Community Detection in Opportunistic Networks using Memory-based Cognitive Heuristics

Matteo Mordacchini, Andrea Passarella, Marco Conti
IIT – CNR
Via G. Moruzzi 1, 56124 Pisa, Italy
Email: firstname.lastname@iit.cnr.it

***Abstract*—In a pervasive networking scenario like the *Cyber-Physical World* convergence, personal mobile devices must assist their users in analysing data available in both the physical and the virtual world, to help them discovering the features of the environment where they move. Mobile Social Networking applications are an example of Cyber-Physical applications, supporting users in their interactions in both worlds (e.g., during physical encounters, as well as during online interactions). It is very important, therefore, that nodes autonomously detect latent and dynamically changing social structures, resulting from common mobility patterns of users and physical co-location events. To this end, in this paper we propose a novel dynamic and decentralised community detection approach, whereby the nodes' behaviour is inspired by that of their human users, if they were exposed to the about physical encounters with other users, and would have to perform the same detection task. Specifically, we use cognitive heuristics, which are simple, low resource-demanding, yet effective, models of the human brain cognitive processes. At each node, the approach proposed in this paper, starting from the observed contact patterns with other nodes, estimates the strength of social relationships and detects social communities accordingly. An initial simulation evaluation shows that nodes are able to correctly identify the social communities that exist in their environment and to efficiently track change of membership due to modifications of the users' movement patterns.**

I. Introduction

The ubiquitous and pervasive presence in the physical world of devices that sense data from the physical environment and produce and exchange information among themselves is leading to a complex information scenario. This is known as the *Cyber-Physical World* (CPW) convergence scenario [1], and it is characterised by a mutual and continuous flow of information between the physical world and the cyber world. Information spread in one world influences the decision taken in the other, and vice-versa.

In this context, Mobile Social Networking applications (MSN) are emerging as one of the most interesting classes of pervasive applications. MSN applications rely on information about the users behaviour in order to deliver social oriented services. Among others, the problem of understanding the dynamic structure of social relationships induced by nodes' mobility is particularly important and challenging. These social relationships and structures are not necessarily the type of stable social relationships that users establish, for example, in Online Social Networks. Rather, they capture implicit structures (possibly unknown to the users themselves) that emerge from similarity of mobility patterns and regular physical co-location. Relevant examples of this class of social communities are groups of people that regularly commute together on a daily basis, people working in the same area of a city, etc. Recent results [2], [3] have shown that physical co-location is highly correlated with the existence of social relationships. Having, therefore, the possibility to dynamically detect this type of social structures is one of the enabler for realising pervasive applications in MSN environments. Opportunistic networks [4] are considered as one of the natural networking technology for supporting MSN applications, and we also consider such a networking environment for our purpose. In opportunistic networks, nodes detect physical contacts between each other in order to exchange information for example for the sake of disseminating data among interested users. Contact detection provided by opportunistic networks is the key enabler we use for our community detection algorithm.

Bringing forward the Cyber-Physical convergence view, in this paper we propose a novel dynamic community detection algorithm whereby nodes behave exactly as their human users would, if they were exposed to the same information about physical contacts and co-location events, and would have to group the other encountered people into a set of social communities. In other words, we see users' mobile devices as proxies of their human users in the cyber world. To this end, we base our algorithm on a set of very simple, yet effective rules known in the cognitive psychology field as *cognitive heuristics* [5]. In general, cognitive heuristics are functional models of how the human brain processes information it is exposed to, and takes decisions using minimal resources and based on partial and possibly noisy information. Algorithms derived from cognitive heuristics have been already successfully exploited to define data and knowledge dissemination schemes (e.g. [6], [7]) in opportunistic networks.

In this paper, we exploit cognitive heuristics to allow nodes to autonomously become aware of the structure of the social environment they are moving in. Such a mechanism should not produce a static picture of the social bindings of one node with the others. Rather, it should allow nodes to be responsive to the continuous changes in the context they are exploring, in order to have a time-varying representation of the existing social groups that form and possibly dissolve over time. To this end, we exploit the concept of memory activation, which is the cognitive model describing the process whereby the brain keeps track of other people we encounter, also taking into account the frequency with which we meet them. We couple this mechanism with other cognitive rules by which the brain groups in the same category people that produce similar memory activation patterns.

