Better Exploration-Exploitation Pace, Better Swarm: Examining the Social Interactions

Marcos Oliveira*, Diego Pinheiro*, Mariana Macedo†, Carmelo Bastos-Filho† and Ronaldo Menezes*
*BioComplex Laboratory, School of Computing
Florida Institute of Technology, Melbourne, Florida, USA
Email: moliveira@biocomplexlab.org, dsilva@biocomplexlab.org, rmenezes@cs.fit.edu
†Department of Computer Engineering
University of Pernambuco, Recife, Brazil
Email: mgmm@ecomp.poli.br, carmelofilho@poli.br

***Abstract*—Swarm-based models have successfully solved real-world problems in the past two decades and yet they continue to exhibit a major shortcoming of premature convergence. Previous research suggests that an appropriate exploitation-exploration balance can prevent premature convergence and different approaches have been proposed to control this balance. Still, despite several references demonstrating the interplay between social interactions and swarm behavior, the majority of works lack a network-based assessment of the level of balance in a swarm. We propose that pacing social interactions is the key to balance exploration-exploitation. Here we examine the impact of the exploration-exploitation balance on the swarm performance by controlling the pace at which the swarm goes from exploration to exploitation. Our results revealed that this pace influences the swarm dynamics and that different problems demand distinct paces. Swarm-based models that are capable of adapting their exploration-exploitation pace have the potential to overcome premature convergence.**

***Index Terms*—social interactions, self-organization, complex systems, network science, swarm intelligence, particle swarm optimization**

I. Introduction

The idea of "intelligence" as in Swarm intelligence emerges from the social interactions between the agents in the swarm. We still fail, however, to understand the impact of these interactions on the swarm dynamics. The ability of the swarm to coordinate and to adapt comes from the rules that define the interactions among individuals as well as between individuals and the environment [1]–[3]. A minimal amount of interactions is necessary for a system to exhibit coordination; nonetheless, adaptability vanishes from a system that exhibits too many interactions [4]. In nature, for instance, ants interact among themselves to avoid random walking while being able to adapt to the environment. In computational models, however, we lack an understanding of such optimal balance which could help us to build better intelligent tools.

Inspired by the flocking of birds, Particle Swarm Optimizers (PSO) consist of a family of swarm-based optimization methods that rely on the interactions among individuals searching for better solutions in a hyper-dimensional search space [5]. These individuals are represented by simple reactive agents (i.e., particles) that move through the search space communicating and acquiring positional information from other particles. Each particle moves in the problem space with inertial energy influenced by two terms: the personal experience and the social experience [2]. Despite different approaches to weigh these terms, the swarm behavior basically results from particles updating themselves based on other particles. Though the technique has been widely used, it still presents a premature convergence problem in which particles reach suboptimal solutions due to early equilibrium states [2]. Such condition occurs when the swarm stagnates on a local optimum usually because of the so-called exploration-exploitation imbalance [2], [3].

The exploration-exploitation balance is a concept used in different fields, ranging from human behavior to computer science, which refers to the strategy of a search: the focus on existing knowledge characterizes exploitation while seeking for new knowledge defines exploration [2], [6]. From the swarm perspective, exploration regards to the swarm's ability to explore the search space broadly, whereas exploitation bounds the search on a particular area—an essential trade-off for a good optimization technique [3]. To understand this balance, researchers usually analyze the particularities of the techniques; in the case of the PSO, most analyses focus on the properties of particles, such as position and velocity—they neglect the social interactions [7], [8].

By overlooking such an essential aspect of swarm intelligence, one misses an opportunity to unveil the very process that leads to a exploration-exploitation imbalance, which starts from the interactions among individuals. Nevertheless, some works have already attempted to understand the swarm behavior by examining the social interactions among particles [9]–[19]. Kennedy and Mendes showed that the infrastructure through which particles interact influences the swarm performance [9]. That is, the limits of social interactions have an impact on the swarm dynamics [9], [12]. Note that such analysis only takes into account the boundaries and neglect the actual social interactions. Indeed, Oliveira et al. were

Marcos Oliveira and Diego Pinheiro thank the Science Without Borders program (CAPES, Brazil) for financial support under grants 1032/13-5 and 0624/14-4. Mariana Macedo thanks to the University of Pernambuco (PFA-UPE) and CAPES (Coordination for the Improvement of Higher-Education Personnel), Brazil, for the partial financial support for this paper. Carmelo J. A. Bastos Filho, IEEE SM, University of Pernambuco (UPE), Recife, Brazil.

the first to examine the social interactions to assess swarm behavior [16]. For that, they proposed a representation of the swarm with a network that nodes are the individuals (i.e., the particles) which are connected if they exchanged positional information in a given iteration; and extended the concept to capture such dynamics over the iterations [17]–[19].

