

A cognitive-based solution for semantic knowledge and content dissemination in opportunistic networks

Matteo Mordacchini*, Lorenzo Valerio*, Marco Conti*, Andrea Passarella*

*IIT-CNR, Pisa, Italy

Email: {matteo.mordacchini,lorenzo.valerio,marco.conti,andrea.passarella}@iit.cnr.it

Abstract—Opportunistic networking is one of the key paradigms to support direct communication between devices in a mobile scenario. In this context, the high volatility and dynamicity of information and the fact that mobile nodes have to make decisions in condition of partial or incomplete knowledge, makes the development of effective and efficient data dissemination schemes very challenging. In this paper we present algorithms based on well-established models in cognitive sciences, in order to disseminate both data items, and semantic information associated with them. In our approach, semantic information represents both meta-data associated to data items (e.g., tags associated to them), and meta-data describing the interests of the users (e.g., topics for which they would like to receive data items). Our solution exploits dissemination of semantic data about the users’ interests to guide the dissemination of the corresponding data items. Both dissemination processes are based on models coming from the cognitive sciences field, named *cognitive heuristics*, which describe how humans organise information in their memory and exchange it during interactions based on partial and incomplete information. We exploit a model describing how semantic data can be organised in each node in a semantic network, based on how humans organise information in their memory. Then, we define algorithms based on *cognitive heuristics* to disseminate both semantic data and data items between nodes upon encounters. Finally, we provide initial performance results about the diffusion of interests among users, and the corresponding diffusion of data items.

Keywords-opportunistic networks; content diffusion; cognitive heuristics, semantic knowledge;

I. INTRODUCTION

In the Cyber-Physical World (CPW) convergence scenario, the interaction between devices and their users contributes to the increasingly complexity of the information environment. This is characterised by information flows that constantly cut across the cyber and physical world. As discussed in detail in [1], applications exploit information coming from the physical world (e.g., from sensors) to optimise their operations in the cyber world. The status of applications in the cyber world can determine actions in the physical world (such as configurations of actors in sensor actor networks, or the users behaviour in social gaming or other social-oriented applications). In this context, mobile devices play an important role because they are the actual representation of their users in the cyber world. In other terms, mobile devices in the role of proxy of their human

counterparts in the cyber world are in charge of discern and collect the important information (for their users) that circulate in the environment. We can consider data dissemination in opportunistic networks as an important specific example of this general concept. Indeed, mobile nodes use direct communication and the store carry and forward paradigm to recognise and disseminate relevant information, and the key challenge is to assess what information is really relevant for the users and, therefore, what data needs to be exchanged during contacts to satisfy the users’ interests.

In this paper we jointly cope with two separate but tightly connected dissemination problems: dissemination of semantic information and of data items. Namely, we consider a scenario where semantic information represents both meta-data associated to data items (e.g., tags associated to them), and meta-data describing the interests of the users (e.g. concepts and ideas for which they would like to receive relevant data items). Our idea is to exploit the semantic data dissemination to drive the spreading of the corresponding data items, in a way that resembles how conversations between humans enable spreading of ideas (i.e., semantic information), which generates interests for specific types of content, and ultimately determine content that people access. Reasonably, two people who meet start talking about a common interest – that we call key-concept – continue with other semantically correlated ones that spontaneously emerge along the conversation. Finally, they can also exchange data items based on their interests (which, in this general example could be, for example, songs or pictures they keep in their smartphones). The amount of information exchanged during a conversation depends on: i) how many concepts related to some topic each person knows, ii) how long the conversation takes place, and iii) how fast the two people are able to retrieve from their memory the next proper concepts on which continuing the conversation.

Our solution aims at mapping this human process into a set of algorithms for exchange of semantic data and data items between users’ mobile devices, exploiting models of the mental processes that drive these exchanges between humans. In particular we address the above three aspects – that we can reformulate as memory representation, concept retrieval, and content selection – through well-established models belonging to the cognitive science field.

