Topology versus Link Strength for Information Dissemination in Networks

Lorenzo Severini, Mahmoud El Chamie, Giovanni Neglia

Topology versus Link Strength for Information Dissemination in Networks

Lorenzo Severini[1] and Mahmoud El Chamie[1][2] and Giovanni Neglia[2]

[1] *Université Nice Sophia Antipolis (UNS) - 28 Avenue Valrose, 06100 Nice, France*
[2] *INRIA Sophia Antipolis - Méditerranée - 2004 Route des Lucioles, B.P. 93 06902 Sophia Antipolis, France*

L'information peut circuler dans un réseau de communication par les liens reliant les nœuds. La topologie du réseau et la force des liens sont deux facteurs qui influent sur la vitesse de propagation de l'information dans le réseau. Dans cet article, nous montrons que la topologie peut avoir un rôle plus important que la force des liens pour la vitesse de propagation de l'information. En particulier, nous considérons un processus itératif de propagation de croyance comme dans les protocoles de consensus moyen où chaque nœud dans le réseau a une certaine croyance (exprimée par un nombre réel), et à chaque itération il met à jour sa croyance en calculant une moyenne pondérée de sa croyance et de celles des ses voisins. Nous montrons que l'ajout de liens peut conduire à une augmentation de la vitesse de convergence du protocole de consensus plus significative que les techniques d'optimisation des poids. Les simulations sont effectuées sur différentes topologies : anneaux, grilles, graphes aléatoires (Erdos Renyi, graphes géométriques aléatoires) et le graphe d'échange de courriels chez Enron.

Keywords: topology optimization, belief propagation, average consensus, weight selection, 2-hop optimization

1 Introduction

Information can disseminate in a network through communication links : virus spread, shared videos, and belief propagation are all examples of information dissemination in connected networks. Three main factors can affect the speed of propagation of a belief in a network. The first factor is the belief's *value* : recent, emotional, and high valued articles are more probable to spread than old uninteresting ones. Second, the *strength of the connection* between nodes can affect how fast information spread. Best friends are more likely to share common beliefs and are more affected by each other's opinions. The third factor is the *network topology* : the more the connections in a network the faster information can spread. We address here optimization techniques applied on networks to speed up the propagation of information.

To study the belief propagation process, we model the network as an undirected graph $G = (V,E)$ of nodes ($V = 1...n$) where any two nodes (or neighbors) that are able to communicate are connected by a link ($E = 1...m$). Each node i has a certain belief and its value at time t is $x_i(t) \in \mathbb{R}$. Let d_i be the degree of a node i in the graph G. The strength of the connections in this graph is represented by a weight matrix W where w_{ij} is the weight given by node i to the strength of the link $l \sim (i,j)$ connecting it to j. We consider an iterative discrete time system where nodes at every iteration t share their opinion with their neighbors and perform a weighted average update for their opinion following the *average consensus protocol* ([RBA05]) :

$$x_i(t+1) = w_{ii}x_i(t) + \sum_{j \in N_i^G} w_{ij}x_j(t), \text{for } i = 1...n, \tag{1}$$

where N_i^G is the set of neighbors of node i in G (so we have $|N_i^G| = d_i$). Under some assumptions on the weights, nodes' beliefs are guaranteed to converge to the average of all initial beliefs (to reach *consensus*), in particular $\lim_{t \to \infty} x_i(t) = x_{ave} \ \forall i$ where $x_{ave} = \frac{1}{n}\sum_i x_i(0)$. In this framework the speed of belief propagation is then characterized by the speed of convergence of $x_i(t)$ to x_{ave}.

Different algorithms for selecting the weights in equation (1) have been proposed in the literature. Simple algorithms do not provide any guarantee on the speed of convergence, and more sophisticated ones are