Here we examine the social interactions within the swarm while controlling the pace at which the swarm goes from exploration to exploitation. We want to understand the impact of the exploration-exploitation pace on the swarm performance. To control this pace, we use a dynamic topology as the infrastructure of communication which presents a parameter that allows us to tune this pace. We found that the exploration-exploitation pace influences the swarm behavior and that distinct problems demand distinct optimal paces. We advocate that regulating social interactions is the key to achieve an adequate exploration-exploitation balance, and as such we need more studies on the interactions among individuals within the swarm.

The remainder of this paper is organized as follows. In Section II, we provide an overview on examining social interactions within the swarm. We present the experimental setup and the obtained results, respectively, in Section III and Section IV. Finally, in Section V, we present a brief discussion about our results.

II. Background

We briefly present the Particle Swarm Optimization algorithm in Subsection II-A and discuss how the communication topology influences the swarm behavior. Then, in Subsection II-B, we describe the procedures to examine the social interactions in the swarm.

A. Particle Swarm Optimization

Particle Swarm Optimization (PSO) is an optimization method inspired by the social behavior in flocks of birds [5]. The technique consists of a swarm of particles in which each particle i is defined by four vectors $\vec{x_i}(t)$, $\vec{p_i}(t)$, $\vec{n_i}(t)$, and $\vec{v_i}(t)$, that play different roles in the meta-heuristic. The current position $\vec{x_i}(t)$ of the particle i is a d-dimensional vector in the search space which represents a candidate solution to the problem. Over the iterations, each particle i stores its best position found in $\vec{p_i}(t)$. The best position found by the neighbors of the particle i is stored in $\vec{n_i}(t)$ and is updated based on the information received via communication between particles. Finally, the velocity $\vec{v_i}(t)$ of a particle i stores the inertia and promotes the movement of the particles in the system.

At each iteration, each particle i updates its current velocity $\vec{v_i}(t)$ and then alters its position based on the current velocity, as following:

$$\vec{x_i}(t+1) = \vec{x_i}(t) + \vec{v_i}(t+1) \ . \tag{1}$$

In the original version of PSO, the particles update their velocity in such way that the so-called explosion state may occur—the velocities increase indefinitely [20]. To overcome this undesired state, Clerc and Kennedy developed the constriction factor χ defined as

$$\chi = \frac{2}{\left|2 - \varphi - \sqrt{\varphi^2 - 4\varphi}\right|} \quad \text{with} \quad \varphi = c_1 + c_2. \tag{2}$$

The factor χ adjusts the influence of the previous particle velocities during the optimization process, so the final equation for the particles' velocities is given by:

$$\begin{aligned}\vec{v_i}(t+1) = \chi \cdot \Big\{ \vec{v_i}(t) \ + \ \vec{r}_1 c_1 \cdot \Big[\vec{p_i}(t) - \vec{x_i}(t)\Big] \\ + \ \vec{r}_2 c_2 \cdot \Big[\vec{n_i}(t) - \vec{x_i}(t)\Big]\Big\} \ , \end{aligned} \tag{3}$$

where $\vec{r}_1$ and $\vec{r}_2$ are random vectors generated from a uniform probability density function within the interval [0,1]. The constants c_1 and c_2 are the cognitive and social acceleration factors, respectively; they are non-negative constants that weigh the cognitive and social components. Thus, each particle i acts depending on the best position $\vec{p_i}(t)$ and the best position $\vec{n_i}(t)$ found by its neighbors.