In this paper we assume a scenario where data items (contents) are equipped with a proper semantic description (concepts/tags). In order to face the memory representation aspect, at the beginning each mobile node builds a local semantic representation of its own contents through a semantic directed weighted graph where, vertices represent the semantic concepts associated to data items, and the edges represent the semantic relationships between concepts. This solution is inspired by the associative network models of human memory (AN) coming from the cognitive psychology field. Upon meeting, under the condition of having common interests, two devices start to exchange their knowledge, selecting concepts by navigating their semantic network, starting from concepts they have in common, and according to an edge ranking algorithm to drive the navigation process. Actually, this is an implementation of the Fluency heuristic (FH), a decision making strategy borrowed from cognitive psychology, that addresses the concept retrieval aspect presented above. Finally, mobile nodes after having enriched their semantic graph with new concepts taken from the encountered nodes, select and share the data items starting from those that have the maximum overlap between the associated tags and the concepts in the semantic networks just exchanged with the encountered node. In cognitive terms this refers to the Tallying heuristic (TH), another cognitive decision strategy used by human brain. A more detailed description about the cognitive foundations of our approach can be found in Sec. II.

In this paper we show that starting from a semantic representation of data in terms of a semantic network according to AN memory models and exploiting FH and TH, nodes are able to successfully disseminate knowledge and the corresponding contents in a very selective way. Even without an explicit subscription to specific channels of interest, nodes interested in some topics can acquire new semantically related concepts and the corresponding new correlated data items. In scenarios where node interests are strongly clustered around some main topics, this is a quite remarkable feature, because with our approach nodes receive only those data items related to their real interests.

II. COGNITIVE PSYCHOLOGY BACKGROUND

In this section we give a short and functional description of the cognitive building blocks we exploited to address the aspects of memory representation, concept retrieval and content selection in order to devise our dissemination approach. According to [10], the most prevalent models of human memory are associative network models, which focus on the patterns and strength of associative linkages among concepts in memory. In associative network models, concepts are represented by nodes that are interconnected by pathways that vary in strength, reflecting the degree of association between each pair of concepts. Two related concepts are connected by links, whose weight represents the activation

level, i.e. how likely it is that one of them is "accessed" after the other has been accessed. Exploration of the network (i.e., reasoning), can be done either sequentially or in parallel, which broadly correspond to depth-first and breadth-first navigation in standard graph theory, respectively [10]. Here, we concentrate our focus on the sequential model (SS) leaving the parallel one for future investigations. A sequential search on AN starts with an activated concept (key-concept) and proceeds node by node along the pathway that connects them. When a node has more than one outgoing path, the one with the strongest activation is selected. If a "dead end" is reached, the search is reinitiated. The SS model relies on the idea of "strongest activation" that in cognitive terms recall a well assessed mental model, the *Fluency Heuristic* that lies in the context of *cognitive heuristics*. Briefly, cognitive heuristics are functional model of mental processes that human brain uses to quickly take actions even in presence of incomplete knowledge. Differently from neural networks, cognitive heuristics do not aim at reproducing the physiology of the brain's processes, but model their functionality. Due to their effectiveness in quickly finding good solutions to infeasible or impractical problems, they have been successfully applied in several fields, such as outcomes of political elections, forecasting purchase, financial decision making, and results of sport events [6].

FH is an inference strategy that can be applied when someone has to choose among two or more alternatives. Among the alternatives that are *recognized* the one perceived as recognized faster (in our case, the link with the strongest activation) is considered to have an higher value w.r.t. the selected criterion. Being recognized means that a given information has been found in the environment a sufficient number of times to let the brain being familiar with it [11].

Finally, content selection is made by exploiting the *Tallying Heuristic* (TH). TH uses m out of a total of M cues, in order to discriminate among alternatives w.r.t. a criterion. For each alternative, it simply counts the number of favourable cues without giving any special weight to any of them. The alternative with the highest number of positive cues is then selected [12]. In our case, cues are tags associated to data items, and a tag is favourable if it matches one of the concepts accessed in their semantic networks by two nodes during an encounter (i.e., if the nodes have "spoken" about a concept that matches the tag).

III. RELATED DATA DISSEMINATION APPROACHES

Many works in literature cope with the problem of the content distribution in pure OppNets. The first attempt was in the PodNet Project [7]. It was based on a cooperative exchange of items driven by the estimated popularity of the channels the items belong to. More advanced solutions, e.g. [8], make use of social information about users. One of the most prominent is ContentPlace[2] in which nodes fill their caches in order to maximise both a local utility

that refers to the interests of the local user, and a global utility that identifies the relevance of data items for the other communities the node is in touch with. An extensive survey about content diffusion in OppNets can be found in [4].

In [5] and [9] the problem is addressed from a completely different perspective. The idea is to mimic in functional terms the human decision making process by exploiting the *Recognition Heuristic* and the *Take the Best Heuristic*, two of the several cognitive heuristic models present in the psychological literature[6]. These works prove the suitability and effectiveness of these heuristics in problems, like data dissemination in OppNets, where every node has only a partial knowledge about its environment.

