

Influence of Shared Memory and Network Topology in the Consensus Dynamics of a Naming Game

Thaís G. Uzun, Reginaldo J. Da Silva-Filho, Matthias R. Brust, Carlos H. C. Ribeiro

Postgraduate Program in Computing and Electronical Engineering – Technological Institute of Aeronautics (ITA) – São José dos Campos, SP – Brazil

thaisgobet@gmail.com, reginaldo@aluno.ita.br,
matthias.brust@gmail.com, carlos@ita.br.

Abstract. *In the Naming Game, individuals or agents exchange pairwise local information in order to communicate about objects in a common environment. The goal of the game is to reach a consensus about naming these objects. In this paper we extend the classical naming game with a globally shared memory accessible by all agents. Although the extended naming game is non-deterministic in its word selection, we show that consensus towards a common vocabulary is reached in diverse network topologies. More importantly, we show the qualitative and quantitative influence of the external source of information, i.e. the shared memory, on the consensus dynamics.*

1. Introduction

The natural emergence of a common language among individuals remains a phenomenon yet to be explained. However, a deeper understanding of the evolutionary processes of language formation is indispensable for developing autonomous multi-agent systems where each agent can potentially have different origins and where no knowledge about the language used in an open-ended environment is provided. To put in a question: How can these agents build a common language through local agreements and reach a consensus about the meaning of their vocabulary on a global scale?

A promising model for a deeper understanding of the common language phenomenon is the *naming game* [Steels 1998]. It describes a model in which individuals can reach a consensus on how to name different objects. All individuals (or agents) exist in the same environment, sense the same set of objects, and are able to invent or create words for these objects. An interaction between two agents is a word transmission from one agent (the *speaker*) to a second agent (the *hearer*), with a resulting outcome (usually success or failure). The goal of the game is to reach, after a number of interactions, an agreement among the agents about the object-word association used for a single object [De Vylder and Tuyls 2006]. Interestingly, a self-organized vocabulary or even a common language with syntactic and semantic levels can be built using such simple local communication process [Steels 1996]. The naming game is therefore a microscopic model for the interaction dynamics among autonomous agents that communicate without any centralized control [Baronchelli et al. 2008]. Such distributed model could be used to understand how large populations reach an agreement with respect to the usage of a

certain word, how new language constructs are established, how rumors and opinions spread, how words propagate in social networks, and even to provide a basis to an emergent communication system where symbol grounding and intentionality are natural outcomes from a particular system dynamics [Steels *et al.* 2007].

Besides its application in modeling the language formation process for individuals, agents or robots (in particular in the field of Artificial Intelligence), the naming game is of relevance to understand the consensus dynamics of collaborative tagging systems of web sites like *delicious* and *flickr* [Marlow *et al.* 2006] that have become increasingly popular in recent years. The users of such sites can attach keywords or tags to provide information (e.g., favorite sites on the Internet). In a recent study [Golder and Huberman 2005], it is shown how collaborative tagging can lead to both regularities regarding users' activities, tag frequency and keyword usage, and stabilities concerning relative proportions between tags for a given URL and strings that define the location of programs or files in the Internet. Although it is potentially possible to have a constantly increasing number of tags, these findings indicate convergence to a name descriptor (the collection of tags) and concept (the contents in the location itself).

Different variants of naming games played by humans can help to overcome one of the most challenging problems for search engines: Image labeling. The ESP game [Ahn and Dabbish 2004] aims to use humans' perceptual abilities in order to create valuable output in the process of image labeling. Two players are shown the same image but they are not able to communicate. They are then asked to describe the image with labels under a given time constraint (e.g. Google Image Labeler uses 2 minutes). As soon as they use a common label, it is saved in the database to index the image, the players earn points accordingly, and the next image is shown. The objective is to get as many points as possible. While the ESP game is initially designed for a two-player game, in a broader context, the label consensus dynamics of the naming game can be directly used to improve the description accuracy of the images.

