Distributed team formation in human networks

Lorenzo Coviello
UC San Diego
lcoviell@ucsd.edu

Massimo Franceschetti
UC San Diego
massimo@ece.ucsd.edu

***Abstract*—This paper describes a model of distributed coordination in which leaders are required to recruit teams of followers according to a minimum team size constraint. The objective is to reach a state in which all possible followers are recruited and all team size constraints are satisfied. We propose a simple, distributed algorithm for group formation and analyze its performance. The algorithm is suitable to model human behavior, being memoryless and with minimum computation requirements. Our results show that while in any network all followers are recruited in a time that is linear in the network size, there are networks in which ensuring a minimum team size constraint may require a number of rounds that is exponentially large in the network size. On the other hand, our algorithm always converges to an *approximate* solution in polynomial time, namely the leaders quickly form teams in which the total number of additional followers required to satisfy all team size constraints is an arbitrarily small fraction of the entire population.**

I. Introduction

The discovery of the algorithmic processes governing humans' behavior over networks is of interest in a variety of contexts, ranging from sociology and economics to tactical scenarios of soldiers deployed in foreign countries. This paper attempts to provide a model of human behavior in a distributed coordination game of team formation. Consider, for example, a set of agents being deployed undercover in a foreign country. These agents cannot communicate among each other and have no access to a central infrastructure. Their objective is each to recruit a team of multiple followers in order to perform a certain task. Hence, each leader aims to recruit a certain number of followers, each follower can join only one leader, and all leaders want to reach a state in which all team size constraints are satisfied. We give a simple mathematical model of the above situation, propose a "human-like" algorithm to solve this coordination task and study its performance in terms of convergence to a stable solution. The proposed algorithm is suitable to model human behavior, being memoryless and with minimum computation requirements. Our approach follows the one of our previous work on matching [2] in which the coordination task was limited to teams of size two (composed of one leader and one follower only). In [2] we also validated theoretical results with experiments on real human subjects connected over a virtual network. The validation of the algorithm proposed here with human experiments is left as part of future work.

Distributed group formation has been studied extensively in the context of robotics and consensus algorithms, see [1], [4]. We distinguish ourselves by being focused on "human-like" algorithms, i.e. algorithms that are computationally simple, memoryless and suitable to model human behavior.

II. Problem definition

Consider a population whose agents belong to two different classes, leaders and followers. Leaders are required to recruit and maintain teams of followers according to minimum size requirements, while followers can only accept or reject leaders' requests. Beyond the individual task of building their own teams, leaders are given the common tasks of achieving collective stability, i.e. a situation in which *all* followers are recruited and *every* leader controls a team that is sufficiently large.

A leader might not be allowed to recruit every follower, as the recruitment process is constrained by a communication network according to whom only a subset of leader-follower pairs can communicate. In addition, the team size requirements may differ between leaders, and followers can transfer from a team to another by accepting leaders' requests.

The communication structure of the population is captured by a bipartite network $G = (L \cup F, E)$ whose nodes' partition represents the set L of leaders and the set F of followers, and where there exists an edge between follower f and leader ℓ if and only if f and ℓ can communicate between each other. For each leader $\ell \in L$, let N_ℓ be the neighborhood of ℓ in G, that is, the set of the followers ℓ can communicate with, and let $c_\ell \geq 1$ be the *minimum size* of the team ℓ is required to form. We also assume that leaders have an infinite capacity to accommodate followers, i.e. there is no upper bound to a team's size.

A subset $M \subseteq E$ is a *matching* of G if for all $f \in F$ there exists at most a single $\ell \in L$ such that $(\ell, f) \in M$. This is consistent with the fact that a follower can be matched to a single leader at a time, while a leader is matched to all followers in his team. A matching M of G uniquely identifies a set of teams $\{T_\ell(M) : \ell \in L\}$, where $T_\ell(M)$ is the team of leader $\ell \in L$ under the matching M. More precisely, given a matching M of G, $T_\ell(M)$ contains all followers that are matched to ℓ in M, i.e. $T_\ell(M) = \{f \in F : (\ell, f) \in M\}$. Given constraints c_ℓ for all $\ell \in L$, it is natural to define the following notion of stability of a matching.

Definition 1. *A matching M of G is stable if and only if all followers are matched in M and $|T_\ell(M)| \geq c_\ell$ for all $\ell \in L$.*

We are only interested in networks which admit stable

matchings. Together with the assumption that $c_\ell \geq 1$ for all $\ell \in L$, this implies that $|L| \leq |F|$.