B. Communication Topology

The swarm topology defines the capability of information exchange among particles. This structure establishes the set of particles from which each particle i can update its $\vec{n_i}(t)$ component at each iteration. In the original PSO paper, the global topology was the first one introduced; with this static structure, the particles are capable of exchanging information with all the particles in the swarm—creating the ability to achieve fast convergence in unimodal problems but causing premature convergence in multimodal problems [5], [21]. In the case of local topologies, the infrastructure bounds the communication by allowing only specific particles to communicate [21]. Particularly, one of the most used local topologies is the ring topology, in which each particle only shares information with exact two other particles [21]. With a ring topology, the communication is more distributed, so it is argued that the swarm tends to explore different regions which leads to better performance in multimodal problems—though global and ring topologies exhibit much variability [22].

These topologies, however, are usually more suited to specific problems, so efforts were made to develop more balanced topologies. Oliveira et al. proposed an adaptive topology in which stagnated particles try to find better particles to be connected [23]. When the swarm perceives that the search is getting stagnated, the swarm modifies the flow of information. For this, each particle contains a new attribute, called p_k-$failure$, that stores the fitness improvement of the particle k. Every time the particle k does not improve its solution, p_k-$failure$ increments, otherwise, p_k-$failure$ sets to zero. If a particle k presents p_k-$failure$ higher than a certain threshold, p_k-$failure^T$, this particle modifies its neighborhood by connecting to a new particle. The particles choose a new neighbor probabilistically using a roulette wheel based on the rank of the particles' fitness (i.e., the best particles have higher chances to be selected as new connections). The threshold of p_k–$failure$ regulates how fast particles attempt to change their neighborhood; this parameter allows us to control the swarm settings. More information about this dynamic topology can be found in [23].

C. Examining the Social Interactions

To examine the social interactions within the swarm, Oliveira et al. defined that a social interaction in the PSO happens when a particle i updates its position based on the position of a particle j (the best neighbor of particle i is the particle j) at a given iteration t. They defined the interaction graph in which the weight of an edge (i, j) in the graph is equal to the number of times the particle i was the best neighbor of the particle j or vice-versa [17]. Additionally, they used a time window to control the recency of the analysis, so the interaction graph at iteration t with window t_w is defined as follows:

$$\mathbf{I}_{ij}^{\mathbf{t_w}}(t) = \sum_{t'=t-t_w+1}^{t} \left[\delta_{i,n_j(t')} + \delta_{j,n_i(t')}\right], \quad \text{with} \quad t \geq t_w \geq 1, \tag{4}$$

where $\delta_{i,j}$ is Kronecker delta. The time window tunes the frequency-recency balance of the analysis. With high t_w, the graph is dominated by the most frequent interactions; while low t_w includes recent interactions.

The interaction graph allows us to measure the diversity in the social interactions. For this, Oliveira et al. proposed to measure how fast the interaction graph can be destroyed (i.e., how fast the number of components increases as we remove edges of the graph) [19]. If the graph is rapidly destroyed, the swarm lacks diversity; otherwise, different information flows are present within the swarm [19]. For a given time window, the area under the destruction curve A_{t_w} measures the diversity in the information flow for a given time window. The communication diversity CD is defined as the following:

$$\mathrm{CD}(t) = 1 - \frac{1}{|T||S|} \sum_{t'_w \in T} A_{t_w = t'_w}(t), \tag{5}$$

where $|S|$ is the number of particles in the swarm and T is a set of time windows. Thus, swarms exhibiting high CD (i.e., low values for A_{t_w}) have the ability to have diverse information flows, while low values for CD imply in swarms with only few information flows (i.e., high value for A_{t_w}). An ideal set T would be one taking into account all time windows (i.e., interactions from $t_w = 1$ until $t_w = t$). Still, this procedure can be computationally expensive given the vast number of possible time windows, and a more reasonable approach is to have a set of time window samples.

III. Experimental Settings

Here we want to assess the behavior of the swarm while controlling the exploration-exploitation pace, thus we simulated the PSO using the dynamic topology with different p_k-$failure^T$. We want to examine the relationship between p_k-$failure^T$ (which regulates the pace at which the swarm goes from exploration to exploitation) and the communication diversity CD (which tells us about the exploration-exploitation with respect to the social interactions). To that end, we performed our experiments setting p_k–$failure$ equal to 2, 25, 50, 75, and from 100 to 1000, with increments of 100; and also used global and ring topologies.