Research on the naming game uses mainly the introduced communication model above and focuses on showing its convergence empirically [Baronchelli *et al.* 2006, Steels 1996, Steels 1998]. Convergence of a deterministic naming game, however, has been mathematically proven [De Vylder and Tuyls 2006]. One common characteristic of these models is that their dynamics are influenced only by the local memories of the agents involved. There is no common access memory, implying that the dynamics of these models is completely uncoupled from any influence of an environment external to the one where the negotiation occurs. The consensus, when reached, is a consensus which belongs to a specific population, and makes sense only in that context. From a sociological point of view, such an arrangement can be plainly artificial, or at least very difficult to establish [Carrington *et al.* 2007].

The variation of the naming game model introduced in this paper differs from these approaches in that it enables agents to access a shared (global) memory with a given probability p (see Fig. 1). The reason for introducing a shared memory originates from the fact that the real world consists of central access points like books, media, and conferences where individuals build a common vocabulary even without a single direct interaction. Additionally, often an individual tends to search for an external reference before even emitting an opinion about a given subject. The shared memory extension might thus be important for modeling and understanding e.g. the influence of the press and media on the consensus of the group of individuals.

Since classical naming games that allow only local negotiations tend to converge [De Vylder and Tuyls 2006], it appears reasonable that an extended version using a shared memory should behave similarly. Although one of our contributions is to show that the extended naming game in fact converges, the focus is on the role and impact of a shared memory on the convergence behavior itself. Knowledge about this influence enables the possibility to control the convergence and, thus, to trigger the outcomes. Another contribution of this paper shows that against common sense expectation, the shared memory is not solely responsible for triggering the consensus word, thus giving importance backing to the importance of local interactions.

Moreover, we also consider that systems as diverse as the World Wide Web are best described as networks with complex topologies [Barabási and Albert 1999]. In fact, a common property of many large networks is that the vertex connectivity follows a scale-free power law distribution. A model based in preferential attachment to nodes that has this property is the Barabási-Albert model (BA) [Barabási and Albert 1999], and the influence of this topology on the consensus dynamics is also analyzed in this paper.

The paper is organized as follows. A detailed description of the shared memory-based model is given in Section II. The model has been implemented as a proof-of-concept prototype. Empirical results are shown and discussed in Section III. Section IV presents a discussion about the influence of the network topology in the dynamics behavior, and also presents empirical results about the model presented in this paper, when applied to BA networks. Finally, Section V presents the main conclusions of the paper.

2. A Non-Deterministic Naming Game Model

This section describes formally the model proposed in Section I. A population of N agents is considered. Each agent has access to a local memory, which can contain potentially any number of words about a given subject. A word can be a composition of alphabetic elements, but also any other kind of unique identifier. Furthermore, all agents have reading access to the common external (shared) memory. This shared memory contains — prior to the beginning of the game — C distinct words ($C \geq 1$). The objective of the game is to reach a steady state (consensus), i.e. a state in which all agents have the same word in their local memories.

At $t=0$, all agents have empty local memories. At each successive time step ($t=1,2,3,\dots$) two agents are randomly selected, one playing the role of the speaker and the other as the hearer. The negotiation dynamics is as follows:

1. The speaker randomly selects one of the words in its own local memory. If the local memory is empty, two actions are possible: a) with probability λ the speaker chooses a word from the shared memory whereby the selected word is added to the local memory of the speaker, or b) with probability $1-\lambda$ a new word is created locally and selected.
2. The speaker transmits the selected word to the hearer.
3. If the hearer does not have the transmitted word in its local memory, the interaction is considered a failure and the hearer adds the transmitted word to the local memory. If the hearer has the transmitted word in the local memory, two actions are possible. In both cases, the negotiation is considered a success: a) with probability λ the agents involved consult the shared memory. If the transmitted word exists in the shared memory, the

speaker and hearer remove all other words from their memories, or b) with probability $1-\lambda$ both agents remove all words, besides the transmitted one, from their local memories.

The model has three inputs: the number of agents N , the probability λ , and the number of words C in the shared memory. The probability λ represents the tendency for the agents to check the shared memory. When $\lambda=0$, the model is reduced to the standard naming game described e.g. in [Baronchelli et al. 2006], i.e. without any external influences which are represented by the shared memory. When $\lambda=1$, the game can be interpreted as a controlled version of the naming game having C possible words.

3. Simulation and Results for MF Topologies

The extended naming game communication model has been implemented as a proof-of-concept prototype. Empirical results obtained by simulation are evaluated and discussed in this section.