In the remaining part of this work, we propose a simple, *distributed* algorithmic model for group formation in leader-follower networks, and we analyze its performance with respect to convergence to stability. Our first result, Theorem 1, states that, in any network with n leaders and m followers, all followers are quickly recruited in $O(m)$ rounds of the algorithm with high probability. A similar guarantee on the algorithm's performance does not hold in general for the other condition of stability of a matching, that is, $T_\ell(M) \geq c_\ell$ for all $\ell \in L$. To show this, we define a family of bipartite networks and prove that our algorithmic model requires exponentially many rounds with high probability to converge to the unique stable matching when starting from a constant fraction of all possible configurations (Theorems 2 and 3). However, we are able to provide performance guarantees on the convergence to an arbitrarily close approximation of a stable matching for any graph with unique stable matching, as show by Theorem 4. In particular, the total number of additional followers required to satisfy all team size constraints is an arbitrarily small fraction of the entire population. Computer simulations corroborate our theoretical results.

III. Algorithmic model

In this section, we propose a synchronized model for team formation, specifying the behavior of leaders and followers. The algorithm for leader $\ell \in L$ is shown in Table 1. At each round of the algorithm, given the current matching M of G, if either ℓ has a team of size $|T_\ell(M)| < c_\ell$ or there exists unmatched follower in N_ℓ, then ℓ attempts to recruit a new follower according to the function Recruit(ℓ). In our implementation, with probability $1 - p$, Recruit(ℓ) does nothing; with probability p, it chooses an unmatched follower in N_ℓ uniformly at random, if any; otherwise, it chooses a follower in $N_\ell \backslash T_\ell(M)$ uniformly at random. In words, leaders always prefer to recruit followers which are currently unmatched over matched ones.

As for the followers, the algorithm for $f \in F$ is represented in Table 2. At each round of the algorithm, if f has incoming recruitment requests, then he chooses the team to join according to the function ChooseTeam(f). In our implementation, ChooseTeam(f) rejects each incoming request independently with probability $1 - q$. One among the pending requests is chosen uniformly at random (and f joins the corresponding team) and all the others are rejected. Observe that a follower is equally likely to join a team when unmatched and to switch team when currently matched. We assume this for the sake of simplicity, and our results hold for more general implementations of the function ChooseTeam(f).

Table 1 Stub for leader $\ell \in L$

if $|T_\ell(M)| < c_\ell$ **or** $\exists$ unmatched $f \in N_\ell$ **then**
  Recruit(ℓ)
end if

Table 2 Stub for follower $f \in F$

if f has incoming requests **then**
  ChooseTeam(f)
end if

It is interesting to observe that the proposed algorithm is *self-stabilizing*, i.e. once a stable matching is reached, leaders stop recruiting followers. Moreover, it is a single-stage algorithm, that is, agents do not change their behavior until stability is reached.

In the model just defined, leaders do not communicate between each other, and the behavior of a leader only depends on the size of his own team and the status of his neighborhood. The collaboration between leaders is therefore minimal and consists in a leader attempting to recruit new followers only when his team size constraint is not met or there are unmatched followers in his neighborhood. This form of *honesty* between leader is justified by the common goal of reaching a stable matching.

IV. Convergence to stable matchings

In this section, we analyze the algorithm's performance with respect to convergence to a stable matching. The following claim, whose proof is straightforward, implies that the first condition of stability of a matching is reached and maintained after all followers are recruited for the first time.

Claim 1. *After follower f is recruited for the first time, he will not remain unmatched after each successive round of the algorithm (i.e. f can only switch team).*

Our first result states that, in any bipartite network with n leaders and m followers, all followers are recruited in a number of rounds that is linear in m.

Theorem 1. *In any bipartite network G of n leaders and m followers which admits a stable matching, all followers are recruited in $O(m)$ rounds with high probability.*

Proof: Recall that, at each round of the algorithm, if there exist unmatched followers in N_ℓ then leader ℓ attempts to recruit one among them uniformly at random, and that there is no maximum allowed team size. Consider any round of the algorithm and let $f \in F$ be any follower who is unmatched at the beginning of this round. As G admits a stable matching, there exists $\ell \in L$ such that $(\ell, f) \in E$ (i.e., f is connected to a leader at least). Moreover $|N_\ell| \leq m$ (i.e. ℓ is connected to m followers at most). Therefore, f is recruited during the current round with probability at least pq/m (as each $\ell \in L$ such that $f \in N_\ell$ proposes to f with probability at least p/m and f joins *a* team with probability q). By independence of successive rounds of the algorithm, the number of rounds needed to recruit f is upper bounded by a Geometric random variable of mean m/pq, and for any $c > 2$ is at most cm with high probability for m large enough. The Theorem follows by observing that the argument above holds for all unmatched

followers and that at the first round of the algorithm all m followers are unmatched. ■