resource consuming because they select the weights solving complex optimization problems [BGPS06]. As mentioned above, the weights are not the only factor to speed up information diffusion, but also the topology plays an important role. In this paper we evaluate if simple changes to the network topology may speed up belief propagation more than complex weight optimization techniques. In particular we compare the performance of the average consensus protocol in the two following scenarios. In the first scenario, the topology is unchanged and weights are selected according to commonly used algorithms, including those that guarantee faster convergence to consensus. In the second scenario, the simplest weights selection algorithms are used, but direct links to 2-hop away nodes are added to the original graph G, then shrinking by 2 the network diameter. We denote by G^2 (the *square graph*) this denser graph. Practically speaking this topological change does not require to really add new links : it can be obtained by forwarding nodes' beliefs 2-hop away, so that a generic node i is aware of all the nodes' beliefs in the extended neighborhood $N_i^{G^2} = \cup_{j \in N_i^G} N_j^G$. In what follows we are going to consider this way to operate. It has also the advantage to allow us to quantify the cost of the topological change in terms of an increase of communication overhead. The comparison is carried on for different graph topologies : rings, square Grids, random graphs (Erdos-Renyi with link existence probability P), Random Geometric graphs (with connectivity radius R), and real world network topologies as Enron internal email exchange network [SA04].

The simulations on these graphs show two main interesting results. The first result is that simple weight selection algorithms can achieve significantly faster convergence on the denser graph G^2 than any weight optimization technique on the original graph G. This improvement comes at the cost of an increase of communication overhead in the network. Our second (less expected) result is that, for a given weight selection algorithm, the convergence is faster on G^2 than on G even when the number of messages is equal. Because of this, simpler weight selection algorithms on G^2 can achieve performance similar to more complex ones on G. These results suggest that topological optimization can have a more important role than weight optimization techniques to speed up information propagation.

The paper is organized as follows : in section 2 we provide the background on the asymptotic convergence speed of consensus protocols and describe different weight selection algorithms. In section 3, we compare the performance in the two scenarios described above. Section 4 concludes the paper.

2   Convergence Speed and Weight Selection Algorithms

We define W to be the weight matrix, whose elements are the weights used by the nodes in (1) (i.e. $(W)_{ij} = w_{ij}$). The authors of [XB04] show that for symmetric weight matrices, the consensus protocol converges if and only if $W\mathbf{1} = \mathbf{1}$ and $\mu(W) < 1$, where $\mathbf{1}$ is a vector of all ones and $\mu(W)$ is the second largest eigenvalue in magnitude of the matrix W. They also show that μ determines the asymptotic convergence speed (for the worst case initial opinions) : the smaller μ the faster the convergence. Then the optimal (symmetric) weight matrix is the solution of the following optimization problem :

$$\begin{aligned}
\text{Argmin}_W \quad & \mu(W) \\
\text{subject to} \quad & W\mathbf{1} = \mathbf{1}, \\
& W \in S_G^n,
\end{aligned} \tag{2}$$

where S_G^n is the set of symmetric $n \times n$ real matrices such that $s_{ij} = 0$ if $(i,j) \notin E$.

Let A be a weight selection algorithm for the average consensus protocol, we denote by $_GW_{(A)}$ the weight matrix generated by A on the graph G. In order to evaluate the effect of weight selection algorithms and of topology on convergence speed, we are going to compare $\mu(_GW_{(A)})$, where A is the optimal weight selection algorithm (that solves problem (2)) or one of its approximations, and $\mu\left(_{G^2}W_{(B)}\right)$, where B is a simpler weight selection algorithm.

In our analysis we have considered the following weight selection algorithms :
– **Max Degree Weights (MD)** [XBK07] : the same weight is selected for all links, it is a function of the maximum degree in the network, i.e. $w_{ij} = w_{ji} = \frac{1}{\Delta+1} \; \forall l \sim (i,j) \in E$, $\Delta = \max_i \{d_i\}$. The advantage of this algorithm is its simplicity and the fact that it works also on dynamic graphs, but it does not offer any guarantees on the speed of convergence.

	RGG $n = 50$		ER $n = 50$		Ring		Grid		Enron
	$R = 0.25$	$R = 0.3$	$P = 0.08$	$P = 0.12$	$n = 50$	$n = 100$	$n = 36$	$n = 64$	
$\mu(_G W_{(FDLA)})$	0.9390	0.8668	0.8511	0.7241	0.9921	0.9980	0.9210	0.9210	0.8287
$\mu(_{G^2} W_{(LD)})$	0.9070	0.8058	0.8328	0.7144	0.9843	0.9961	0.8523	0.9155	0.8208

TAB. 1: The effect of Graph Density versus Weight Optimization on the speed of convergence of consensus protocols. The table shows the comparison on different graph topologies between the speed of convergence of : **(1)** the simple weight selection algorithm (LD) on the graph G^2 quantified by $\mu(_{G^2} W_{(LD)})$ and **(2)** the best weight selection algorithm (FDLA) on the graph G quantified by $\mu(_G W_{(FDLA)})$.

