

The Dynamics of Two Cognitive Heuristics for Coordination on Networks

Iyad Rahwan

¹Masdar Institute of Science & Tech., UAE²Massachusetts Institute of Technology, USA³University of Edinburgh, UK

Fatemeh Jahedpari

British University in Dubai, UAE

Sherief Abdallah

¹British University in Dubai, UAE²University of Edinburgh, UK

Abstract—This paper is about how groups solve global coordination problems such as the distributed graph coloring problem. We focused on scenarios in which agents are not able to communicate explicitly, but can rely on observing the momentary choices of their immediate neighbors in a social network. It has been reported that humans use two cognitive heuristics when solving such problems: (i) the *frequency-based* heuristic, where people make choices that minimize conflicts with neighbors; and (ii) the *degree-based* heuristic, where people avoid conflict with well-connected neighbors. In this paper, we present a model capable of capturing these cognitive heuristics to varying intensities. Then, through simulation, we shed light on the behavior of these heuristics under different classes of social networks. Our analysis generally speaks in favor of both heuristics, provided they are used in moderation, and illustrates the utility of taking social status (connectivity) into account.

I. INTRODUCTION

Networks capture fundamental relationships and interactions in biological, economic, social and artificial systems [1], [2], [3]. It is widely recognized that the structural properties of networks can significantly influence individual and collective behavior, such as consensus finding [4], short path identification [5], the evolution of cooperation [6], and the efficiency of market trading [7], [8], [9].

In a recent seminal paper in *Science*, Kearns *et al* conducted a laboratory study of the effect of network structure on the ability of humans to achieve global coordination through local coordination [10]. They used the distributed graph coloring (DGC) problem, which abstracts a wide range of collective coordination problems, ranging from scheduling classrooms, to social differentiation among peers. Each individual controls the color of a vertex in a network, and can only see its immediate neighbors. The collective goal is to color the entire network, using a minimum number of colors, such that no two neighbors have identical color.

Kearns *et al* studied the DGC problem under different network topologies (Figure 1), and captured the group's performance (measured in total number of color changes) in successfully coloring different kinds of networks. They reported that people often employ a *frequency-based* heuristic, “choosing colors that will result in the fewest local conflicts.” When people were given additional information about the connectivity of their neighbors (i.e. their degree), many reported using a *degree-based* heuristic, “attempting to avoid conflicts

with neighbors with high connectivity” [10]. The frequency-based heuristic reflects bias towards avoiding popular choices, while the degree-based heuristic reflects bias towards avoiding popular individuals.

Our aim in this paper is to shed more light on the frequency-based and degree-based heuristics in DGC problems. These heuristics are particularly relevant when agents do not have the ability to communicate explicitly (e.g. exchanging ‘no good’ messages to propagate their constraints [11] or relying on mediators [12]). The inability to exchange messages may be due to time constraints. For example, in a time-critical scenario, people may not have the time to communicate explicitly to decide how to coordinate, and may have to rely on observing and reacting to each other's actions instantaneously. People may also be unable to exchange messages due to the nature of the problem. For example, when one develops an expertise to differentiate oneself from peers in an organization, this person has no authority to ask others to avoid certain choices of training, but can only signal his/her own choice.

Our first contribution is a model that captures both the frequency-based and the degree-based heuristics, together with parameters that control their intensity.

The second contribution of this paper is shedding light (through extensive simulations) on the effectiveness of these naturally occurring heuristics on coordination in different classes of networks. We find that increasing the intensity of the above heuristics can often be useful. However, there are network structures in which this may be severely counter-productive. We identify a ‘safe’ range of intensities which seem to work well in a variety of networks. Furthermore, our model provides a possible explanation of the observation, by Kearns *et al*, that people's coordination can sometimes worsen when more information is provided to them.

Our results have implications on both artificial and natural coordination. By shedding light on the interplay between network structure and the effectiveness of naturally occurring heuristics, we provide means for identifying possible causes of mis-coordination in natural systems. On one hand, our results inform the design of automated local coordination heuristics (e.g. used by agents when coordinating with people). On the other hand, our work informs the design of organizational structures in a way that ensures successful coordination, *given* that people use particular local heuristics.

II. BACKGROUND

In this section, we formally define the Constraint Satisfaction Problem (CSP) and its distributed version. We then summarize the recent behavioral experiments, conducted by Kearns *et al* [10], to observe how people solve these kinds of problems.