A similar guarantee on the algorithm's performance does not hold in general for the other condition of stability of a matching M, that is $T_\ell(M) \geq c_\ell$ for all $\ell \in L$. First, we define a sequence of networks such that $c_\ell = 1$ for each leader ℓ and show that, from a specific configuration, convergence to the stable matching is exponential in the network's size with high probability (Theorem 2). Then, we show that the configurations from which convergence is exponential constitute a constant fraction of all possible configurations.

For $n \geq 1$, let $G_n = (L_n \cup F_n, E_n)$ be the bipartite network, shown in Figure 1, with n leaders and n followers (i.e. $L = \{\ell_1, \ldots, \ell_n\}$ and $F = \{f_1, \ldots, f_n\}$), with edges $E_n = \{(\ell_i, f_j) : 1 \leq i \leq n, j \leq i\}$, and constraint $c_{\ell_i} = 1$ for all $1 \leq i \leq n$. G_n has a unique stable matching $M^*_n = \{(\ell_i, f_i) : 1 \leq i \leq n\}$.

Theorem 2. *Consider the matching M'_n of G_n defined by*

$$M'_n = \{(\ell_i, f_{i-1}) : 2 \leq i \leq n\} \cup \{(\ell_n, f_n)\}.$$

Then, the number of rounds required to converge to the stable matching M^*_n starting from M'_n is exponential in n with high probability.

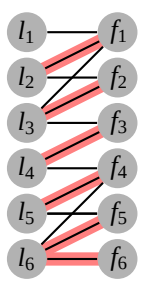

Fig. 1. The network G_n for $n = 6$. The matching M'_n is highlighted.

Here we only provide a sketch of the proof, which follows the same argument of [2, Theorem 3]. Given a matching M, we say that a leader ℓ is *poor* if $T_\ell(M) < c_\ell$, *rich* if $T_\ell(M) > c_\ell$, and *stable* otherwise. Observe that, under M'_n, leader ℓ_1 is poor, leader ℓ_n is rich, while all other leaders are stable. According to the algorithm, as all followers are matched in M'_n, only ℓ_1 attempts to recruit new followers. Due to the network structure, ℓ_1 can only recruit f_1 (currently in ℓ_2's team). If f_1 accepts, then ℓ_1 becomes stable and ℓ_2 becomes poor (and can in turn attempt to recruit either f_1 or f_2). After each round of the algorithm all followers are matched (by Claim 1) and there exists a unique poor leader until the stable matching is reached. The stochastic process tracking the position of the poor leader is not a classical random walk on L_n and its transition probabilities depend on the current matching. Following the approach in [2], we can show that the random process defining the position of the poor leader *and* the current matching is equivalent to a random walk on a tree H_n whose size is exponential in n. Observing that the time needed to reach the stable matching of G_n is equivalent to the hitting time of the root of H_n when starting from the furthest leaf, we can derive the exponential bound of Theorem 2.

Theorem 2 only considers a single matching of G_n and raises the question whether the exponential convergence is proper of M'_n exclusively. The next result states that the convergence to the stable matching of G_n requires exponentially many steps when starting from a constant fraction of all possible configurations (i.e. matchings of G_n in which there exists a unique leader with an empty team).

Theorem 3. *The algorithm requires exponentially many rounds with high probability to converge to the stable matching of G_n when starting from a constant fraction of all matchings of G_n with a single unmatched leader.*

The proof, which uses the same argument in [2, Theorem 4] and therefore is not reported here, consists in counting all configurations which correspond to nodes of H_n at a given distance from the root (that are a constant fraction of all nodes of H_n). Theorems 2 and 3 alone do not imply that the number of rounds to convergence to the stable matching of G_n is exponential in n with constant probability (and therefore in expectation), but simulations' results suggest that this is indeed the case.

V. Convergence to approximate stable matchings

Despite the negative results of Theorems 2 and 3, we are still able to provide performance guarantees on the convergence to arbitrarily close approximations of a stable matching. In our result, stated by Theorem 4, the notion of approximation of a stable matching reads in term of the *surplus* of a matching, that we define as follows.