In the next section, we briefly describe some of the existing approaches for the community detection problem. In Sec. III we describe how cognitive heuristics can be used to derive a categorisation of contacts in the physical world into communities of socially related nodes. In Sec. IV we give an initial evaluation of the proposed solution in a simulated environment, while Sec. V concludes the paper.

II. Related Work

There exists a vast literature about the community detection problem [8]. However, many of the solutions used in complex networks analysis assume the exploitation of global knowledge and/or global communication capabilities (e.g. [9], [10], [11], [12]). Therefore, they are unlikely to fit a dynamic, real-time scenario as the one envisioned in this paper. On the other hand, other algorithms assume that each node in the network can rely only on its own local, limited view to build a picture of the social network. For example, some proposals make use of the contact duration between nodes to derive a representation of the social ties existing among peers in the network ([13], [14], [15]).

In this paper, we propose to face the problem from a different perspective. Specifically, we propose to exploit some of the cognitive processes that are used by the human brain to rapidly evaluate its social strength with another interacting person [16], and take decisions accordingly [17]. Cognitive heuristics have been used in opportunistic networks to successfully design effective data and knowledge dissemination schemes [6], [18], [19], [20], [7]. The main features of cognitive heuristics is that they usually require only minimal information to be effective. This is the key advantage of our approach with respect to other distributed solutions proposed in the literature. In fact, as shown in detail in the next section, our algorithm requires that each node keeps track of only the most recent contact times with other nodes and do not require the exchange of data between nodes upon contact, in order to compute their social affinity. Moreover, other approaches in the literature require that nodes exchange their list of most similar other peers (e.g. [13]) and/or have to continuously perform much complex (and computationally intensive) operations (like the operations over affinity matrices and the computation of eigenvalues) upon contact in order to derive the right social categorisations of other encountered nodes. Finally, the cognitive heuristic rules used in our proposal are designed to rapidly detect and adapt to changes in the environment, making them a suitable candidate for a complex task, like the community detection, in a dynamic scenario, like the CPW convergence scenario.

III. Community Detection in OppNets using Cognitive Heuristics

Recent results in the cognitive science field highlight the fact that patterns of encounters between humans can be used by the brain to predict the probability of future meetings with other people [16], [21].

Specifically, Pachur et al. [16] , observe that there are three main factors that influence the response of the memory when a meeting between persons occur. These are:

- *frequency* effects, i.e. the rate at which previous encounters happened in the past
- *recency* effects, i.e. the time passed from the last meeting
- *spacing* effects, i.e. how previous contacts where distributed over time, e.g. they all happened in a short time (*massed* contacts) or they where separated by each other by a long amount of time (*spaced* contacts)

A model that aims at describing the memory response to the patterns of encounter with other subjects should take into account all these effects. To this end, Pachur et al. [16] exploit the fact that humans maintain a *declarative* memory about the other people they met in the past. In this memory, for each social peer i, the brain keeps a record where it stores the times when previous encounters with i happened. When a new meeting with i happens, it produces an *activation* in memory. This activation measures the strength with which the brain is able to recall i, and it is derived from the pattern of occurrence of meetings with i. Activation about i can then be computed as:

$$A_i(t) = \ln\left[\sum_{j=1}^{k}(t-t_j)^{-d}\right] \quad (1)$$

where t is the time of the new meeting, $t_1 \ldots t_k$ are the times of the previous encounters with i that the brain is able to recall, and d is a decay parameter. Note that, using Equation 1, the relevance of a previous encounter at a time t_j decays exponentially over time. This makes the memory activation value more adaptive and responsive to changes in the environment, like a change in the relationship with another person, resulting in a different contact pattern. As a consequence of this fact, the brain does not need to keep track of all the previous meetings with other people.

In addition to this, Pachur et al. [16] also observe that the dynamics of physical contacts between humans seem to follow general regularities, where frequent contacts occur with only a very small number of other persons, while with most people contact is relatively rare. They argue that this is due to the fact that contact patterns are intrinsically related to the social ties that exist between individuals. Therefore, the social distance between two persons is highly associated with their probability of contact. This correlation has been also observed in other works in the ICT field [3], [2]. Pachur et al.[16] then suggest that the human brain can exploit, as a simple cognitive heuristic, the regularities observed in the contact patterns to estimate the degree of relationship that could exists with members of its social network. Specifically, they argue that the statistical structure of social contacts should be reflected in the ease with which information about these contacts is retrieved from memory. Upon contact with another person, the human memory is activated to a level that depends on the regularities experienced in other meetings with that person. Since those patterns are correlated with the social relationship that exists with that given individual, memory activation, computed as per Equation 1, can be used as an estimator of the importance of social ties.