This paper also exploits this approach. However, here for the first time we exploit both cognitive models for the representation of concepts and ideas, and cognitive models for selecting which information is exchanged by nodes upon encounter. We thus use a much more complete set of cognitive models *together*, achieving a more complete translation of the mental processes used by humans to exchange knowledge and data. Note that, with respect to the popular topic based approaches described above, the use of semantic networks to represent concepts that are then linked to the data tags can be seen as a generalisation of the use of topics to categorise data, much closer to the way humans categorise information in their memory.

IV. COGNITIVE HEURISTICS FOR DATA AND CONTENT DISTRIBUTION IN OPPNETS

In the following, we show how concepts coming from the cognitive science area can be exploited to design schemes for disseminating both the semantic knowledge and its related content items in an opportunistic network. Firstly, we show how a cognitive model of memory is used to represent the semantic concepts owned by a user. Then, we describe how simple cognitive heuristics like *fluency* and *tallying* can be used to select the most relevant information to pass from one node to the other upon contact.

A. Semantic Network Creation

In Sec. II, we stated that the semantic concepts of a node and the relations between them are described using an AN model. More formally, each user’s semantic network is defined as a *dynamic weighted graph* $G = \{V, E, f(e, t)\} : t \in T$, where t is the time, V is the set of vertices (i.e. semantic concepts) in the graph and E is the set of edges (i.e. the connections between semantic concepts). In an AN, links are associated with a strength value that reflects the degree of association between each pair of concepts in memory. In G , this value is computed by a *remembering function* $f(e, t)$. Being t_0 the creation time of the edge e_{ij} , we have the initial setting $f(e_{ij}, t_0) = 1$. At any other time instant $t > t_0$, $f(e_{ij}, t)$ decreases exponentially depending on the length of the interval $[t', t]$, where t' is the last time e_{ij}

was “activated” (i.e. used in interactions with other peers, and thus “refreshed” in memory). Thus, we have:

$$f(e_{ij}, t) = e^{-\beta_{ij}(t-t')} \quad (1)$$

where β_{ij} is the “speed of forgetting”, taken in accordance with the experimental curve obtained in [13]. In real human memory, rarely accessed information is more likely forgotten than frequently used one, hence, $\beta_{ij} = \frac{\gamma}{p_{ij}^t}$, where γ is a speed regulator parameter and p_{ij}^t is the “popularity” of edge e_{ij} , i.e. the number of times it was used during the encounters of a specific user with other peers until time t . If the value of $f(e_{ij}, t)$ falls below a *remember threshold* f_{min} , e_{ij} is removed from G . In order to create the initial graph G at $t_0 = 0$, we assume that the data available locally at each user is associated with a set of tags, as in real social networks like Flickr, Twitter, Instagram, etc. As shown in the example of Fig. 1, for each data item, its tags are linked together in order to form a *clique*. Then, each set of vertices carrying the same label (i.e. they were created from tags having the same name) is merged together, forming a single vertex. In the example, the two cliques are merged using the common vertex “lake” as a pivot.




Figure 1. Creation process of a user Semantic Network.

B. Semantic knowledge dissemination

We now describe how a node, when meeting another peer, is able to retrieve from its memory the most relevant semantic information to be exchanged. The main role in this process is played by the *fluency* heuristic. In this phase, we call the node’s semantic network $G = (V, E, f(e, t))$ the *donor network*, while the other peer’s semantic network $G' = (V', E', f'(e', t))$ is termed the *recipient network*. The two nodes swap roles (donor and recipient) to realise a bidirectional exchange of information. The subgraph $C = (\bar{V}, \bar{E}, \bar{f}(\bar{e}, t))$ selected from the donor network to be passed to the recipient one is called the *contributed network*. In the following, we assume that resource consumption constraints limit the number of exchangable concepts to a value *tag_limit*. The subsequent description follows the specification given in Alg. 1 and 2. Like in a real human communication, we assume that the dialogue starts from a set of common semantic concepts, i.e. a set of *key vertices* $K = \{v_k | v_k \in V \cap V'\}$ (line 3 of Alg. 1). The relevance of a vertex is increased every time it is included in K during information exchanges with other peers. In our system, the *key vertices* relevance is augmented by increasing the

popularities p_{ij}^t of all the edges attached to them (lines 5–7). In order to compute the contributed network, vertices and edges are selected from the donor network by first ordering the key vertices by their relevance in memory. This is done by summing up the weights of their incoming edges (line 8). Taking the sorted key vertices one at a time, edges and vertices are visited and passed from the donor network to the contributed one using Alg. 2, based on *fluency*, here applied to evaluate whether to follow an edge e_{ij} or not. *Fluency* favours *recognized* edges (i.e. the one seen more than a given amount of times, see Sec.II) against unrecognized one, assuming the former are more relevant than the latter. Hence, we start by excluding all unrecognized edges, i.e. those whose popularity is below a recognition threshold θ_{rec} (line 5 of Alg. 2). The subsequent discrimination made by