Definition 2. *Given a bipartite network $G = (L \cup F, E)$ with constraints $c_\ell \geq 0$ for all $\ell \in L$, and a matching M of G, the surplus $s(M)$ of M is defined as*

$$s(M) = \sum_{\ell \in L} \max\{0, |T_\ell(M)| - c_\ell\}.$$

In words, $s(M)$ sums the numbers of followers each leader ℓ has in excess of its constraint c_ℓ, and each leader ℓ such that $|T_\ell(M)| \leq c_\ell$ does not contribute to $s(M)$. Observe that, if $s(M) = 0$ and all followers are matched in M then M is a stable matching[1].

For a network G, let $\tau_0(G)$ be the first round of the algorithm after which all followers are recruited (notice that $\tau_0(G)$ is the quantity upper bounded by Theorem 1). The following claim, whose proof is straightforward, states that when all followers are recruited the surplus never increases and is upper bounded by the number m of followers (as $c_\ell \geq 1$ for all $\ell \in L$).

Claim 2. *Given a network G with n leaders and m followers admiting a stable matching, let $M(t)$ be the matching of G at round $t \geq 0$. Then, for all $t \geq \tau_0(G)$, $s(M(t))$ is non-increasing in t, and $s(M(t)) < m$.*

[1]However, in general it is not true that if M is stable and all followers are matched in M then $s(M) = 0$.

We focus our attention on bipartite networks such that each stable matching has zero surplus. Under this assumption[2], the surplus of a matching constitutes some sort of *distance* from this matching to the set of the stable matchings. We are now ready to state our result about the convergence to arbitrarily close approximations of a stable matching.

Theorem 4. *Let $G = (L \cup F, E)$ be a bipartite network with n leaders and m followers, and such that $s(M) = 0$ for each stable matching M of G. Then, for any $\varepsilon > 0$, a matching with surplus strictly smaller that εm is reached in a number of rounds that is polynomial in m with high probability.*

For the proof of Theorem 4 we need to introduce the novel concept of *surplus-decreasing* path, which is the adaptation of the notion of augmenting path to our setup.

Definition 3. *Given a bipartite network $G = (L \cup F, E)$ and a matching M of G, a cycle-free path $P = \ell_0, f_1, \ell_1, \ldots, f_k, \ell_k$ (of even length $2k$) is a surplus-decreasing path relative to M if $(\ell_i, f_i) \in M$ for all $1 \leq i \leq n$, ℓ_0 is a poor leader, ℓ_n is a rich leader and ℓ_i is a stable leader for all $1 < i < n$, and $\ell_i \neq \ell_j$ for all $i \neq j$, $f_i \neq f_j$ for all $i \neq j$.*

In words, a surplus-decreasing path starts from a poor leader and an unmatched edge, ends at a rich leader and a matched edge, and alternates unmatched and matched edges containing stable leaders and followers. Observe that, if $s(M) > 0$ and P is a surplus-decreasing path relative to M, by switching each unmatched edge of P into a matched edge, and vice versa, we obtain a new matching M' such that $s(M') = s(M) - 1$.

We say that two paths are *follower-disjoint* if they do not share any follower (however, they might share some leader). The following lemma, needed to prove Theorem 4, is an adaptation of [3, Theorem 1] to our setup. The symmetric difference of two sets A and B is denoted by $A - B$.

Lemma 1. *Let $G = (L \cup F, E)$ be a bipartite network with n leaders and m followers, and such that $s(M') = 0$ for each stable matching M' of G. Let M be a matching of G with surplus $s(M) > 0$ and let M^* be a stable matching of G. Then, in $M - M^*$ there are $s(M)$ follower-disjoint surplus-decreasing paths relative to M.*

The proof is similar to the one in [3], with two simple adjustments. First, our definition of matching allows for a leader being part of multiple edges. Second, given a matching M with $s(M) > 0$, for each unit of surplus we can connect the corresponding rich leader to an auxiliary unmatched follower (to whom the leaders matches as soon as it looses a follower from his team). In this way, reducing the surplus by one corresponds to increasing the matching size by one. We also make use the following corollary of Lemma 1.

Corollary 1. *Let $G = (L \cup F, E)$ be a bipartite network such that $s(M') = 0$ for each stable matching M' of G. Let M be a matching of G with surplus $s(M) \geq \varepsilon m$, for some $\varepsilon > 0$. Then, there exists a surplus-decreasing path relative to M of length at most $2\lfloor 1/\varepsilon \rfloor$.*

Proof: By Lemma 1, if $s(M) \geq \varepsilon m$ and M^* is a stable matching of G then $M - M^*$ contains εm follower-disjoint surplus-decreasing paths relative to M, of cumulative length at most $2m$. Necessarily, there is one of length at most $2\lfloor 1/\varepsilon \rfloor$. ■

We are now ready to present the proof of Theorem 4.

