematic systems," <http://www.trafficmaster.net>.
- [5] 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.
- [6] R. Morris, J. Jannotti, F. Kaashoek, J. Li, and D. S. J. De Couto, "CarNet: A scalable ad hoc wireless network system," in *the 9th ACM SIGOPS European workshop: Beyond the PC: New Challenges for the Operating System*, Kolding, Denmark, Sept. 2000.
- [7] G. Polychronopoulos and J. Tsitsiklis, "Stochastic shortest path problems with recourse," *Networks*, vol. 27, pp. 133–143, 1996.
- [8] G. Andreatta and L. Romeo, "Stochastic shortest paths with recourse," *Networks*, vol. 18, pp. 193–204, 1988.
- [9] A. Orda, R. Rom, and M. Sidi, "Minimum delay routing in stochastic networks," *IEEE/ACM Transactions on Networking*, vol. 1, no. 2, pp. 187–198, 1993.
- [10] S. Waller and A. Ziliaskopoulos, "On the online shortest path problem," *Networks*, vol. 40, no. 4, pp. 216–227, 2002.
- [11] "Dynamit: a simulation-based system for traffic prediction," <http://web.mit.edu/its/dynamit.html>.
- [12] Y. Shavitt and A. Shay, "Optimal routing in gossip networks," Tel Aviv University, Israel, Tech. Rep., 2003, <http://www.eng.tau.ac.il/Utils/reportlist/reports/reports2003.html>.
- [13] D. Bertsekas, *Dynamic Programming and Optimal Control*, 2nd ed. Athena Scientific, 2000.