In order to assess the swarm, we need to validate their performance across benchmark functions designed with specific features that pose distinct challenges to these models such as multi-modality. In this sense, we employed the following CEC'2005 functions: Rosenbrock, Rastrigin, Griewank, Ackley, Schwefel, and Sphere [24]. In all experiments, the number of simulations is 30, the number of dimensions D is set to 1000, the maximum number of iterations is 6000, the swarm is composed by 100 particles that are updated according to Eq. 3 with $c_1 = 2.05$, $c_2 = 2.05$, which guarantee the convergence behavior of the algorithm [20]. The initial topology of the dynamic topology is defined as the ring topology.

IV. Results

Our results are organized into three parts. First, we show that different functions require distinct exploration-exploitation paces. Second, we demonstrate that different topologies promote distinct patterns of social interactions that are capable of regulating how the swarm adapts its pace. Finally, we present the different schemes of adaptation provided by the dynamic topology.

A. The Appropriate Exploration-Exploitation Pace

Swarm-based models are generally designed assuming that functions require the same exploration-exploitation pace and thus tend to under-perform on specific problems due to either late or premature convergence. Our results reveal that different functions require distinct exploration-exploitation paces and best solutions are found when the swarm adjusts its pace accordingly (see Fig. 1). When comparing the performance of PSO in different functions, ring and global topologies alternate the functions on which they outperform each other—confirming results in [22]. In this setting, since ring is the least connected topology and global is the highest connected topology, they regulate PSO at the lowest and highest paces, respectively. Given the time constraints, the global topology regulates PSO at exploitation pace and outperforms the ring topology on functions requiring faster paces such as Rosenbrock, Schwefel, and Sphere. In the case of Rastrigin, Griewank, and Ackley, the ring topology regulates PSO at exploration pace and outperforms the global topology

The dynamic topology regulates intermediate paces of exploration-exploitation and allows the swarm to find better solutions than the ring and global topologies in each function. This regulation is controlled by the parameter p_k-$failure^T$ which determines the extent in which connections of an initial ring topology are rewired. The higher this parameter (e.g., 900 and 1000), the less frequent connections are rewired and the

