

Robustness of Self-Organizing Consensus Algorithms: Initial Results from a Simulation-Based Study

Alexander Gogolev¹ and Christian Bettstetter^{1,2}

¹ Institute of Networked and Embedded Systems, University of Klagenfurt, Austria

² Lakeside Labs GmbH, Klagenfurt, Austria

Abstract. This short paper studies distributed consensus algorithms with focus on their robustness against communication errors. We report simulation results to verify and assess existing algorithms. Gacs-Kurdyumov-Levin and simple majority rule are evaluated in terms of convergence rate and speed as a function of noise and network topology.

Keywords: Networked systems, self-organization, consensus, robustness

1 Introduction

The overall objective of our research is to analyze *robustness* characteristics of *features of self-organization* in networks. We recognize any feature of a network as feature of self-organization if it is capable of performing actions aimed at maintaining the network function in a completely distributed manner. Each node has only local view, and simple rules are applied. Systems presenting such features are recognized as “having an ability of self-organization.” We use the term robustness as ability of the system (or an algorithm as part of the system) to maintain a certain level of network function despite changes or noise in the environment and/or changes in system structure, including partial failure.

This short paper is our first step in this research direction. It studies self-organizing *consensus* algorithms in networks with focus on robustness toward communication errors between nodes. Consensus algorithms can be applied to a broad spectrum of tasks, ranging from distributed version management over analysis of social networks to forecasting of crowd behavior.

Our work is inspired by an article of Moreira *et al.* [1]. That paper studies the impact of noise leading to communication errors on distributed consensus. Gacs-Kurdyumov-Levin (GKL) [2] and simple majority rule (SMR) [3, 4] are tested for convergence under different noise levels in a small-world network [5]. It was shown that noise can actually improve performance of simple algorithms in certain setups [1]. The work at hand slightly extends that work using almost the same modeling assumptions but investigating a larger set of parameters.

2 Modeling Assumptions

2.1 Setup and Notation

The system is modeled as a cellular automata graph [6]. Given is a one-dimensional grid of n nodes where each node is connected to its k closest neighbors. At a time $t = 0$, nodes are initially assigned with a random binary state σ whose value is either $+1$ (white color) or -1 (black color). At every time step $t > 0$, all nodes update their state following a given consensus rule based on their own state and their neighbors' states. The network is expected to converge within $2n$ time steps to a single value, which corresponds to the initial majority of value distribution [1]. The state of node i at time t is called $\sigma_i[t]$; the state that node i receives for the state of node j is called $\tilde{\sigma}_j^i[t]$.

2.2 Consensus Algorithms

Let us describe the GKL and SMR rules. According to GKL, a node i computes its new state $\sigma_i[t + 1]$ using its current state $\sigma_i[t]$ and the current state of the first and third nearest neighbors to the left or right from the node. The current value defines the side from which the neighbors are chosen. We have [1]

$$\sigma_i[t + 1] = \begin{cases} G(\tilde{\sigma}_i^i[t] + \tilde{\sigma}_{i-1}^i[t] + \tilde{\sigma}_{i-3}^i[t]) & \text{if } \tilde{\sigma}_i^i[t] = -1 \\ G(\tilde{\sigma}_i^i[t] + \tilde{\sigma}_{i+1}^i[t] + \tilde{\sigma}_{i+3}^i[t]) & \text{if } \tilde{\sigma}_i^i[t] = +1 \end{cases}, \quad (1)$$

where the nodes $i - 1$ and $i - 3$ are the first and third neighbors to the left (nodes $i + 1$ and $i + 3$ are to the right). The update function G is [1]

$$G(x) \equiv \begin{cases} -1 & \text{for } x < 0 \\ +1 & \text{for } x > 0 \end{cases}. \quad (2)$$

Using SMR, a node calculates its new value on basis of its current state and the state of its closest k neighbors from both sides. i.e.,

$$\sigma_i[t + 1] = G\left(\sum_{j=i-k}^{i+k} \tilde{\sigma}_j^i[t]\right). \quad (3)$$

2.3 Noise

Noise is quantified by a noise parameter $\eta \in [0, 1]$. Noise-free environments are modeled using $\eta = 0$; random dynamics is $\eta = 1$ [1]. Communication errors caused by noise are then modeled as follows [1]:

$$\tilde{\sigma}_j^i[t] = \begin{cases} \sigma_j[t] & \text{with probability } (1 - \eta/2) \\ -\sigma_j[t] & \text{with probability } \eta/2 \end{cases}. \quad (4)$$

3 Performance Study

3.1 Contributions and Methodology

We extend some parts of the results reported in [1] by using a larger set of system parameters. As network topologies, we apply the commonly used Watts-Strogatz and Newman-Watts-Strogatz models [7] on a one-dimensional grid and vary the number of neighbors from 2 to 5. The noise parameter $\eta/2$ is varied from 0 to 1. Moreover, we consider $k = 2 \dots 5$ neighbors in SMR. Every set of conditions is run 1 000 times with both synchronous and asynchronous update functions. The simulation engine is built with the programming language Python and uses the NetworkX software package (`networkx.lanl.gov`).