Using all these observations, in this section, we propose an adaptive scheme that exploits the simple cognitive memory activation rule, defined in Equation 1, to let nodes in an opportunistic network infer the social tie between nodes that meet over time. In order to include a forgetting process in the memory of a mobile node, given a previous meeting of node n with a contact i happened at time t_j, in case $\frac{(t-t_j)^{-d}}{\sum_{j=1}^{k}(t-t_j)^{-d}} < \theta_{forget}$, $\theta_{forget} \in [0,1]$, we consider that the contribution given by the encounter occurred at t_j is not more relevant in the memory activation computation. Thus, a node n can *forget* about that previous meeting and deletes that data from its memory.

Starting from the memory activation defined as in Equation 1, with the forgetting mechanism described above, we show in the next section how other cognitive strategies can be used by each node to easily group the other peers into communities of socially-related entities.

A. Assigning nodes to communities

Using the memory activation value, each node estimates how strong is the social tie between itself and any other encountered peer. We now describe how activation, together with other cognitive models, can also be used to group encountered nodes into social communities.

The *similarity heuristic* [17] is a fast and frugal rule used by the human brain to evaluate the likelihood that a stimulus received from the environment belongs to one category rather than another. This cognitive heuristic assigns membership to categories by simply exploiting the degree to which the perceived stimulus is similar to others in that category. Specifically, given an observed data (e.g. the activation in memory of a social contact), the similarity heuristic compares that data against a function of the other data currently belonging to that category (in cognitive terms, a series of subjective hypotheses about the categories), and assigns the observed data to the most similar hypothesis.

In order to use the similarity heuristic in our algorithm, nodes must formulate hypotheses about possible categories to which other encountered nodes in the environment may belong to. Hypotheses should be formulated in such a way that the activation (computed as in Equation 1) of an encountered node can be compared against them, in order to assign each node to its most similar category.

To this end, note that an activation in memory reflects both the social strength with encountered nodes and the regularities in the contacts with the same nodes. As already pointed out in the previous section, results in the cognitive science field correlate the strength of social connections with the dynamics of physical contacts. Therefore, peers belonging to the same social category should show similar contact patterns, and, as a consequence, they should induce similar activation values in memory upon meeting.

In a similar way as proposed in the cognitive sciences by Raviv et al. [22], we assume that the observation of a series of similar values produces a trace in the memory of a node that gives a sort of digest of all these previously observed similar stimuli (i.e. memory activations). Since nodes with similar social behaviour produce similar memory activation values, these traces can be regarded as the basis for formulating the hypotheses needed by the similarity heuristic. Thus, they can be used for assigning newly observed memory activations to social categories.

More in details, the procedure used for the classification of encountered nodes using the memory activation process described in Equation 1 is described in Algorithm 1.

Algorithm 1 Classification of a node i upon contact at time t

1: Let $a = A_i(t)$
2: Let S be the set of social category exemplars
3: **if** $S = \emptyset$ **or** $d(a, s) > \theta_{split}\ \forall s \in S$ **then**
4: $\quad S = S \cup \{a\}$
5: $\quad$ Return IndexOf(a)
6: **else**
7: $\quad$ Let $s^* = \min_{s \in S} d(a, s)$
8: $\quad s^* = \alpha s^* + (1 - \alpha)a$
9: $\quad$ Return IndexOf(s^*)
10: **end if**