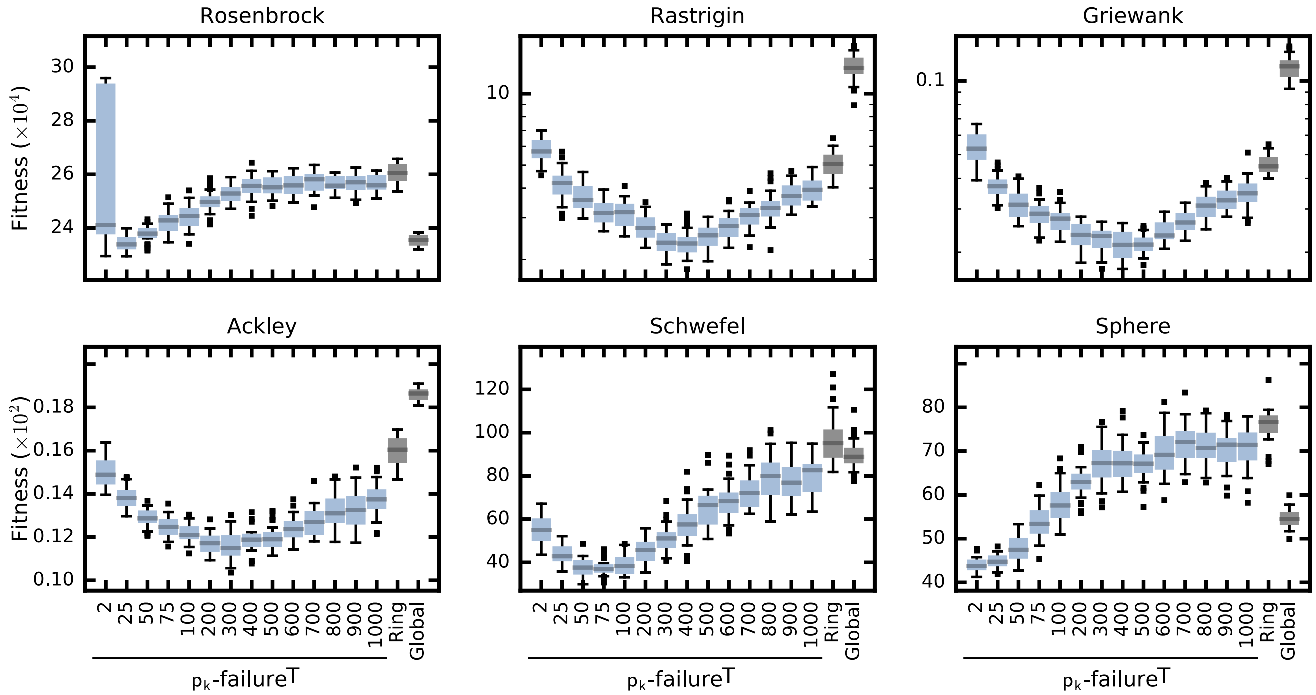

Fig. 1. Different functions require a distinct exploration-exploitation pace at which better solutions can be found. Each plot shows the fitness of final solutions found by PSO with exploration-exploitation paces regulated by the dynamic topology (first 14 leftmost box-plots) as well as ring and global topologies (two rightmost box-plots).

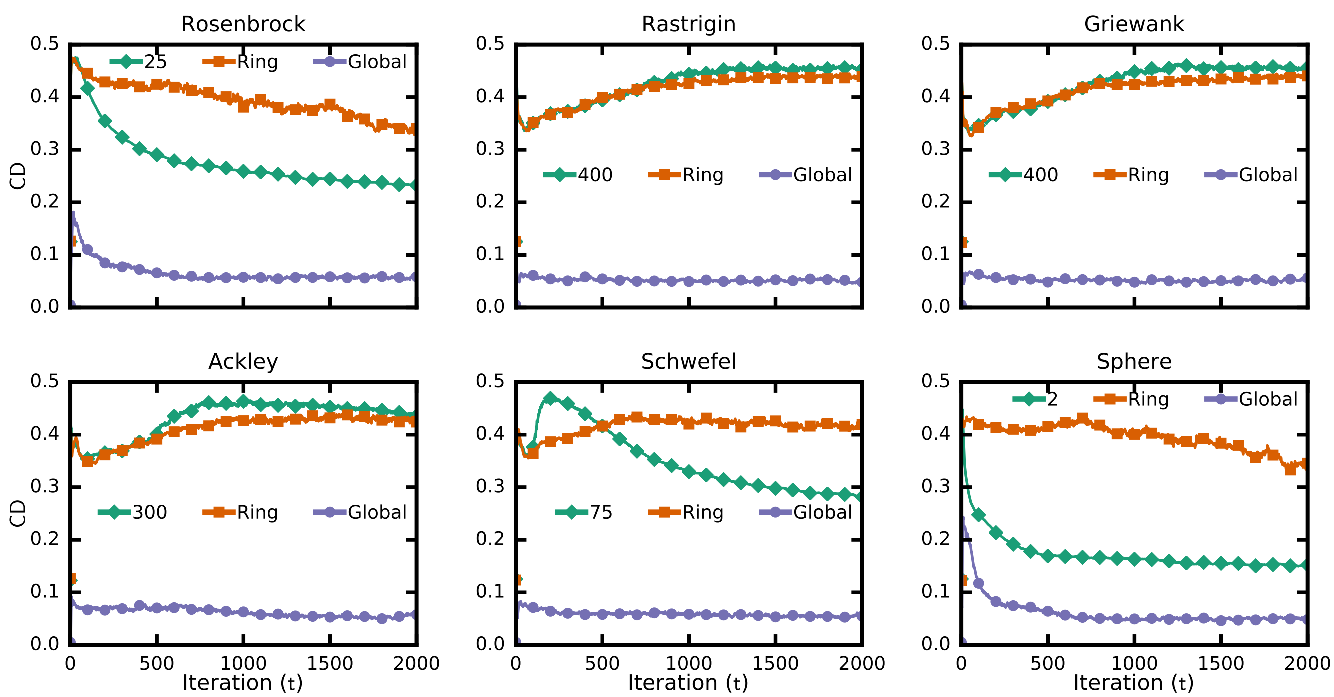

Fig. 2. The interplay between exploration-exploitation pace and communication diversity. Different topologies promote distinct patterns of social interactions that regulate how the swarm adapts its pace towards one required by each function. Each plot shows the communication diversity CD of PSO along the iterations for the ring and global topologies as well as for the dynamic topology with $p_k{-}failure$ threshold equals to the value for which the best fitness was found (see also Fig.1).

