Topological determinants of convergence rate for gossip algorithms

Christel Sirocchi[1] and Alessandro Bogliolo[1,2]

[1]Dipartimento di Scienze Pure e Applicate, Università degli Studi di Urbino, 61029 Urbino, Italy
[2]DIGIT srl, 61029 Urbino, Italy

Keywords— gossip algorithms, distributed aggregation, convergence rate, social networks, wireless sensor networks

Abstract

Distributed data aggregation has drawn considerable attention in automated control, signal processing, robotics, and communication. It applies to all settings characterised by given operational conditions: an underlying network restricting communication, nodes interacting exclusively with their neighbours, lack of a centralised entity, and network topology not wholly known to the individual nodes. Notable examples are Wireless Sensor Networks (WSN), made of devices equipped with measuring sensors and low-range communication capability, and Social Networks (SN), consisting of individuals who hold opinions and share them with their peers. Distributed average is an instance of distributed computation aimed at calculating the global average by iterating local calculations until the desired level of convergence. To this aim, average consensus algorithms realise simultaneous updates of all nodes at discrete time points while gossip algorithms apply pairwise asynchronous and randomised interaction schemes. The latter is more suited to model real networks but is harder to characterise mathematically due to the added randomness of neighbour selection. Much effort has been put toward characterising the efficiency and improving existing algorithms by deriving upper and lower bounds and designing quicker algorithms. In contrast, our work focuses on isolating the effect of topological network features on convergence rate, so that faster convergence can be engineered by acting on the underlying network rather than on the averaging algorithm.

The study analyses the convergence rate of gossip algorithms in four network families: Erdos-Renyi (ER) and Geometric Random Graphs (GR) are standard models for WSNs, while Small World (SW) and Scale Free (SF) networks are more suited for SNs. We deploy simulations to evaluate, for each graph family, the topological limiting factors, the most predictive graph metrics, and the most efficient random algorithm with respect to convergence. A regression model built on network metrics predicts the convergence rate with high accuracy and confirms that topological features determine the behaviour.

In all experiments, the communication network is a connected undirected graph where each node is assigned an IID gaussian random variable whose interpretation (environmental measure, opinion, etc.) depends on the considered setting. Each node initiates an interaction at the times of a rate 1 Poisson process, according to the asynchronous time model. The interaction entails that both nodes set their values equal to the average of their current values. The gossip scheme, represented by a stochastic matrix, defines the probability p that each node interacts with a given neighbour. The study compares four algorithms: (1) random selection and (2) ordered selection, with p equal for all neighbours, (3) degree selection, with p proportional to the neighbours' degree, and (4) distance selection, with p inversely proportional to the number of shared neighbours. The accuracy is expressed as the variance of the values normalised by the initial variance. Established that, after a transient phase, the logarithm of the accuracy is inversely proportional to the number of interactions, the convergence rate is calculated as the slope of this linear stationary regime. Thirty metrics are computed for each graph, grouped into i) local metrics, ii) global metrics, and iii) spectral metrics. Graphs are obtained by varying continuously and within a suitable range each parameter of the generating model to explore the space of all possible graphs. We use the Watts-Strogatz model (with rewiring probability p_r) for SW graphs and the Holme and Kim algorithm (with clustering probability p_c) to generate SF graphs. All convergence measures are derived experimentally through event-driven simulations.

In ER graphs, the average degree d determines the convergence rate (Fig. 1) modelled by an exponential plateau with maximum Y_M. For $d < 20$, increasing the number of edges speeds up the convergence by orders of magnitudes, while for $d > 30$, maximum convergence has been achieved and adding edges is a waste of resources. Distance metrics (diameter, radius, eccentricity) are most predictive of the algorithm performance (R-squared > 0.80). In SW graphs, the average degree affects the convergence rate according to an exponential plateau where Y_M depends linearly on p_r for $p_r \in [0, 0.5]$. For these graphs, p_r is the limiting factor, and heterogeneity metrics (degree standard deviation and entropy) are most predictive, so rewiring existing edges is generally more effective than adding new ones in speeding up convergence. SF graphs show a convergence rate similar to ER graphs for $p_c < 0.8$, best predicted by its algebraic connectivity. However, high clustering ($p_c = 1$) substantially reduces the convergence rate due to the redundancy of the information shared. In ER and SF graphs, the ordered selection is on average 10% faster than random selection for $d < 30$ despite having an identical probability matrix, while degree selection is the slowest. In SW and GR graphs, distance selection is significantly faster, but convergence largely depends on the initial position of the values assigned to nodes. A Multilinear Regression Model built on all metrics displayed high predictive power (R-squared > 0.95) and allowed predicting convergence rates in randomly generated graphs.

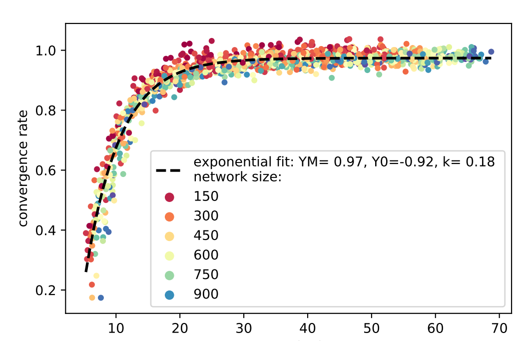

Figure 1: Simulations on 1000 ER graphs.