Using the description given in the algorithm, we assume that each node has a set S of social categories. When the node meets another peer i at a time t, it computes the memory activation of i, $A_i(t)$ (line 1 of Alg. 1). In case the set S is empty (line 3), the memory activation value is stored as the first of the social category exemplars. Otherwise, the exemplar with the lowest distance from the actual activation is chosen as the category hypothesis that best fits the observed activation value (line 7). In this paper, we consider that the distance function is simply the absolute value of the difference between the actual activation and the category representatives, i.e. $d(a, s) = |a - s|$, $\forall s \in S$, where a is the observed memory activation. Moreover, as suggested by Raviv et al. [22], we assume that the chosen category representative s^*, in response to the presentation of a new, related stimulus is updated in a way that linearly combines its current value with the observed activation. Specifically, upon a new activation a related to the category represented by s^*, the node updates its reference value for this category using a *smoothed average* (line 8):

$$s^* = \alpha s^* + (1 - \alpha)a \tag{2}$$

where $\alpha \in [0, 1]$ Moreover, in case the actual activation a has a distance from each of the exemplars stored in S greater than a threshold θ_{split} (line 3), a is considered to represent a new, previously unobserved category. Therefore, a new category is added to S, using a as the seed of the value that represents it. This behaviour is similar to what is described in the cognitive sciences by Anderson [23]. According to this work, stimuli are clustered by the brain on the basis of their similarity. When a new stimulus differs too much from the members of the existing clusters, a new group is formed.

The procedure described in Alg. 1 returns an index that uniquely identifies the chosen category representative within S, and it can then be used by the node to decide the social category the encountered peer belongs to.

Note that by using Alg. 1, the information that two nodes have to exchange upon meeting is made up of just their own IDs. Knowing which is the other interacting party, each node is able to retrieve the data about previous encounters and, using the memory activation value, can decide which is the social category to which the other node belongs to.

IV. Simulation Evaluation

In this section, we report initial, yet significant results about the algorithm performance obtained by simulation using the HCMM mobility model [24]. HCMM is a mobility model that integrates temporal, social and spatial notions in order to obtain an accurate representation of real user movements. In order to achieve this goal, its design is inspired by results in the sociology and complex networks literature. One of its main features is the ability to reproduce statistical properties of real user movement patterns, such as inter-contact times and contact durations. In HCMM, the simulation space is divided in cells. Each cell could host a group of nodes, that represents a social community. Different groups can be connected by special nodes, called "travellers", that move across different communities, thus bridging them. Each social group is initially assigned to a cell (its *home-cell*) avoiding that two groups are physically adjacent (no edge contacts between groups) or in the same cell. This allows us to eliminate physical shortcuts between groups. Therefore, nodes can communicate only due to social mobility and not due to random colocation. Table I shows the values of the main simulation parameters used to obtained the results reported in this section.

TABLE I. MAIN SIMULATION PARAMETERS

Simulation Parameters	
Simul. Area	1000m^2
Grid	4x4
Numb. of Communities	3
Numb. of Nodes	45 (15 per comm.)
Node speed	unif. in $[1, 1.86]m/s$
Transm. range	20m
Simulation time	50000s
α	0.9
θ_{split}	1.5
d	0.5

In the first set of experiments, we use a configuration with 3 distinct communities of 15 nodes each. In this scenario, a community has a traveller toward each of the other groups. Each traveller is allowed to move from its home community to just one other group. In these experiments, we force the system to make each single node recognise no more than two distinct communities. This lead to classify the other nodes into *friends* and *familiar strangers*, which is a typical classification performed by other reference community detection approaches (see, e.g., [13], [14]). Each node considers as friends the peers with a memory activation associated to the hypothesis with the highest value.

Fig. 1 reports the results of an experiment where we want to investigate the ability of the proposed solution to let nodes other than travellers (i.e., nodes "internal" to their social community) to recognise the members of their own community, i.e. their *friends*. The curves in Fig. 1 are obtained with three different mean sojourn times of the travellers in their home and external communities, respectively. This curves correspond to the following scenarios:

- Evenly distributed (ED) sojourn times: travellers spend in their home cell and in the external community almost the same amount of time;
- Home community shorter (HS): travellers spend more time in the external community they visit than in their home community;
- Home community longer (HL): travellers spend more time in their home community than that spent in the external ones.

The values of the mean sojourn times of these three scenarios are reported in Table II.