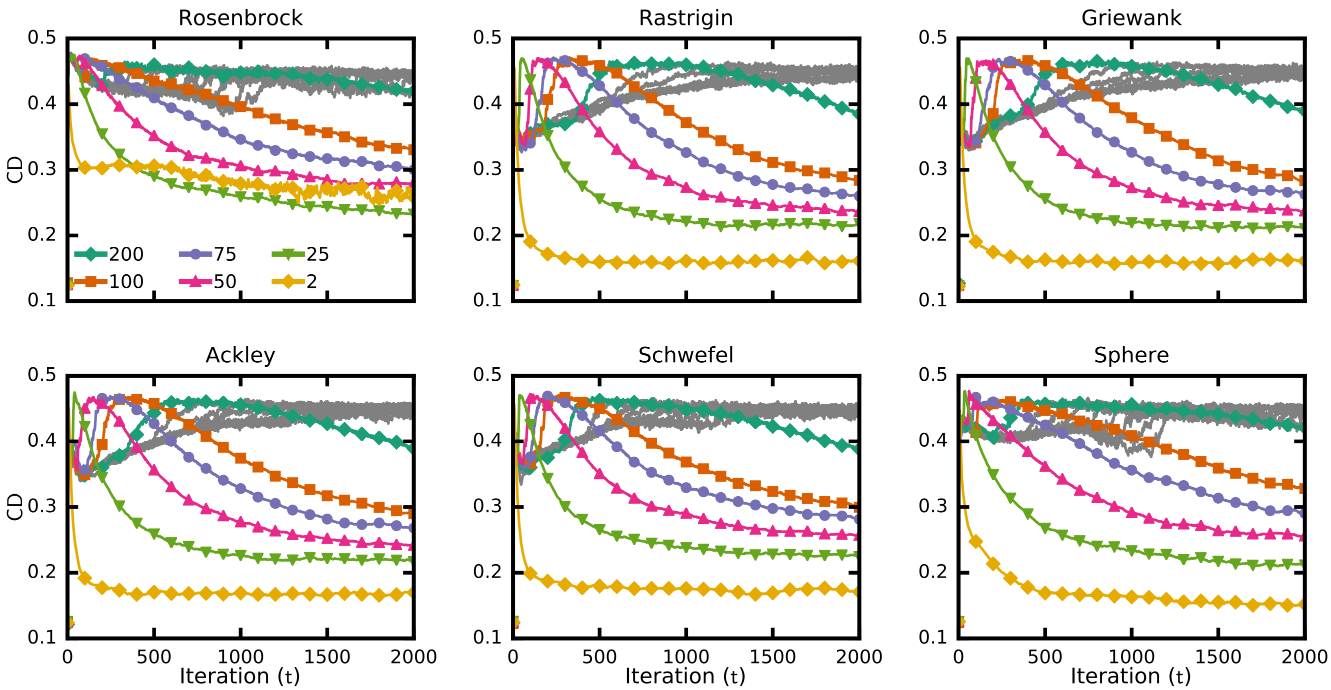

Fig. 3. Schemes of adaptive pace provided by the dynamic topology. Each plot shows the communication diversity CD of PSO along the iterations for the dynamic topology with different values of $p_k{-}failure^T$ {2, 25, 50, 75, 100, 200}.

more the dynamic topology resembles its initial ring topology. Similarly, the smaller this parameter (e.g., 2 and 25), the more frequent connections are rewired and the more the dynamic topology resembles a scale-free structure.

Further comparing the fitness found by PSO, functions appear to require faster paces when the global is better than the ring, and slower paces, otherwise. Nevertheless, PSO finds best solutions when using the dynamic topology with intermediate values of $p_k{-}failure$: lower values when functions require faster paces and higher values when functions require slower paces. In the function *Ackley*, for instance, PSO finds the best solution at the exploration-exploitation pace regulated by the dynamic topology with $p_k{-}failure$ threshold equals to 300.