Formally, a CSP is defined as follows [11], [12].

Definition 1 (CSP): A *Constraint Satisfaction Problem* consists of:

- A set of variables $X = \{X_1, \dots, X_n\}$;
- A domain Dom_i for each variable X_i ;
- A set of constraint predicates of the form $P_k(X_{k1}, \dots, X_{kj})$ defined on the Cartesian product $Dom_{k1} \times \dots \times Dom_{kj}$, and returning `true` exactly for variable assignments that satisfy the constraint.

Solving a CSP amounts to finding an assignment of values to variables that satisfies all constraints. The CSP is NP-complete in general, making the use of heuristics inevitable.

In a distributed CSP (or DisCSP), each variable is controlled by a single agent. By making a local assignment to its variable, an agent hopes to contribute to a global solution. Most algorithms for solving DisCSPs rely on some form of explicit communication among agents. For example, agents may communicate their current legal values to other agents, facilitating the so-called arc consistency. Alternatively, they may communicate suggested values and “no good” messages to implement asynchronous backtracking [11]. In this paper, we are interested in situations in which agents cannot communicate such messages, but instead, can only observe each others’ momentary variable assignments, which may change as they search for a solution.

The inability to exchange messages may be due to time constraints. For example, in a time-critical scenario, people may not have the time to communicate explicitly to decide how to coordinate, and may have to rely on observing and reacting to each other’s actions instantaneously. People may also be unable to exchange messages due to the nature of the problem. For example, when one develops an expertise to differentiate oneself from peers in an organization, this person has no authority to ask others to avoid certain choices of training, but can only signal his/her own choice [10].

In this paper, we focus on a specific DisCSP, namely the distributed graph coloring (DGC) problem, where each individual controls the color of a vertex in a network. The collective goal is to color the entire network, using a given minimum number of colors, such that no two neighbors have identical color. The DGC is an example of a *binary* DisCSPs, in which each constraint involves exactly two variables.

Kearns *et al* [10] studied the DGC problem, with human participants, under different network topologies (Figure 1). They begin with a simple cycle of 38 nodes. Then, they add 5 or 20 randomly chosen arcs, while preserving a chromatic




Fig. 1. Network classes we study (adopted from [10]), from top-left to bottom-right: simple cycle, 5-chord cycle, 20-chord cycle, 2 leader cycle, preferential attachment with $v = 2$ and $v = 3$ links initially added to each new node.

number of 2, resulting in small-world networks.¹ Finally, they investigate behavior on *scale-free* networks generated via *preferential attachment* [14], in which highly-connected individuals are more likely to receive further connections, leading to so-called *hubs*.

In their study, Kearns *et al* ran experiments with three informational conditions.

- 1) In the *low-information* condition, each participant could see (through a computer interface) his/her own color as well as that of each of its immediate neighbor.
- 2) In the *medium-information* condition, each participant could also see the degree of each of its neighbors (i.e. their connectivity).
- 3) In the *high-information* condition, participants could see the entire network.

Our focus in this paper is with modeling behavior in the first two conditions.

Kearns *et al* reported that participants self-reportedly employed two behavioral heuristics:

- 1) *Frequency-based Heuristic*: Many people reported “choosing colors that will result in the fewest local conflicts.”
- 2) *Degree-based Heuristic*: When given additional information about the connectivity of their neighbors (i.e. their degree), many reported “attempting to avoid conflicts with neighbors with high connectivity.”

The frequency-based heuristic reflects bias against popular choices, while the degree-based heuristic reflects bias against popular individuals.

Our aim is to understand the interplay between these behavioral biases and the group’s ability to achieve global coordination.

¹A small-world network has low average shortest-path between its nodes, while having a high clustering coefficient [13]. Most real world networks, including social friendship networks, are reported to have this property.

III. MODEL OF HEURISTIC BEHAVIOR

We model node behavior as follows. A node retains its color if it is locally consistent. Otherwise, if there are one or more unused color, one such color is selected randomly. In the case of irreconcilable conflict, the node chooses a new color c , from the set C , with probability $p(c)$ inversely proportional to each color's importance $I(c)$.

$$p(c) = \left(I(c) \sum_{c_i \in C} \frac{1}{I(c_i)} \right)^{-1} \quad (1)$$

To capture the frequency-based and degree-based heuristic, let N be the set of neighbors, and $N(c)$ be the set of neighbors with color c . Let $d(i)$ be the degree (i.e. number of neighbors) of node i . The importance of color c is:

$$I(c) = \left(\sum_{n_i \in N(c)} (d(n_i))^K \right)^Q \quad (2)$$

Parameter K captures the intensity of the degree-based heuristic, while Q captures the intensity of the frequency-based heuristic. Note that if $K = 0$, then $I(c) = |N(c)|^Q$; that is, the node ignores degree information when deciding on the importance of the colors in its neighborhood. That is, the importance of a color depends solely on the frequency of that color among neighbors, ignoring information about the degree of neighbors. This can be used to capture behavior under the low-information view.