Proof of Theorem 4: Consider a bipartite network G with n leaders and m followers. Let $M(t)$ denote the matching of G at time $t \geq 0$. Recall that we defined $\tau_0(G)$ as the first round after which all followers are recruited, and by Theorem 1, $\tau_0(G)$ is $O(m)$ with high probability. Moreover, by Claim 2, $s(M(t))$ is nonincreasing in t for all $t \geq \tau_0(G)$. For all $0 < \varepsilon < 1$, let

$$\tau_\varepsilon = \min\{t \geq \tau_0 : s(M(t)) < \varepsilon m\}$$

be the random variable representing the first round of the algorithm during which the surplus becomes strictly less than εm.

Consider any round $t_1 \geq \tau_0(G)$. By Claim 2, $s(M(t_1)) < m$, and therefore there exists[3] $0 < \varepsilon' < 1$ such that $s(M(t_1)) = \varepsilon' m$. Observe that $\tau_{\varepsilon'}$ can be equivalently defined as

$$\tau_{\varepsilon'} = \min\{t > t_1 : s(M(t)) < s(M(t_1))\},$$

i.e. it is the fist round after t_1 during which the surplus decreases. We show that $\tau_{\varepsilon'}$ is polynomial in the network size with high probbility.

Let $\Delta = \max\{|N_\ell| : \ell \in L\}$ be the maximum degree of a leader in G ($1 \leq \Delta \leq m$). Let $2h(t) \geq 2$, be the length of the shortest surplus-decreasing path relative to $M(t)$. By Corollary 1, $h(t) < \lfloor 1/\varepsilon' \rfloor$ for all $t_1 \leq t < \tau_{\varepsilon'}$. We distinguish the cases of $h(t_1) = 1$ and $h(t_1) > 1$.

Let $h(t_1) = 1$. We show that with probability at least pq/Δ the surplus decreases of at least one unit during round t_1. Consider a rich leader $\ell \in L$. ℓ contributes with $u \geq 1$ units to the surplus, and at most $u \leq \varepsilon' m$ poor leaders $\ell_1, \ldots, \ell_u$ are at distance 2 from ℓ along distinct surplus-reducing paths. If follows that with probability at least pq/Δ one among $\ell_1, \ldots, \ell_u$ takes a follower f away from $\ell's$ team (by asking f with probability at least p/Δ and having f accept with probability at least q). Observe that followers in ℓ's team can be recruited only by poor leaders at distance 2 from ℓ, and that the surplus still decreases by one even if each of $\ell_1, \ell_2, \ldots, \ell_u$ is at distance 2 form ℓ and successfully recruits a distinct follower from ℓ's team.

Let now $h(t_1) \geq 1$. As $h(t) < \lfloor 1/\varepsilon' \rfloor$ for all $t_1 \leq t < \tau_{\varepsilon'}$, we track the length $2h(t)$ of the shortest surplus-decreasing path. We claim that, with probability at least pq/Δ, $h(t)$ decreases of at least one unit during round t for all $t_1 \leq t < \tau_{\varepsilon'}$. Let $P = \ell_0, f_1, \ell_1, \ldots, f_{h(t)}, \ell_{h(t)}$ be a shortest surplus-decreasing path relative to $M(t)$, of length $2h(t) \leq 2\lfloor 1/\varepsilon' \rfloor$. With probability at least pq/Δ, during round t, the poor leader ℓ_0 attempts to recruit f_1 and f_1 accepts, and thus ℓ_1 becomes poor. Observe

[2]This assumption is for sake of simplicity and our result holds also in a more general formulation.

[3]We assume $\varepsilon > 0$, the case $\varepsilon = 0$ being trivial.

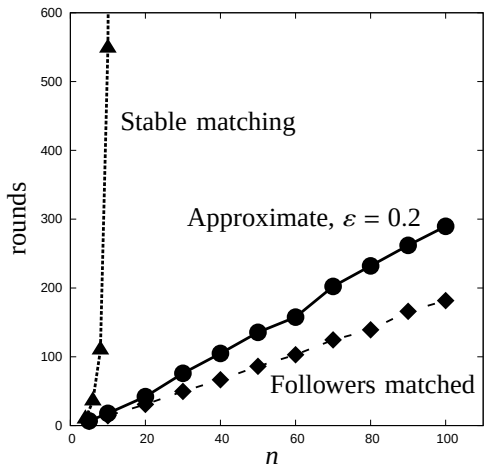

Fig. 2. Simulation results on the sequence of networks G_n.