- **Local Degree or Metropolis Weights (LD)** [XB04] : the weights on the links depend only on the degrees of the incident nodes : $w_{ij} = w_{ji} = \frac{1}{\max\{d_i, d_j\}+1} \ \forall l \sim (i,j) \in E$. In addition to the advantages of the MD, this algorithm can be implemented in distributed way without nodes being aware of global information (as the maximum degree for example). It does not provide any guarantee on the speed of convergence.
- **Optimal Constant Weights (OC)** [XB04] : is the solution of (2) with the additional constraint that all the weights are equal, i.e. $w_{ij} = w_{i'j'}, \ \forall l \sim (i,j), l' \sim (i',j') \in E$. In this case the weights depend on the eigenvalues of the Laplacian L of the graph : $w_{ij} = w_{ji} = \frac{2}{\lambda_1(L)+\lambda_{(n-1)}(L)} \forall l \sim (i,j) \in E$.
- **Trace Minimization-p(TM-p)** [ECNA12] : approximates (2) considering as objective function $Tr(W^p)$ instead of $\mu(W)$. The authors show that the parameter p trades off approximation quality (and then convergence speed) with complexity.
- **Fastest Distributed Linear Averaging (FDLA)** [XB04] : this algorithm solves exactly the problem (2) using a semi-definite programming (SDP) formulation. It is then optimal in the sense that $\mu(_G W_{(FDLA)}) \leq \mu(_G W_{(A)})$ for any graph G and any weight selection algorithm A. Its main disadvantage is the high complexity, the problem becomes untractable already for graphs having few thousands of links.

3   Graph Density versus Weight Optimization

We first compare the performance (the asymptotic speed of convergence) of the consensus protocol on the denser graph G^2 when weights are selected according to the *LD* algorithm with the performance on the original graph G when the optimal weight selection algorithm FDLA is used. Results in Table 1 show that on all the topologies considered $\mu(_{G^2} W_{(LD)}) < \mu(_G W_{(FDLA)})$ and then $\mu(_{G^2} W_{(LD)}) < \mu(_G W_{(A)})$ for any algorithm A. Then the higher graph density provides a more significant improvement than the optimal choice of links weights.

We now evaluate the communication overhead of the two approaches in terms of the number of messages sent. Equation (1) requires that nodes at each iteration t use the beliefs of their neighbors (node i uses $x_j(t)$ for all $j \in N_i^G$). Therefore, each node must receive at every iteration these beliefs and the total number of messages M sent in the system will be $M = 2 \times m$ where m is the number of links in the graph. On G^2, N_i^G in equation (1) is replaced by $N_i^{G^2}$. As we mentioned above, it is possible to mimic the consensus protocol on G^2 using only the links in G. In this case the operation requires 2 steps. First each node broadcasts its belief to its neighbors in N_i^G. Then, each node sends another broadcast message to its neighbors in N_i^G with all the beliefs that it has collected during the first step. In this way every node gets to know the beliefs of all the nodes in $N_i^{G^2}$. The total number of messages is then twice as larger than in the first scenario[†]. For this reason, we decided to compare the speed of convergence in the two scenarios when the number of messages being equal. This corresponds to consider that the consensus protocol on G performs two weighted linear iterations according to (1) for each linear iteration on G^2. Another possible way to interpret this comparison is that if the duration of an iteration is determined by the time needed to transmit one message on a link, then a consensus protocol iteration on G^2 requires twice as much time than one on G. It

---

[†] We observe here that the messages sent in the second step have usually a larger data payload than those sent in the first step, because they carry many belief values. Here we assume that the number of messages is an adequate metric to evaluate the performance, as for example is the case if the packet header is much larger than the data payload for this application.