B. Social Interactions Regulate Exploration-Exploitation Pace

Our results show that different topologies promote distinct patterns of social interactions which in turn regulate the exploration-exploitation pace (see Fig. 2). These distinct patterns of social interactions are revealed by different levels of communication diversity CD within the swarm. By assessing the components of the interaction network, CD measures the number of different information flows occurring in the swarm: the higher the CD, the more exploration, whereas the lower the CD, the more exploitation. Overall, the global topology exhibits the lowest CD which slightly changes along the iterations; the ring and dynamic topologies demonstrate significantly higher CD that varies along the iterations towards the exploration-exploitation pace required by each function.

The CD of the ring and dynamic topologies demonstrate their adaptive capabilities given that they increase and decrease their exploration-exploitation pace when functions require faster paces (e.g., Schwefel) and slower paces (e.g, Ackley), respectively. When compared to the ring and global topologies, the dynamic topology shows a higher adaptive capability to adapt its pace: by rewiring swarm connections to promote different patterns of social interactions, it more rapidly reaches values of CD that are more appropriate to the required pace by each function.

C. Schemes of Adaptive Pace

Finally, we show that topologies with higher adaptive capabilities such as the dynamic topology can promote patterns of social interactions that effectively regulate the exploration-exploitation pace at which best solutions can be found (see Fig. 3). Each adaptive scheme regulates distinct patterns of social interactions that promote different exploration-exploitation paces along the iterations.

Using the lowest and highest values of $p_k{-}failure$, the dynamic topology provides adaptation schemes that regulate patterns of social interactions closely resembling those of global and ring topologies, respectively. However, the dynamic topology is flexible to also provide intermediate adaptation schemes that regulate balanced patterns of social interactions and thus promote a better exploration-exploitation pace.

V. Conclusion

The essence of swarm intelligence are the social interactions. Depending on their diversity, the swarm can exhibit high exploitation and prematurely converge. To avoid premature convergence, the dynamic topology controls the exploration-exploitation balance by promoting patterns of social interactions with distinct levels of communication diversity. This mechanism makes the swarm to adapt itself towards an appropriate exploration-exploitation balance. Such optimal pace, however, depends on the particular problem in hand. Therefore, future approaches must be able to adapt the pace.

By assessing social interactions instead of analyzing individuals' properties, our method and results have the potential to be applied to different swarm techniques. Such generality opens the possibility to comprehensive analyses of swarm intelligence, providing the means to build better swarm-based tools.

All analyses performed here can be accessed at http://github.com/macoj/network_swarm/ and the same code can be applied to other swarm-based techniques.

Acknowledgment

Marcos Oliveira and Diego Pinheiro would like to thank the Science Without Borders program (CAPES, Brazil) for grants under number 1032/13-5 and 0624/14-4. Mariana Macedo would like to thank to the University of Pernambuco (PFA-UPE) and CAPES (Coordination for the Improvement of Higher-Education Personnel), Brazil, for the partial financial support for this paper.

References

[1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm intelligence: from natural to artificial systems*, 1st ed. New York, NY, USA: Oxford University Press, 1999.

[2] J. Kennedy and R. C. Eberhart, *Swarm Intelligence*, 1st ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2001.

[3] M. Dorigo and T. Stützle, *Ant colony optimization*, 1st ed. Cambridge, Massachusetts: MIT Press, 2004.

[4] H. Chaté and M. A. Muñoz, "Viewpoint: insect swarms go critical," *Physics*, vol. 7, p. 120, 2014.

[5] J. Kennedy and R. Eberhart, "Particle swarm optimization," vol. 4, 1995, pp. 1942–1948 vol.4.

[6] T. T. Hills, P. M. Todd, D. Lazer, A. D. Redish, and I. D. Couzin, "Exploration versus exploitation in space, mind, and society," *Trends in Cognitive Sciences*, vol. 19, no. 1, pp. 46–54, jan 2015.

[7] O. Olorunda and A. P. Engelbrecht, "Measuring exploration/exploitation in particle swarms using swarm diversity," in *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, Jun. 2008, pp. 1128–1134.

[8] Y. Shi and R. Eberhart, "Monitoring of particle swarm optimization," *Frontiers of Computer Science in China*, vol. 3, no. 1, pp. 31–37, 2009.