TABLE II. MEAN SOJOURN TIMES

	Home Comm.	*Ext. Comm.*
ED	554.60s	549.93s
HS	479.82s	1704.07s
HL	619.99s	175.49s

In order to evaluate the similarity between the community detected by each node and the community structure with which we have configured HCMM, we use the Jaccard index. For each node n, it is defined as:

$$\sigma(n) = \frac{|C_f(n) \bigcap T_f(n)|}{|C_f(n) \bigcup T_f(n)|}$$

where $C_f(n)$ is the community of friends detected by n, while $T_f(n)$ is its true set of peers belonging to the node's home community. In principle, the target of a community detection algorithm should be to achieve a value of 1 for this index. In practice, this might not be the case. Remember that in our scenario travellers spend time both inside and outside their home community. According to our definition, nodes belong to the same community if they spend sufficient time together and meet frequently enough. Therefore, nodes internal to a community may perceive travellers either as being part of the community or not, depending on how much time they spend together. In other words, we may expect that our algorithm *correctly* achieves similarity values lower than 1, in cases where the HCMM association of a traveller to its home community does not reflect the amount of time it spends inside it. Results in Fig. 1 are the average over time for all the internal nodes of all the three communities used in the experiments. In all these cases, $\theta_{forget} = 0.01$. In all the presented results, detection of travellers is the discriminating point in the performance figure. In fact, the best results are achieved in the HL case, when travellers spend most of their time inside their home community, rather than outside it. In the HS case, the Jaccard similarity stabilises around a value given by the fact that all the other internal nodes are seen as friends, while all the home community travellers are regarded as familiar strangers. This behaviour can be judged positively, since travellers usually stay a long time away from their home community, thus having a more strict relation with the nodes of the external group. In the ED case, note that travellers stay in both the communities they visit an almost identical amount of time. The effect of this behaviour is to lead an internal node to consider travellers (either of its own home community or coming from outside) alternatively as familiar strangers or friends, in case they have just come back from another group, or they have already spent a sufficient amount of time in the community, respectively.

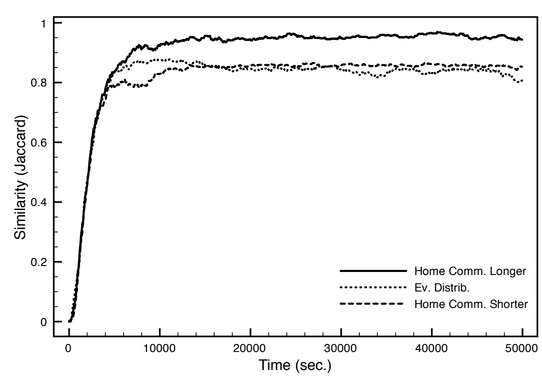

Fig. 1. Familiar nodes recognition for a internal peer, with different traveller sojourn times

In Fig. 2, we analyse the impact of forgetting on the algorithm performance. Since the most complex scenario for an internal node is the ED one, we focus our attention on this configuration of the system. As expected, higher values of θ_{forget} degrade the performance of the proposed solution. Anyway, note that making the θ_{forget} value three times higher than that used in Fig. 1 does not cause great variations in the Jaccard similarity values. On the other hand, increasing θ_{forget} up to 0.1 (ten times the value of the previous experiments) produces a strong degradation of the performance. In fact, in this case, too few information is retained by a node for each of its social contacts. As a consequence, other internal nodes that are not seen for a while can rapidly become *familiar strangers*. Note that higher similarities around the beginning of the simulation are due to the fact that in that phase the relative weight of each contact with respect to the whole activation value is not below the threshold, because there are still few contacts, and each of them has a significant relative weight. As contacts accumulate over time, their relative weight becomes lower and lower, and they are progressively discarded, resulting in a decrease of the similarity value.

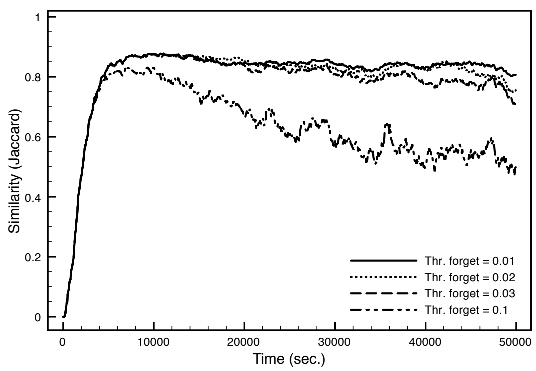

Fig. 2. Variation of the Jaccard similarity metric for the *friends* set of an internal node, with different θ_{forget} values

Fig. 3 shows the average number of previous contacts that are maintained in the memory of an internal node, using the same forgetting parameters of the previous figure. We can observe that the number of stored previous encounters decreases linearly with respect to the value of θ_{forget}. Thus, the number of contacts maintained for $\theta_{forget} = 0.02$ is half the number used with $\theta_{forget} = 0.01$, and so on. As seen in Fig. 2, we do not have a great variation in the similarity metrics using $\theta_{forget} = 0.03$ instead of $\theta_{forget} = 0.01$, while, at the same time, we use a third of the information about previous meetings. Using $\theta_{forget} = 0.1$, the number of previous encounters (5.5 on average) becomes to low to achieve a good performance.