	RGG $n = 50$		ER $n = 50$		Ring		Grid		Enron
	$R = 0.25$	$R = 0.3$	$P = 0.08$	$P = 0.12$	$n = 50$	$n = 100$	$n = 36$	$n = 64$	
$\mu(_G W^2_{(MD)})$	0.9665	0.9274	0.9036	0.8327	0.9894	0.9974	0.8957	0.9401	0.9761
$\mu(_{G^2} W_{(MD)})$	0.9319	0.8577	0.8967	0.7923	0.9843	0.9961	0.8730	0.9240	0.9057
$\mu(_G W^2_{(LD)})$	0.9493	0.8951	0.8591	0.7572	0.9894	0.9974	0.8876	0.9364	0.9726
$\mu(_{G^2} W_{(LD)})$	0.9070	0.8058	0.8328	0.7144	0.9843	0.9961	0.8523	0.9155	0.8208
$\mu(_G W^2_{(OC)})$	0.9378	0.8677	0.8363	0.7276	0.9843	0.9960	0.8662	0.9239	0.9534
$\mu(_{G^2} W_{(OC)})$	0.8761	0.7543	0.8177	0.6650	0.9751	0.9937	0.7919	0.8776	0.8277
$\mu(_G W^2_{(TM-2)})$	0.9419	0.8800	0.8334	0.6749	0.9894	0.9974	0.8857	0.9359	0.9143
$\mu(_{G^2} W_{(TM-2)})$	0.8900	0.7565	0.7078	0.4590	0.9843	0.9961	0.8403	0.9119	0.5568
$\mu(_G W^2_{(FDLA)})$	0.8817	0.7513	0.7244	0.5243	0.9843	0.9960	0.8482	0.9126	0.6868
$\mu(_{G^2} W_{(FDLA)})$	0.7591	0.5478	0.5219	0.3098	0.9691	0.9922	0.7241	0.8343	-

TAB. 2: Asymptotic speed of convergence on G and G^2 given the same communication overhead.

is easy to evaluate the speed of convergence of the "accelerated" consensus protocol that performs two linear iterations every time unit. In fact it can be checked that this corresponds to use as a weight matrix $(_G W_{(A)})^2$ [BGPS06]. Then, the asymptotic speed of convergence is determined by $\mu(_G W^2_{(A)})$. Note that the following equation holds : $\mu(_G W^2_{(A)}) = (\mu(_G W_{(A)}))^2$. Simulation results in Table 2 show that $\mu(_{G^2} W_{(A)}) < \mu(_G W^2_{(A)})$ for any algorithm A of previous sections. Then the denser topology leads to faster convergence speed even when the number of messages is equivalent. For this reasons, simple weight selection algorithms as LD on G^2 can still outperform more complex ones like $TM - 2$ or OC (the results from Table 2 show that $\mu(_{G^2} W_{(LD)}) < \mu(_G W^2_{(TM-2)})$ and $\mu(_{G^2} W_{(LD)}) < \mu(_G W^2_{(OC)})$ on most of the topologies) and also achieve in some cases results very similar to FDLA (e.g. on the grid).

4 Conclusion

In this paper we consider consensus protocols as a general model for belief propagation. The speed of convergence of consensus protocols is the indicator used to measure the speed of propagation of beliefs. We show that a higher graph density can improve convergence speed more than an accurate weight selection. This result suggests that topology may play a more important role for information spread than the strength of the links.

Références

[BGPS06] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized Gossip Algorithms. *IEEE Trans. Inf. Theory*, 52 :2508–2530, June 2006.

[ECNA12] Mahmoud El Chamie, Giovanni Neglia, and Konstantin Avrachenkov. Distributed weight selection in consensus protocols by schatten norm minimization. INRIA Research Report RR-8078, INRIA, October 2012.

[RBA05] Wei Ren, R.W. Beard, and E.M. Atkins. A survey of consensus problems in multi-agent coordination. In *Proceedings of the American Control Conference, 2005*, pages 1859 – 1864, vol. 3, June 2005.

[SA04] J. Shetty and J. Adibi. The enron email dataset database schema and brief statistical report. Technical report, University of Southern California, Information Sciences Institute Technical Report, 2004.

[XB04] Lin Xiao and Stephen Boyd. Fast linear iterations for distributed averaging. *Systems and Control Letters*, 53(1) :65 – 78, 2004.

[XBK07] Lin Xiao, Stephen Boyd, and Seung-Jean Kim. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1) :33–46, 2007.