With $Q = 1$, the relative frequency of colors is mirrored in their importance, but with $Q > 1$, the relative importance of a color increases exponentially as relative frequency increases. With $K \geq 1$, the relative importance of a color increases exponentially as the total degree of neighbors with that color increases (e.g. for $K = 2$ and $Q = 1$, red appearing in a single neighbor with degree 4 is eight times as important as blue appearing in two neighbors of degree 1 each).

The node's behavior is summarized in Algorithm 1, where:

- *conflict* is true if the node's current color is common with at least one of its neighbors.
- *previousColor* refers to the current color of the node.
- *unusedColors* refers to the set of colors which are not used with any of the nodes' neighbors.

IV. BEHAVIOR IN LOW INFORMATION VIEW

We begin by studying what Kearns *et al* refer to as the *low information view*, in which agents do not have access to information about their neighbor's degree, but only its current color [10]. In this setting, agents can only apply the frequency-based heuristic. In our model, this corresponds to assigning $K = 0$ in Equation 2, and varying the value of Q to capture the intensity of the frequency-based heuristic.

We simulated the heuristic on six different classes of networks, of 38 nodes each following Kearns *et al* (see Figure 1), with $K = 0$ and varying values of Q over 0, 0.125, 0.25, 0.5, 1, 2... 10. In each graph, we ran 100 random initial colorings. In the case of scale-free and chord

Algorithm 1 Node Behavior

Output: new action α

```

1: if  $\neg$ conflict then
2:    $\alpha \leftarrow$  previousColor
3: else if conflict  $\wedge$  (unusedColors  $\neq \emptyset$ ) then
4:    $\alpha \leftarrow$  some color  $c \in$  unusedColors selected uniformly at random
5: else if conflict  $\wedge$  (unusedColors =  $\emptyset$ ) then
6:   for each  $c \in C$  do
7:      $I(c) = \left( \sum_{n_i \in N(c)} (d(n_i))^K \right)^Q$ 
8:      $p(c) = \left( I(c) \sum_{c_i \in C} \frac{1}{I(c_i)} \right)^{-1}$ 
9:   end for
10:   $\alpha \leftarrow$  some color  $c \in C$  selected randomly according to  $p(\cdot)$ 
11: end if

```

networks, the graph structure was also generated randomly. Figure 2 summarizes our results, with the number of color changes under each value of Q averaged over 100 runs.

As expected, in the case of simple cycles, frequency bias has no effect (Figure 2(a)). In the leader cycle network, higher values of Q perform consistently better (Figure 2(d)). This may be due to the fact that as soon as one of the two hubs causes a significant portion ($> 50\%$) of its neighbors to fixate on one color, the hub itself is likely to fixate on the opposite color thanks to the frequency heuristic, which in turn forces all its remaining neighbors to coordinate correctly.

In all other networks, we observe that performance initially increases (i.e. average number of color changes decreases) as the value of Q increases. However, there appears to be a critical point (i.e. value of Q), after which performance worsens. The exact value of the critical point changes from one network structure to another. Indeed, in the case of 20-chord cycles, this worsening is not significant (Figure 2(c)). In PA and 5-chord cycles, however, performance worsens significantly as Q goes above 3. Indeed, in 5-chord cycles, performance can be even worse than dropping the frequency heuristic altogether (Figure 2(b)). Overall, the value $Q = 3$ appears to capture a consistently successful intensity of the frequency heuristic across all networks.