We assume that the agents are in a fully connected network, where each agent can communicate with all others. This topology is referred in Statistical Mechanics as a Mean Field topology (MF), and is also assumed in the communication models described in [Baronchelli et al. 2006, De Vylder and Tuyls 2006, Steels 1996, Steels 1998]. The number of agents N is set to 100 in all simulations. The values of λ vary from 0.0 to 1.0, and we tested values of C as 1, 5, 10, 50, 100 and 500. For each combination (λ, C) the game was executed 1000 times. The results shown are averages over these runs.




Figure 1. Illustration of the extended naming game using a shared memory that agents are able to access with probability λ .

There are three default measurements for the naming game, see e.g. [Baronchelli et al. 2007]. The first one is the variation of the total number of words as a function of time $N_w(t)$. For a given time step t , the value of $N_w(t)$ is the sum of the number of words in the local memories of all agents. Second, we define $N_d(t)$ as a function that gives the number of different words at time t , i.e. it is the number of elements of the set containing all the words in the model at time t . Third, we define the success rate $S(t)$ as follows: In a given interaction between two agents, the value 1 is assigned if the

interaction is a success and 0 if it is a failure. It is important to note that for a given execution the success rate $S(t)$ can only have values either 0 or 1.

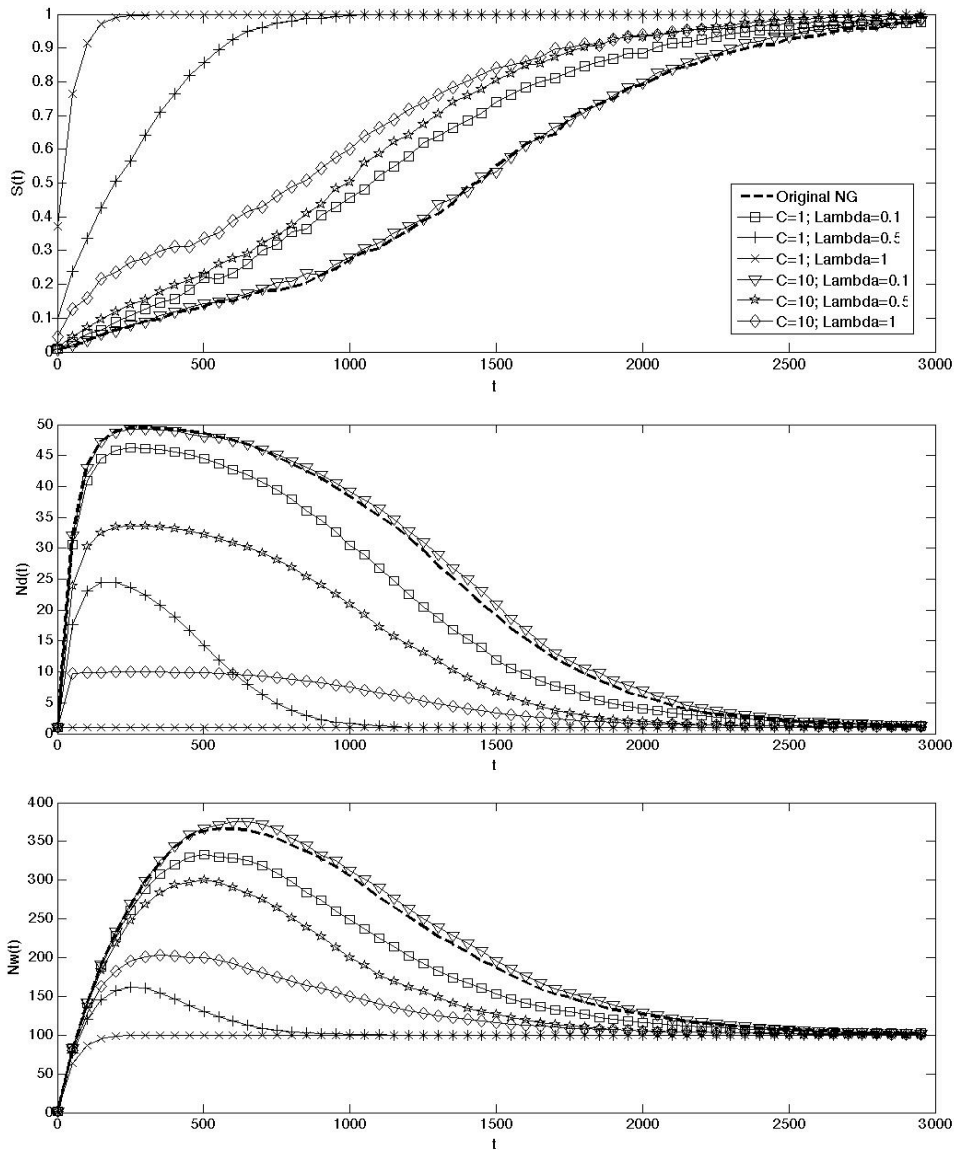


Figure 2. Curves for N_w , N_d and S as a function of time for $N=100$ agents playing the game described by the introduced model for a MF topology.

Figure 2 is an overview of the behavior of the basic properties of the system using the introduced shared memory in a MF topology. It shows that the system dynamics is influenced by both λ and C . The dark blue curves show the results for the standard naming game [Baronchelli et al. 2006], which occurs when $\lambda = 0$. It is also possible to see that the system clearly undergoes a disorder/order transition. At the beginning, the total number of words in the system $N_w(t)$ grows smoothly, indicating that unsuccessful interactions occur, a fact that can be confirmed by the low value of $S(t)$. On the other hand, the number of different words $N_d(t)$, grows significantly, quickly reaching its maximum value. This means that new words are introduced. Still at the beginning, $N_d(t)$ begins to decrease, although somewhat moderately, while the value of $N_w(t)$ is still increasing. This means that although successful interactions start to