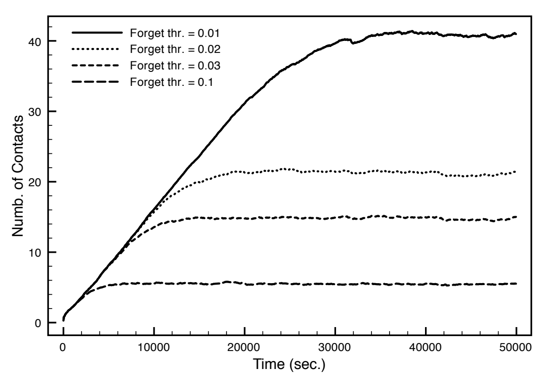

Fig. 3. Number of previous encounters maintained by an internal node, with different θ_{forget} values

While in the previous experiments we observed the behaviour of internal nodes, we now move our attention to the performance of the travellers. Since travellers move from one community to another, one of the main results we want to achieve is to let them autonomously recognise that the environment around them changes and, as a consequence, modify their perception of friends and familiar strangers accordingly.

Fig. 4 and 5 show how the composition of the *friends* set of a tagged traveller changes over time. Specifically, the two figures show the number of nodes of the tagged traveller's home and external communities that are regarded as friends by the traveller. The results are obtained with $\theta_{forget} = 0.01$ in Fig. 4, and with $\theta_{forget} = 0.1$ in Fig. 5. For the sake of readability, the figures show only a part of the data of the whole experiment (from sec. 42000 to sec. 48000). Figures show a series of pairs of vertical bars. In each pair, the first bar denotes the time instants when the traveller exits from the previous community, and the second the time when it enters the next one.

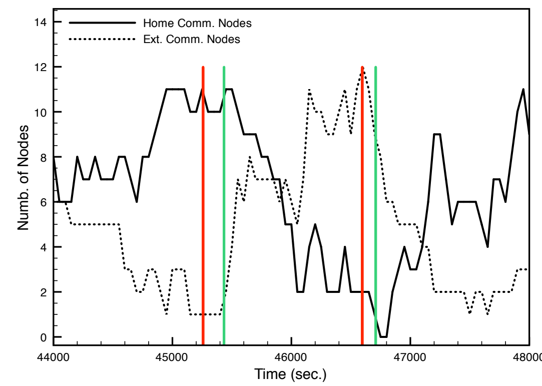

Fig. 4. Variation in the *friends* set composition for a tagged traveller, $\theta_{forget} = 0.01$

In both the figures, it is possible to observe that the number of "friends" of the traveller coming from its home community decreases, as a consequence of a change of community done by the tagged traveller. On the other hand, the number of friends belonging to the external community increases accordingly. Typically, when the number of friends of one community reaches its peak, the number of nodes of the other community in the friend set achieves its minimum. This point is even more noticeable in Fig. 5, highlighting the fact that higher forgetting could be more beneficial for a traveller, since it becomes even more responsive to changes in its surroundings.

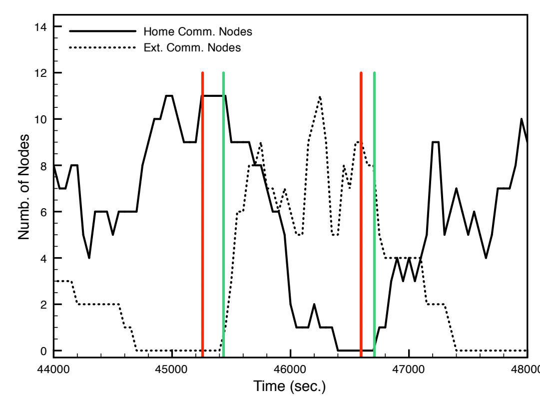

Fig. 5. Variation in the *friends* set composition for a tagged traveller, $\theta_{forget} = 0.1$