Network Class	Correlation coefficient	p -value	N
All	-.350**	.000	3600
Cycle Family	-.341**	.000	2400
PA Family	-.396**	.000	1200
Simple Cycle	0.006	0.875	600
5 Chord Cycle	-.330**	.000	600
20 Chord	-.480**	.000	600
2 Leader Cycle	-.486**	.000	600
PA (V=2)	-.415**	.000	600
PA (V=3)	-.376**	.000	600

TABLE I
CORRELATION COEFFICIENTS BETWEEN FREQUENCY BIAS (Q) AND AVERAGED SOLUTION TIME FOR $Q \leq 3$, $K = 0$. SIGNIFICANT CORRELATIONS ARE MARKED BY **

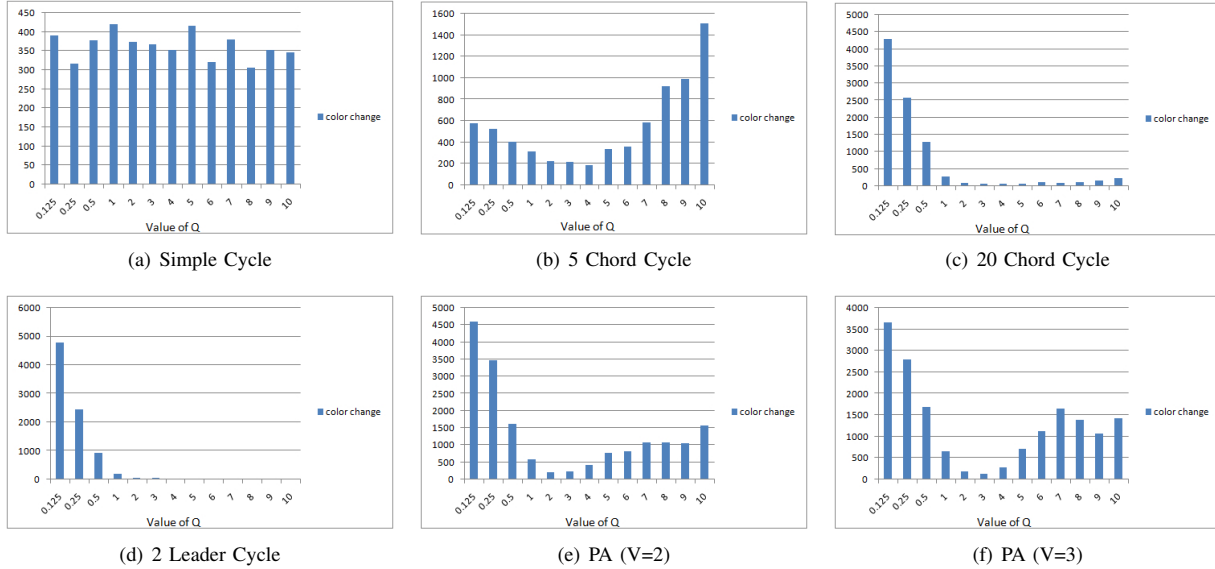


Fig. 2. Average number of color changes with $K = 0$ and varying value of Q

To confirm our observation of this special value of Q , we performed Pearson correlation tests between the intensity of frequency bias (i.e. varying values of Q with fixed $K = 0$) and the solution time. Results of the correlation tests for $Q \leq 3$ and $Q \geq 3$ are shown in Tables I and II, respectively. Significant correlations are marked with **, and N denotes the number of cases with non-missing values.

Network Class	Correlation coefficient	p -value	N
All	.157**	.000	4800
Cycle Family	.129**	.000	3200
PA Family	.215**	.000	1600
Simple Cycle	-0.04	0.262	800
5 Chord Cycle	.259**	.000	800
20 Chord	.147**	.000	800
2 Leader Cycle	-.410**	.000	800
PA (V=2)	.212**	.000	800
PA (V=3)	.218**	.000	800

TABLE II
CORRELATION COEFFICIENTS BETWEEN FREQUENCY BIAS (Q) AND AVERAGED SOLUTION TIME FOR $Q \geq 3$, $K = 0$. SIGNIFICANT CORRELATIONS ARE MARKED BY **

For $Q \leq 3$, Table I shows significant *moderate negative* correlation between frequency bias and average solution time in all networks except simple cycle. On the other hand, for $Q \geq 3$, Table II shows significant *weak positive* correlation between frequency bias and average solution time in PA networks, 5-chord cycles, and 20-chord cycles. Again, as expected, there is no significant correlation within simple cycle networks. Interestingly, leader cycles continue with *negative* correlation, but as Figure 2(d) shows, this is not a large improvement to what has already been achieved with $Q = 3$. Indeed, if we group cycle-based and PA network groups together, correlation is positive, meaning that intensities of

the frequency heuristic above $Q = 3$ are generally counter-productive.