occur, failures are still predominant. After $N_d(t)$ reaches its maximum, no new words are further created. Instead, the initially created words spread all over the network. The difference compared to the initial phase is in the fact that the rate for creating new words is steadily decreasing. The value of $S(t)$ grows moderately at first, but when the existing words are propagated to the majority of the agents, some become very popular, and the success rate starts to grow at a faster pace. With the more frequent occurrence of successful interactions, both the total number of words and the number of different words decrease, eventually leading to a consensus state, where $N_d=1$ and $N_w=N$.

An important issue is to analyze how the input parameters λ and C influence the behavior of the system. In other words, it is important to verify how the shared memory affects the game dynamics. The most clearly affected property is the maximum value of $N_d(t)$, $\max(N_d)$, which is the maximum number of distinct words in the system. For a fixed value of C , $\max(N_d)$ decreases for increasing values of λ . On the other hand, for a fixed value of λ , $\max(N_d)$ increases for increasing values of C . Figure 3 shows the behavior of $\max(N_d)$ with respect to λ and C .

Observe that checking of shared memory by the agents as described in Section II is done with probability λ . It can potentially happen in two situations: (a) an agent receives a transmitted word or (b) an agent is selected as speaker and does not have a word in its local memory. In the latter case, the agent can choose one of the words of the shared memory with probability λ or invents a new word with probability $1-\lambda$.

In the classical naming game with N agents [Baronchelli et al. 2006], $\max(N_d)$ is approximately $N/2$, meaning that on average half of the agents invent new words. This happens because the inventing agents were chosen as speakers while their local memories were empty. With the introduction of the shared memory, this behavior is expected as well, so that on average $N/2$ agents are chosen as speakers while their local memories are empty. Amongst these agents, $\lambda N/2$ choose a word from the shared memory for transmission, while $(1-\lambda)N/2$ will introduce (invent) new words, ideally distinct ones. Then, the average maximum number of distinct words expected in the system obeys $\max(N_d(N,\lambda,C)) \leq (1-\lambda)N/2 + NC_d(N,\lambda,C)$, where $NC_d(N,\lambda,C)$ represents the maximum possible number of words chosen by the $\lambda N/2$ agents amongst the C words of the shared memory, in other words $NC_d(N,\lambda,C) = C$ if $\lambda N/2 > C$ and $NC_d(N,\lambda,C) = \lambda N/2$, if $\lambda N/2 \leq C$.