As performance metrics we measure the convergence rate (fraction of converged networks out of 1 000 simulations) and convergence speed (number of steps until network converges to a single state normalized by the overall number of steps within one simulation).

There appear “noise artifacts” at certain noise levels. These artifacts are ignored in [1], i.e., a network is considered as converged even if not all nodes have the same value. We only consider the consensus to be converged if really all nodes have the same state. Therefore our simulations assume somewhat more difficult conditions.

3.2 Simulation Results

For illustration, Figs. 1 and 2 show the evolution of the nodes’ states over time for given starting conditions. These plots basically resemble previous results [1].




Fig. 1. SMR rule in a network of $n = 100$ nodes with Watts-Strogatz connectivity and zero rewiring probability. Both networks do not convergence in the given time. a) Clustering with stable domain boundaries. b) Clustering with unstable domain boundaries.

Figs. 3 and 4 show the convergence rate and speed with different parameters and network models. The curves can be interpreted as follows:

- Using GKL, low and medium noise levels lead to a convergence rate of about 50%, while no and high noise tend to prevent consensus (Fig. 3a).




Fig. 2. GKL rule under noisy conditions in a network of $n = 100$ nodes with Watts-Strogatz connectivity and zero rewiring probability. a) No convergence. b) Convergence.

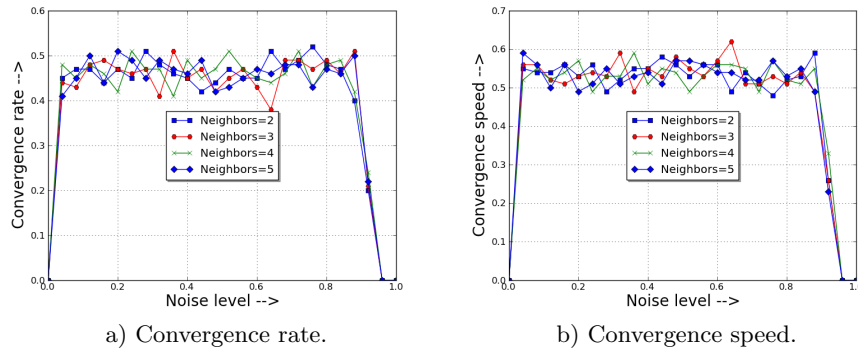


Fig. 3. GKL rule in a network of $n = 100$ nodes with Watts-Strogatz connectivity and rewiring probability 0.5 for topologies with different number of neighbors.




Fig. 4. SMR rule in a network of $n = 100$ nodes with rewiring probability 0.9.

- With Newman-Watts-Strogatz networks (Fig. 4b), there is an extra peak in convergence rate for high noise levels using SMR with synchronous update.¹
- Both GKL and SMR have weaker performance than in [1], which is due to the fact that we did not mitigate “noise artifacts” in measuring convergence.

In future work, the simulation engine will be used for robustness analysis with respect to other types of faults and in two-dimensional topologies.

Acknowledgments

This work was performed within the Erasmus Mundus Joint Doctorate in “Interactive and Cognitive Environments,” which is funded by the EACEA Agency of the EC under EMJD ICE FPA n 2010-0012. The work of A. Gogolev is supported by Lakeside Labs, Klagenfurt, with funding from the ERDF, KWF, and the state of Austria under grant 20214/21530/32606.

References

1. Moreira, A. A., Mathur, A., Diermeier, D., Amaral, L.: Efficient system-wide coordination in noisy environments. In: Proc. National Academy of Sciences of the USA, vol. 101, no. 33, pp. 12085–12090, (2004)
2. Gacs, P., Kurdyumov, G. L., Levin, L. A.: One dimensional uniform arrays that wash out finite islands (in Russian). *Problemy Peredachi Informatsii*, vol. 14, pp. 92–98, (1978)
3. Boyd, R., Peter J. R.: Why does culture increase human adaptability? *Ethology and Sociobiology*, vol. 16, pp. 125–143, (1995)
4. Heyes, C. M., Galef, B. G. Jr., eds.: *Social Learning in Animals: The Roots of Culture*. San Diego: Academic Press, (1996)
5. Watts, D. J.: *Small Worlds: The Dynamics of Networks Between Order and Randomness*, Princeton University Press, (1999)
6. Wolfram, S.: *A New Kind of Science*, Wolfram Media, Champaign, IL (2002)
7. Newman, M.: The structure and function of complex networks. *SIAM Review*, vol. 45, no. 2, pp. 167–256, (2003)

¹ As the Newman-Watts-Strogatz model represents social networks and excessive noise might correspond to a crowd-like message delivery, these extra peak could maybe be interpreted as potential ability of achieving consensus in social group with faulty message delivery.