Note that while our analysis above suggests that a value of approximately $Q = 3$ is a good choice in general, finding the precise *optimal* value of Q can be done using standard optimization techniques (e.g. using gradient ascent on Q to find the value that maximizes correlation in a particular network class or across all classes). We leave this to future work.

V. BEHAVIOR IN MEDIUM INFORMATION VIEW

In the previous section, we investigated the frequency-based heuristic in the low information view, in which agents only saw the current color of their immediate neighbors. In a different condition, Kearns *et al* allowed people to see the *degree* of each neighbor (i.e., the number of peers that neighbor had). This *medium information view* allows people to differentiate between their neighbors based on their connectivity, which can be seen as a measure of the neighbor's importance. In this condition, participants in Kearns *et al*'s behavioral experiment reported using a degree-based heuristic, in which they tend to avoid conflict with highly-connected neighbors. Note that this heuristic is used in conjunction with the frequency-based heuristic we studied earlier.

A. Varying Degree Bias with Moderate Fixed Frequency Bias

Our first aim is to understand if degree bias is an effective coordination heuristic, given a fixed (and moderate) frequency bias. To investigate this, we simulated the degree-based heuristic on six different graph classes of 38 nodes each following Kearns *et al* (see Figure 1), varying the value of K over 0, 0.125, 0.25, 0.5, 1, 2... 10, and fixing the frequency bias parameter to $Q = 1$. In each graph, we ran 100 random initial colorings –in the case of scale-free and chord networks, the graph structure was also generated randomly.

Our results speak in favor of the degree-based heuristic. As shown in Table III, with the exception of the simple and 5-chord cycle families, there is significant negative correlation between K and the average number of color changes. This means that taking into account the social importance of one's neighbors is generally a useful strategy towards achieving global coordination.

Network Class	Correlation coefficient	p -value	N
All	-.002	.853	8400
Cycle Families	.100**	.000	5600
PA Families	-.344**	.000	2800
Simple Cycle	.004	.884	1400
5 Chord Cycle	.292**	.000	1400
20 chord	-.267**	.000	1400
2 Leader cycle	-.499**	.000	1400
PA ($V=2$)	-.338**	.000	1400
PA ($V=3$)	-.350**	.000	1400

TABLE III
CORRELATION COEFFICIENTS BETWEEN DEGREE BIAS (K) AND AVERAGED SOLUTION TIME (WITH $Q = 1$). SIGNIFICANT CORRELATIONS ARE MARKED BY **

Figure 3 provides more detail, showing effect of varying K with fixed $Q = 1$. It is interesting to note that, as was the case with parameter Q , the value $K = 3$ also appears to be a reasonable choice across all networks. It is also interesting to observe that, unlike the case with Q , higher values of K are not counter-productive in PA networks. This may be due to the fact that the degree-based heuristic naturally exploits the highly-connected hubs found in PA networks, something the frequency-based heuristic misses completely.

B. Varying Both Heuristics Simultaneously

To shed more light on our model, it is important to understand the interplay between the frequency-based and degree-based heuristics more thoroughly. In particular, it is important to vary both their parameters simultaneously. Thus, we simulated the heuristics on the six different graph classes of 38 nodes, varying the values of *both* parameters K and Q as follows: 0, 0.125, 0.25, 0.5, 1, 2 . . . 10. In each graph, we ran 100 random initial colorings –in the case of scale-free and chord networks, the graph structure was also generated randomly. All figures are plotted as heat maps with the two axes denoting K and Q . Figures are shown in gray-scale, with lighter colors denoting a higher number of color changes in solving the CSP, averaged over 100 runs.²

Variations in K and Q have no significant effect on performance in the simple cycle (Figure 4), which is expected since the symmetry of neighbors renders the heuristic useless.

In preferential attachment (PA) networks, the average number of color changes needed to find a solution decreases sharply as either K or Q increase. Figure 5 shows the case of PA networks (with $v = 3$). PA networks with ($v = 2$) and 2-leader cycles follows a similar pattern. This shows that either

²We note that all patterns observed were similar when using the median instead of average.




Fig. 4. Heat map of the average number of color changes required for solving simple cycles under varying values of parameters K and Q .




Fig. 5. Heat map of the average number of color changes required for solving PA networks (with $v = 3$), under varying values of parameters K and Q . PA networks with $v = 2$, leader cycles and 20-chord cycles follow a similar pattern.

heuristic (or combination thereof) would be effective in these types of networks.