Figure 4 shows the variation of the time in which the number of distinct words in the system reaches its maximum value $t_{\max(N_d)}$. For a fixed C , $t_{\max(N_d)}$ decreases for increasing values of λ . For a fixed value of λ , $t_{\max(N_d)}$ increases for increasing values of C . Figure 5 shows the behavior of the average convergence time t_{conv} for the game. We say that the system has converged when every agent has exactly one word, which is the same for all of them, that is, when $N_w=N$ and $N_d=1$. For $C=1$ (only one word in the shared memory), the convergence time always decreases when λ increases. For other values of C , the convergence time is maximum for some λp , increasing in the interval $[0, \lambda p)$ and decreasing in $(\lambda p, 1]$. In general, for a fixed λ , t_{conv} increases for increasing C . When $\lambda=0$, the convergence time obviously does not depend on C and its value is approximately 2,500, in fact the same registered in [Baronchelli et al. 2006] and, thus, indirectly validating the implementation.

The curves for the maximum number of words in the system, $\max(N_w)$ are shown in Figure 6. For a fixed value of λ , $\max(N_w)$ increases for increasing values of C .

When $C=1$, the value of $\max(N_w)$ always decreases when λ increases. For other values of C , $\max(N_w)$ also reaches its maximum value for some λp , increasing in the interval $[0, \lambda p)$ and decreasing in $(\lambda p, 1]$. Figure 7 shows the behavior of the property $t_{\max(N_w)}$, the time in which the total number of words in the system $N_w(t)$ reaches its maximum value. For a fixed value of λ , $t_{\max(N_w)}$ always increases for increasing values of C .

For all simulations executed, convergence was observed to a state in which all the agents have the same word, i.e. a steady state. Interestingly, the resulting *consensus word* is not always amongst the C words in the shared memory. To analyze this result we define the parameter P_{shared} as the quotient between the number of executions in which the consensus word is also in the shared memory and the total number of executions. It means the probability that a system with inputs N , λ and C converges to a word in the shared memory. The behavior of P_{shared} is shown in Figure 8. Remarkably, the shared memory only contains the consensus word in all executions when $\lambda > 0.5$. For $\lambda < 0.5$ the ratio depends on the number of words C in the shared memory whereby more words mean a lower ratio.




Figure 3. Variation of the maximum number of distinct words in the system $\max(N_d)$ with respect to λ and C for the MF topology.




Figure 4. Time when $N_d(t)$ reaches its maximum value $t_{\max(N_d)}$ as a function of λ for the MF topology.




Figure 5. Average convergence times for the proposed model for the MF topology.




Figure 6. Maximum number of words in the system $\max(N_w)$ for the MF topology.




Figure 7. Time in which the total number of words in the system is maximum ($t_{\max(N_w)}$) for the MF topology.




Figure 8. Behavior of P_{shared} as function of λ and C .






Figure 9. Curves for $N_w(t)$, $N_d(t)$ and S as function of time for $N=100$ agents playing the game described by the introduced model for a BA topology.




Figure 10. Behavior of the maximum number of distinct words in the system $\max(N_d)$ with respect to λ and C for the BA topology.

In order to explain this phenomenon, we consider that invention of words only occur in the very beginning of the game. When λ increases, the number of agents that choose words from the shared memory instead of inventing new ones increases at the same time. Thus, some of the words which were initially exclusively in the shared memory become popular from the very beginning of the game. This explains the fact that the consensus word has a relatively high probability of belonging to the shared memory.

4. Simulation and Results for BA topologies

In the previous section, we assumed that the agents were in a fully connected (MF) network. However, several works report on how the underlying topology influences the consensus behavior [Baronchelli et al. 2006, Baronchelli et al. 2007, Brust et al. 2008, McIntyre and Steels 1999]. In this section, we investigate how the network topology affects the properties of the extended naming game.

Many topologies of large networks, from the WWW to citation patterns in science, display that, independently of its constituents, the probability $P(k)$ that a vertex in the network interacts with k other vertices decays as a power law that follows $P(k) \sim k^{-\gamma}$. This property is called scale-free [Barabási and Albert 1999]. We consider here the scale-free BA model from [Barabási and Albert 1999], which has become one of the most used models for complex heterogeneous networks. A BA topology is a very simple construct. Starting from a small set of m interconnected nodes, new nodes are introduced one by one. Each new node selects m older nodes for connection according to the preferential attachment rule, i.e., the probability of connecting to a node is proportional to its degree. When a predefined network size is reached, this procedure stops. It can be shown that the obtained network follows a power law distribution $P(k) \sim k^{-\gamma}$, with $\gamma = 3$ [Barabási and Albert 1999]. We thus consider here BA networks with $N=100$ agents and average degree $\langle k \rangle = 15$. The values of λ and C are the same as in Section III and, for each combination (λ, C) , the game was executed 500 times. The results were averaged over these runs. Figure 9 shows the behavior of the parameters $N_d(t)$, $N_w(t)$ and $S(t)$. The system dynamics behaves as in the MF case, influenced by the values of both λ and C .