In order to have a deeper understanding of the ability of travellers to dynamically adapt their vision to the structure of the local social context, in Fig. 6 we present the results obtained in a more complex scenario. In this case, we still have three distinct communities. However, just one of them has a traveller. This traveller bridges its home community with all the remaining groups. Moreover, we do not impose to classify the other nodes in just two sets, i.e. *friends* and *familiar strangers*. We give the system the full flexibility to autonomously decide the proper number of social categories needed to classify the other encountered peers (i.e., we use exactly Algorithm 1, where new categories are generated according to the split parameter, without imposing a maximum number of communities). The results presented in Fig. 6 are obtained with exactly the same parameters used for all the previous experiments and with $\theta_{forget} = 0.1$.

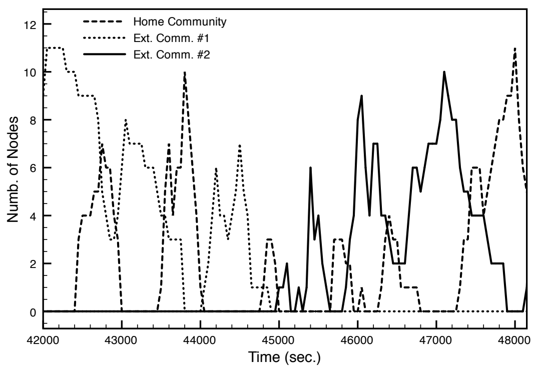

Fig. 6. Variation in the composition of the social set with the highest activation

For the same reasons used for the two previous figures, we limit the view of the results to a portion of all the data, from sec. 42000 to sec. 48500. During the experiment, the traveller created *three* different social categories, that is consistent with the moving patterns of the traveller among three different communities. In Fig. 6 we show the number of nodes of each of the three existing communities inside the set associated with the highest activation values, i.e. the nodes considered to be more socially connected with the traveller. Again, the traveller shows its ability to adapt its view of its social ties. Analogously to the previous results, a peak in the number of most socially related nodes coming from a community corresponds a minimum in the number of nodes in this set belonging to the other groups.

V. Conclusion

In this paper, we propose a community detection scheme for nodes in an opportunistic networking environment that is inspired by memory-based human cognitive solutions. A memory activation-based method is used by nodes to evaluate the social strength existing with another encountered node and assign it to a proper social community. We give an initial evaluation of the proposed approach in a simulated environment. Our method shows to be able to detect the right communities when nodes are restricted to move inside just one single group, letting them also discover which are the nodes that act as connectors with other communities (i.e. travellers). Moreover, when nodes, like travellers, experience sudden changes in the environment in which they move, they are able to change the composition of the social groups they perceive, autonomously adapting their view to the new context.

Acknowledgments

This work is funded by the EC under the FET-AWARENESS RECOGNITION (FP7-257756) and FIRE EINS (FP7-288021) projects and by the EIT ICT Labs Emergent Social Mobility Project.

References

[1] M. Conti et al., “Looking ahead in pervasive computing: Challenges and opportunities in the era of cyber–physical convergence,” *Pervasive and Mobile Computing*, vol. 8, no. 1, pp. 2–21, 2012.

[2] D. Wang, D. Pedreschi, C. Song, F. Giannotti, and A.-L. Barabási, “Human mobility, social ties, and link prediction,” in *KDD*, C. Apté, J. Ghosh, and P. Smyth, Eds. ACM, 2011, pp. 1100–1108.

[3] S. Scellato, A. Noulas, and C. Mascolo, “Exploiting place features in link prediction on location-based social networks,” in *ACM SIGKDD*, 2011.

[4] L. Pelusi, A. Passarella, and M. Conti, “Opportunistic networking: data forwarding in disconnected mobile ad hoc networks,” *Communications Magazine, IEEE*, vol. 44, no. 11, pp. 134–141, 2006.

[5] G. Gigerenzer and P. M. Todd, *Simple Heuristics that Make Us Smart*. New York: Oxford University Press, 1999.

[6] M. Conti, M. Mordacchini, and A. Passarella, “Design and performance evaluation of data dissemination systems for opportunistic networks based on cognitive heuristics,” *ACM Trans. Auton. Adapt. Syst.*, vol. 8, no. 3, pp. 12:1–12:32, 2013.