Fig. 6. Heat map of the average number of color changes required for solving 5-chord cycles.

Surprisingly, in 5-chord cycles (simple cycle with 5 randomly added connections), the pattern is completely reversed –bias towards degree or frequency leads to good average performance up to a point, after which it becomes severely

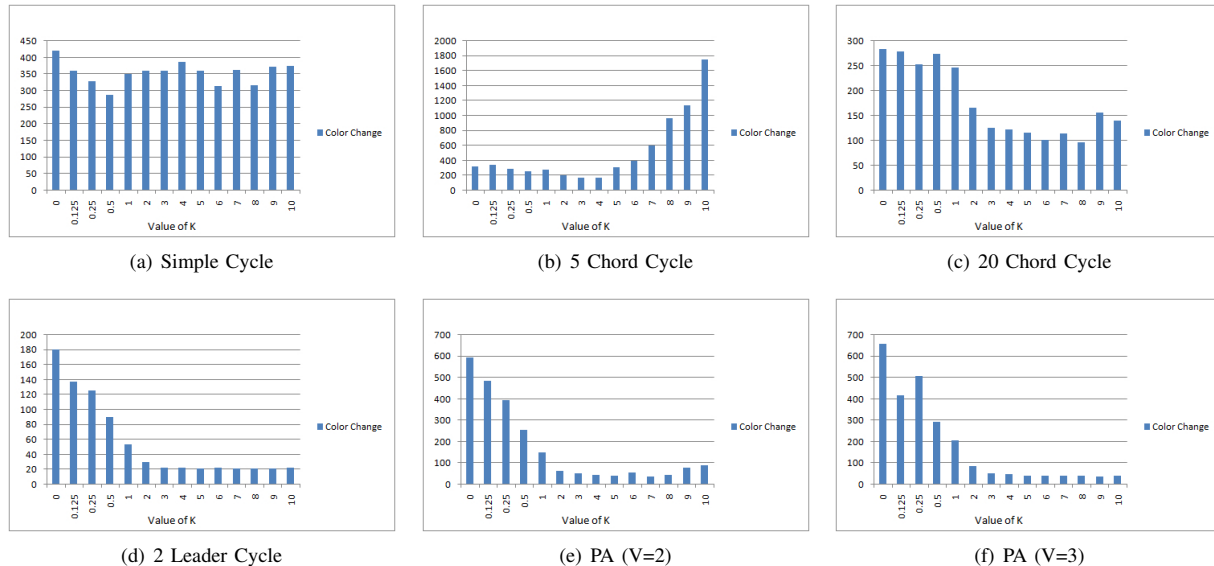


Fig. 3. Average number of color changes with $Q = 1$ and varying value of K

counter-productive (Figure 6).

However, as more random links are added in the case of 20-chord cycles, the pattern reverts back, with the problem becoming consistently easier with higher K and Q . One possible explanation is that as more hubs emerge, successful local coordination becomes more indicative of good global coordination.

VI. EXPLAINING PA NETWORK DIFFICULTY

As a side exercise, we were interested in explaining a peculiar observation by Kearns *et al* [10]. Kearns *et al* observed that people find it more difficult to solve PA networks, when they are shown degree information. This observation is counter-intuitive, as one would expect the additional information available to individuals to be useful or, if not useful, easily ignored.

It is clear that our model does not generate this pattern on average (Figure 5). However, we offer one plausible explanation of this phenomenon. When we plotted the worst-case performance (i.e. highest number of color changes for any of the 100 network in the class) for PA networks under each combination of K and Q values (see Figure 7), we observe that with sufficiently high Q , increasing K increases difficulty. This suggests that the difficulty observed by Kearns *et al* may be due to a particularly unfortunate choice of networks they studied.³ Upon closer inspection, we observed that the difficulty distribution of PA networks is heavy-tailed (Figure 8), which supports our explanation.

The lesson learned: the degree-based and frequency-based heuristics are effective, up to a point, after which they may be counter-productive in some networks.

³It is not clear how many of the 6-7 trials per-network class involved a medium-information view [10], hence it is not immediately clear if their observation was statistically significant across the entire PA network class.




Fig. 7. Heat map of the worst-case in PA networks (with $v = 3$) under each combination of K and Q values.




Fig. 8. Histogram showing distribution of PA ($v = 3$) network difficulty, measured as the number of color changes needed to find a solution, using $K = 0$ and $Q = 1$.