The most clearly influenced parameter is the maximum value of $N_d(t)$, $\max(N_d)$. For a fixed value of C , $\max(N_d)$ decreases for increasing values of λ . Also, for a fixed value of λ , $\max(N_d)$ increases for increasing values of C , as in the MF case. A comparison between Figure 9 and Figure 2 shows that the maximum total number of words is smaller than in the MF case, while the number of different words remains almost the same. This is due to the fact that the network topology does not influence the creation of new words, but the network has an average connectivity smaller than the MF network. The nodes have access only to the words of their neighbors' inventories, whereas in the MF case all nodes have potential access to all words in the game. In this way, the average memory size of any agent is smaller than in the MF case, as it has more limited access to the existing words. The behavior of $\max(N_d)$ with respect to λ and C is shown in Figure 10. Figure 11 shows the behavior of the convergence time t_{conv} . Comparing Figures 11 and 5, one can notice that t_{conv} for BA networks is larger than in the MF case. We can explain this fact also by the smaller average connectivity of the BA networks, resulting in a more local spreading of words, thus taking longer to reach consensus. Figure 12 shows the curves for the maximum number of words in the system, $\max(N_w)$. We can notice, comparing once again with the same parameter in the MF case (Figure 6), that the BA case has lower memory use, due to the fact that each node has restricted access to the existing words in the network, as already mentioned in this section. The curves for the time in which $N_w(t)$ reaches its maximum, $t_{\max(N_w)}$ are shown in Figure 13. The behavior of the parameter P_{shared} , defined in the last section, is shown in Figure 14. In the BA case, words are propagated in a more local way than in the MF case. For BA networks, when $\lambda > 0.4$, the shared memory contains a consensus word with high probability, while in the MF this occurs for λ larger than 0.5. This difference happens because, in the BA case, the words are propagated in a more localized form than in the MF case.

When, in a BA network, an agent invents a random word, it can only be present in various network nodes if it is propagated step by step during various interactions. The word presence in different nodes increases the chances of it being the consensus word. On the other hand, when more than one agent selects a same word of the shared memory, few interactions are necessary (relatively to when the word in question was invented) for the word to be in various different network nodes. Thus, in the BA network, the difference between a probability of consensus on a word of the shared memory and the probability of consensus on a random invented word is larger than the same difference for a MF network. In a MF network, an invented word may be in various nodes of the network after few interactions, because every node is connected to all the others. Thus, there is less difference (relatively to the same pattern in a BA network) in the consensus probability of a random word and of a word in the shared memory.

As a final remark, we observe that, regarding the variation of the size of the shared memory C and probability λ , the model maintains qualitatively the same characteristics found in the previous section, with respect to the variation of $\max(N_d)$, t_{conv} , $\max(N_w)$, $t_{\max(N_w)}$ and P_{shared} .

5. Discussion and Conclusion

In this paper, we introduce a shared external memory into the original naming game model [Steels 1998] and analyzed the resulting consensus dynamics both for completely connected (MF) and scale-free (BA) topologies. The memory can be interpreted as the

role of a dictionary, a popular reference (book, encyclopedia, etc.), the press, or a search engine, in a social network of simple communicating agents.