[7] M. Mordacchini, A. Passarella, M. J. Chorley, G. B. Colombo, and V. Tanasescu, “Making mobile users' devices aware of the surrounding physical environment: an approach based on cognitive heuristics,” in *Proc. of SASO 2013*. IEEE, 2013, pp. 199–208.

[8] S. Fortunato, “Community detection in graphs,” *Physics Reports*, vol. 486, no. 3, pp. 75–174, 2010.

[9] U. N. Raghavan, R. Albert, and S. Kumara, “Near linear time algorithm to detect community structures in large-scale networks,” *Physical Review E*, vol. 76, no. 3, p. 036106, 2007.

[10] A. Lancichinetti, F. Radicchi, J. J. Ramasco, and S. Fortunato, “Finding statistically significant communities in networks,” *PloS one*, vol. 6, no. 4, p. e18961, 2011.

[11] M. Giatsoglou and A. Vakali, “Capturing social data evolution using graph clustering,” *Internet Computing, IEEE*, vol. 17, no. 1, pp. 74–79, 2013.

[12] P. Hui, J. Crowcroft, and E. Yoneki, “Bubble rap: Social-based forwarding in delay-tolerant networks,” *Mobile Computing, IEEE Transactions on*, vol. 10, no. 11, pp. 1576–1589, 2011.

[13] T. Hossmann, T. Spyropoulos, and F. Legendre, “Know Thy Neighbor: Towards Optimal Mapping of Contacts to Social Graphs for DTN Routing,” *2010 Proceedings IEEE INFOCOM*, pp. 1–9, 2010.

[14] P. Hui, E. Yoneki, S. Y. Chan, and J. Crowcroft, “Distributed community detection in delay tolerant networks,” in *Proceedings of 2nd ACM/IEEE international workshop on Mobility in the evolving internet architecture*, ser. MobiArch '07. New York, NY, USA: ACM, 2007, pp. 7:1–7:8.

[15] E. Borgia, M. Conti, and A. Passarella, “Autonomic detection of dynamic social communities in opportunistic networks,” in *Ad Hoc Networking Workshop (Med-Hoc-Net), 2011 The 10th IFIP Annual Mediterranean*. IEEE, 2011, pp. 142–149.

[16] T. Pachur, L. J. Schooler, and J. R. Stevens, “When Will We Meet Again? Regularities of Social Connectivity and Their Reflections in Memory and Decision Making,” in *Simple Heuristics in a Social World*, R. Hertwig, U. Hoffrage, and ABC Research Group, Eds. Oxford University Press, Nov. 2012.

[17] D. Read and Y. Grushka-Cockayne, “The Similarity Heuristic,” *Journal of Behavioral Decision Making*, no. 24, pp. 23–46, 2011.

[18] R. Bruno, M. Conti, M. Mordacchini, and A. Passarella, “An analytical model for content dissemination in opportunistic networks using cognitive heuristics,” in *Proc. of MSWIM 2012*. ACM, 2012, pp. 61–68.

[19] M. Conti, M. Mordacchini, A. Passarella, and L. Rozanova, “A semantic-based algorithm for data dissemination in opportunistic networks,” in *Proc. of IWSOS 2013*. Springer, 2013.

[20] M. Mordacchini, L. Valerio, M. Conti, and A. Passarella, “A cognitive-based solution for semantic knowledge and content dissemination in opportunistic networks,” in *Proc. of AOC 2013*. IEEE, 2013.

[21] M. M. Krasnow, A. W. Delton, J. Tooby, and L. Cosmides, “Meeting now suggests we will meet again: Implications for debates on the evolution of cooperation,” *Scientific reports*, vol. 3, 2013.

[22] O. Raviv, M. Ahissar, and Y. Loewenstein, “How recent history affects perception: the normative approach and its heuristic approximation.” *PLoS computational biology*, vol. 8, no. 10, pp. 1–10, Jan. 2012.

[23] J. R. Anderson, “The adaptive nature of human categorization.” *Psychological Review*, vol. 98, no. 3, pp. 409–429, 1991.

[24] C. Boldrini and A. Passarella, “Hcmm: Modelling spatial and temporal properties of human mobility driven by users' social relationships,” *Comput. Commun.*, vol. 33, pp. 1056–1074, June 2010.