VII. DISCUSSION & CONCLUSION

This paper was concerned with how groups solve coordination problems such as the distributed graph coloring problem. We focused on scenarios in which agents are not able to communicate explicitly, e.g. by exchanging ‘no good’ messages to propagate their constraints among each other [11], or by relying on mediators [12]. In such situations, agents must simply rely on observing the momentary variable assignment behavior of its immediate neighbors, which may change as they collectively search for a solution.

We investigated two human-inspired heuristics for making use of such social information, as observed in recent behavioral studies of humans solving the DGC problem [10]. We presented a model capable of generating these heuristics to varying intensities. We then used the model to understand the behavior of the heuristics under different network structures.

Our work on coordination is in the spirit of recent work on the cognitive modeling of other human collective behaviors, such as contagion and imitation [15]. Our results have implications on both artificial and natural coordination. By shedding light on the interplay between network structure and the effectiveness of naturally occurring heuristics, we provide means for identifying possible causes of mis-coordination in natural systems. On one hand, our results inform the design of automated local coordination heuristics (e.g. used by agents when coordinating with people). On the other hand, our work informs the design of organizational structures in a way that ensures successful coordination, *given* that people use particular local heuristics.

There are many directions of future work. We focused on the average time needed to find a solution, i.e. on the final outcome. To complement this, we are interested in investigating the dynamics of different heuristics *during* their search for a solution.

Another direction of future research is investigating human behavior when they are allowed a more expressive communication language, such as those used in algorithmic approaches to the DisCSP. For example, some approaches allow agents to communicate suggested values and no-good messages [11]. There are no empirical results on the kinds of heuristics people would use in such scenarios, so new behavioral experiments are needed.

REFERENCES

- [1] M. O. Jackson, *Social and Economic Networks*. Princeton NJ, USA: Princeton University Press, 2008.
- [2] J. Kleinberg, “The convergence of social and technological networks,” *Communications of the ACM*, vol. 51, no. 11, pp. 66–72, 2008.
- [3] M. E. J. Newman, “The structure and function of complex networks,” *SIAM Review*, vol. 45, no. 2, pp. 167–256, 2003.
- [4] M. Kearns, S. Judd, J. Tan, and J. Wortman, “Behavioral experiments on biased voting in networks,” *Proceedings of the National Academy of Science*, vol. 106, no. 5, pp. 1347–1352, 2009.
- [5] J. Kleinberg, “Navigation in a small world,” *Nature*, vol. 406, p. 845, August 2000.
- [6] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak, “A simple rule for the evolution of cooperation on graphs and social networks,” *Nature*, vol. 441, no. 25, pp. 502–505, May 2006.
- [7] G. Charness, M. Corominas-Bosch, and G. R. Fréchet, “Bargaining and network structure: An experiment,” *Journal of Economic Theory*, vol. 136, no. 1, pp. 28–65, 2007.
- [8] J. S. Judd and M. Kearns, “Behavioral experiments in networked trade,” in *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*. New York, NY, USA: ACM, 2008, pp. 150–159.
- [9] F. C. Santos, M. D. Santos, and J. M. Pacheco, “Social diversity promotes the emergence of cooperation in public goods games,” *Nature*, vol. 454, no. 10, pp. 213–216, July 2008.
- [10] M. Kearns, S. Suri, and M. Montfort, “An experimental study of the coloring problem on human subject networks,” *Science*, vol. 313, pp. 824–827, 2006.
- [11] M. Yokoo and K. Hirayama, “Algorithms for distributed constraint satisfaction: A review,” *Autonomous Agents and Multi-Agent Systems*, vol. 3, no. 2, pp. 185–207, September 2000.
- [12] R. Mailler and V. Lesser, “Asynchronous Partial Overlay: A New Algorithm for Solving Distributed Constraint Satisfaction Problems,” *Journal of Artificial Intelligence Research*, vol. 25, pp. 529–576, April 2006. [Online]. Available: <http://mas.cs.umass.edu/paper/397>
- [13] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 409–10, 1998.
- [14] A.-L. Barabási and R. Albert, “Emergence of scaling in random networks,” *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [15] N. Fridman and G. A. Kaminka, “Towards a cognitive model of crowd behavior based on social comparison theory,” in *Proceedings of the 22nd national conference on Artificial intelligence (AAAI)*. AAAI Press, 2007, pp. 731–737.