Results show that if the agents follow the communication rules described in the extended naming game consensus is always reached, *i.e.*, the agents reach an agreement on the vocabulary about the objects in their environment. That happens without centralized control, as the agents only have reading access to the shared memory. We have empirically shown the degree of impact of the external shared memory on the consensus dynamics. It is noteworthy observing that, although the shared memory does not completely determine the consensus word (a possible indication of the importance of local interactions), it has enormous influence in defining it, even for low access probabilities. For further investigations we can consider the idea that some entity has total access to the shared memory (reading and writing). This entity could to a certain extent determine and manipulate the outcome of the game, bringing to the fore an interesting discussion on related media control phenomena for which the model could be understood as a first simplification. As the value of λ does not even have to be the same for all agents, it might be more realistic, as far as a study on social communications is concerned, to consider that each agent has its own value of λ , simulating different likelihoods of being influenced by an external source of information. The characteristics of the external shared memory (including but not limited to the value of C) are in a certain way determined by the entities that control the external memory. In other words, if these entities allow agents to access only a limited set of possible words, this will result in lowering C . On the other hand, with a low value of C , it is easier to predict the outcome of the game: the consensus word will very likely be amongst those words that are “interesting” to the entities that control the memory, even if the system has small λ . Further studies will consider socially-related concepts such as trust and reputation in the line initiated in [Brigatti 2007], and the inclusion of utility measures as considered in the multi-agent games literature.




Figure 11. Average convergence times for the proposed model for the BA topology.




Figure 12. Maximum number of words $\max(N_w)$ for the BA topology.




Figure 13. Average time in which the total number of words in the system is maximum $t_{\max(N_w)}$ for the BA topology.




Figure 14. P_{shared} versus λ and C for the proposed model in BA networks.

6. References

Ahn, L.V. and Dabbish, L. "Labeling images with a computer game," *Conference on Human Factors in Computing Systems*, 2004.

- Barabási, A. and Albert, R. "Emergence of scaling in Random Networks", *Science*, 286, pp. 509-512, 1999.
- Baronchelli, A., Barrat, A., Dall'Asta, L. and Loreto, V. "Non-equilibrium dynamics of language games on complex networks", *Physical Review E*, vol. 74, 2006.
- Baronchelli, A., Barrat, A., Dall'Asta, L. and Loreto, V. "The role of topology on the dynamics of the Naming Game," *The European Physical Journal Special Topics*, vol. 143, pp. 233-235, 2007.
- Baronchelli, A., Felici, M., Loreto, V., Caglioti, E. and Steels, L. "Sharp transition towards shared vocabularies in multi-agent systems," *Journal of Statistical Mechanics: Theory and Experiment*, 2006.
- Baronchelli, A., Loreto, V. and Steels, L. "In-depth analysis of the Naming Game dynamics: the homogeneous mixing case," vol. 19, 2008.
- Brigatti, E. "Consequence of reputation in an open-ended Naming Game," *Physical Review E*, 78(4):046108, 2008.
- Brust, M.R., Ribeiro, C.H. and Mesit, J. "Avoiding Greediness in Cooperative Peer-to-Peer Networks," *Collaborative Computing: Networking, Applications and Worksharing*, E.B. Joshi and J.B. D., Orlando, USA: Springer Berlin Heidelberg, pp. 370-378, 2008.
- Carrington, P., Scott, J. and Wasserman, S. *Models and Methods in Social Network Analysis*, Cambridge Univ. Press, 2007.
- De Vylder, B. and Tuyls, K. "How to reach linguistic consensus: a proof of convergence for the naming game," *Journal of theoretical biology*, vol. 242, pp. 818-31, 2006.
- Golder, S.A. and Huberman, A. "The Structure of Collaborative Tagging Systems," Social Computing Lab, 2005.
- Marlow, C., Naaman, M., Boyd, D., and Davis, M. "HT06, tagging paper, taxonomy, Flickr, academic article, to read," *Conference on Hypertext and Hypermedia*, 2006.
- McIntyre, A., and Steels, L. "Spatially Distributed Naming Games," *Advances in Complex Systems*, vol. 1, pp. 301-323, 1999.
- Steels, L. "Self-Organizing Vocabularies," *Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, C.G. Langton and K. Shimohara, Nara, Japan, 1996.
- Steels, L. "The Origins of Ontologies and Communication Conventions in Multi-Agent Systems," *Autonomous Agents and Multi-Agent Systems*, vol. 1, pp. 169 - 194, 1998.
- Steels, L., Loetzsch, M. and Spranger, M., "Semiotic Dynamics Solves the Symbol Grounding Problem", Available from Nature Precedings <<http://hdl.handle.net/10101/npre.2007.1234.1>>, 2007.