[9] J. Kennedy and R. Mendes, "Population structure and particle swarm performance," in *Evolutionary Computation, 2002. CEC '02. Proceedings of the 2002 Congress on*, vol. 2, 2002, pp. 1671 –1676.

[10] ——, "Neighborhood topologies in fully-informed and best-of-neighborhood particle swarms," in *Soft Computing in Industrial Applications, 2003. SMCia/03. Proceedings of the 2003 IEEE International Workshop on*, june 2003, pp. 45 – 50.

[11] R. Mendes, J. Kennedy, and J. Neves, "Watch thy neighbor or how the swarm can learn from its environment," in *Swarm Intelligence Symposium, 2003. SIS '03. Proceedings of the 2003 IEEE*, april 2003, pp. 88 – 94.

[12] W.-B. Du, W. Ying, and G. Yan, "The Impact of Population Structure on Particle Swarm Optimization: A Network Science Perspective," in *Advances in Swarm Intelligence*, 2016, pp. 341–349.

[13] C. Liu, W.-B. Du, and W.-X. Wang, "Particle Swarm Optimization with Scale-Free Interactions," *PLoS ONE*, vol. 9, no. 5, p. e97822, may 2014.

[14] J. Janostik, M. Pluhacek, R. Senkerik, I. Zelinka, and F. Spacek, "Capturing Inner Dynamics of Firefly Algorithm in Complex Network—Initial Study," in *Proceedings of the Second International Afro-European Conference for Industrial Advancement AECIA 2015*, ser. Advances in Intelligent Systems and Computing, A. Abraham, K. Wegrzyn-Wolska, A. E. Hassanien, V. Snasel, and A. M. Alimi, Eds. Cham: Springer International Publishing, 2016, vol. 427, pp. 571–577.

[15] M. Pluhacek, J. Janostik, R. Senkerik, I. Zelinka, and D. Davendra, "Pso as complex network—capturing the inner dynamics—initial study," 2016, pp. 551–559.

[16] M. Oliveira, C. J. Bastos-Filho, and R. Menezes, "Assessing particle swarm optimizers using network science metrics," in *Complex Networks IV*, ser. Studies in Computational Intelligence, G. Ghoshal, J. Poncela-Casasnovas, and R. Tolksdorf, Eds. Springer Berlin Heidelberg, 2013, vol. 476, pp. 173–184.

[17] M. Oliveira, C. Bastos-Filho, and R. Menezes, "Towards a Network-based Approach to Analyze Particle Swarm Optimizers," in *Swarm Intelligence (SIS), 2014 IEEE Symposium on*, dec 2014, pp. 166–173.

[18] M. Oliveira, C. J. A. Bastos-Filho, and R. Menezes, "Using network science to assess particle swarm optimizers," *Social Network Analysis and Mining*, vol. 5, no. 1, pp. 1–13, 2015.

[19] M. Oliveira, D. Pinheiro, B. Andrade, C. Bastos-Filho, and R. Menezes, *Communication Diversity in Particle Swarm Optimizers*. Cham: Springer International Publishing, 2016, pp. 77–88.

[20] M. Clerc and J. Kennedy, "The particle swarm - explosion, stability, and convergence in a multidimensional complex space," *Evolutionary Computation, IEEE Transactions on*, vol. 6, no. 1, pp. 58–73, 2002.

[21] D. Bratton and J. Kennedy, "Defining a standard for particle swarm optimization," *Swarm Intelligence Symposium, 2007. SIS 2007. IEEE*, pp. 120 –127, april 2007.

[22] A. P. Engelbrecht, "Particle swarm optimization: Global best or local best?" in *2013 BRICS Congress on Computational Intelligence and 11th Brazilian Congress on Computational Intelligence*, Sept 2013, pp. 124–135.

[23] M. Oliveira, C. J. A. Bastos Filho, and R. Menezes, *Using Network Science to Define a Dynamic Communication Topology for Particle Swarm Optimizers*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 39–47.

[24] P. N. Suganthan, N. Hansen, J. J. Liang, K. Deb, Y. P. Chen, A. Auger, and S. Tiwari, "Problem definitions and evaluation criteria for the cec 2005 special session on real-parameter optimization," Nanyang Technological University, Singapore, Tech. Rep., 2005.