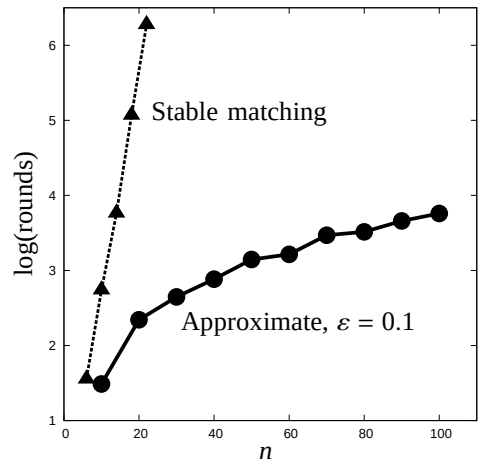

Fig. 3. Simulation results on the sequence of networks G_n, logarithmic scale.

that, as long as $h(t) > 1$, $\ell_{h(t)}$ remains rich during round t because any other surplus-decreasing path ending at $\ell_{h(t)}$ would have length at least $2h > 2$. Therefore, $h(t)$ decreases by one with probability at least pq/Δ during round t. By independence of successive rounds of the algorithm (and the upper bound on $h(t)$ for $t_1 \leq t < \tau_{\varepsilon'}$), $h(t)$ decreases of one unit in a number of rounds that is upper bounded by a Geometric random variable with mean m/pq (as $\Delta \leq m$), and therefore, for any $c > 2$, is at most cm/pq with high probability for m large enough.

Given the result for $h(t) = 1$ and that $h(t) < \lfloor 1/\varepsilon' \rfloor$ for all $t_1 \leq t < \tau_{\varepsilon'}$, it follows that the process $\{h(t), t \geq t_1\}$ hits 1 and the surplus decreases of at least one unit in at most $\lfloor 1/\varepsilon' \rfloor cm/pq$ rounds with high probability for any $c > 2$ and for m large enough. Together with $\tau_0(G) = O(m)$, this show that τ_ε is polynomial in m (and therefore in the network size, as $m \geq n$) with high probability for all $0 < \varepsilon < 1$. ■

VI. Simulations

To validate our theoretical results, we ran our algorithmic model on the sequence of networks G_n defined in Section IV. Recall that G_n has n leaders and n followers, and observe that its unique stable matching has zero surplus. Figure 1 shows that the average[4] number of rounds to converge to the stable matching of G_n is exponential in n (triangles joined by dotted line), while on average all followers are recruited in a time that is linear in the network's size (diamonds joined by dashed line). Moreover the average time to reach a matching with surplus at most εm, for $\varepsilon = 0.2$, is clearly upper bounded by a polynomial in n (circles joined by continuous line).

Figure 2 compares, in logarithmic scale, the average time to converge to the stable matching of G_n (triangles joined by dotted line) and the average time to reach a matching with surplus at most εm), for $\varepsilon = 0.1$ (circles joined by continuous line). The curve corresponding to the approximate matching clearly shows a polynomial growth.

[4]Each point in each figure is obtained as the average of 100 simulations.

VII. Discussion

In this work we proposed a "human-like" algorithmic model of distributed team formation in networks of leaders and followers. The model is computationally simple, memoryless and suitable to model human behavior, but, despite its simplicity, it provides performance guarantees in the attainment of arbitrarily close approximations of a stable matching in any network. Notice that, in this model, there is no communication between leaders, who act only in response of their own status and of the status of their neighborhoods of followers. The only collaboration between leaders consists in a form of honesty deriving from the common goal of achieving stability.

Understanding how communication between leaders affects their coordination and characterizing the tradeoff between performance guarantees and model complexity is left for future work, as well as the validation of the model through laboratory experiments on human subjects

References

[1] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009.

[2] L. Coviello, M. Franceschetti, M.D. McCubbins, R. Paturi, and A. Vattani. Human matching behavior in social networks: an algorithmic perspective. *Under submission, http://acsweb.ucsd.edu/ lcoviell/,* 2011.

[3] J.E. Hopcroft and R.M. Karp. An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. *SIAM Journal on Computing*, 2:225, 1973.

[4] N.A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, 1996.