Algorithm 1 Contributed Network computation at time t^*

- 1: Let $G = (V, E, f(e, t))$ be the donor network;
 - 2: Let $C = (\bar{V}, \bar{E}, \bar{f}(\bar{e}, t))$ be the contributed network;
 - 3: Let K be the set of *key vertices*, $K \subseteq V$
 - 4: **for** each $v_i \in K$ **do**
 - 5: **for** each $e_{ij} \in E$ **do**
 - 6: increase popularity of e_{ij}
 - 7: **end for**
 - 8: Let $rel_{ij} = \sum_{e_{ij} \in E} f(e_{ij}, t^*)$
 - 9: **end for**
 - 10: **for** each $v_i \in K$ taken in desc. order w.r.t. rel_{ij} **do**
 - 11: $C = visit(v_i, 1, t^* - t)$
 - 12: **end for**
 - 13: Send C to the other node
-

fluency is based on the perceived speed of retrieval from memory. In order to replicate this fact in our system we assume that: (i) the highest the remembering value of an edge, the most relevant the edge is; (ii) the relevance of an edge decreases as long as we get farther from a key vertex; (iii) The longer the contact time, the longer the time available to retrieve a concept from memory after another one close by in the network has been retrieved, and thus the higher the relevance of the connecting edge. We take into account all these observations by computing, for each outgoing edge e_{ij} of a vertex v_i , a *retrieval weight* quantity

$$w(e_{ij}, n, t^* - t) = f(e_{ij}, t^*) \frac{1 - e^{-\tau(t^* - t)}}{n} \quad (2)$$

where t, t^* are the times at which the contact starts and ends, n is the the number of hops in the shortest path to the nearest key vertex and τ is a “speed” factor that regulates the dependancy of this value on the communication duration ($t^* - t$). We sort the edges w.r.t. their retrieval weight value. Taking them one at a time in descending order, we include the selected edge in the contributed network and continue the donor network exploration from this connection (lines 6–14). All the edges whose retrieval value is below

a threshold w_{min} are not considered (line 7). Note that the remembering value (eq. 1) in memory of selected edges is set to 1, since inclusion in the exchanged data corresponds to an “activation” in memory of those connections. When

Algorithm 2 Function $visit(v_i, n, t^* - t)$

- 1: Let $G = (V, E, f(e, t))$ be the donor network;
 - 2: Let $C = (\bar{V}, \bar{E}, \bar{f}(\bar{e}, t))$ be the contributed network;
 - 3: **if** $|\bar{V}| < tag_limit$ **then**
 - 4: $\bar{V} \cup = v_i$
 - 5: Let $R = \{e_{ij} \in E | p_{ij}^{t^*} \geq \theta_{rec}\}$
 - 6: **for** each $e_{ij} \in R$ in desc. order w.r.t. $w(e_{ij}, n, t^* - t)$ **do**
 - 7: **if** $w(e_{ij}, n, t^* - t) \geq w_{min}$ **then**
 - 8: $\bar{E} \cup = e_{ij}$
 - 9: $f(e_{ij}, t^*) = 1$
 - 10: $\bar{f}(\bar{e}, t^*) = 1$
 - 11: $C \cup = visit(v_j, n + 1, t^* - t)$
 - 12: **end if**
 - 13: **end for**
 - 14: **end if**
 - 15: Return C
-

$|\bar{V}| = tag_limit$ and/or no other paths (i.e. edges) can be selected from the donor network, the contributed network computation ends and the resulting graph is passed to the recipient node. This peer merges the contributed network to the recipient one by simply adding all the missing vertices and edges. This process correspond to an enrichment of the semantic knowledge of the recipient peer in terms of both concepts (i.e. vertices) and relationships between them (i.e. edges). All the edges received from the donor network (new or already present) set their remembering value in memory to 1 (i.e. they are “activated” by the “conversation”).

C. Semantic data dissemination

The previous selection of the most relevant semantic concepts, with respect to the current interaction, drives the next step in the data exchange process: the selection of relevant data items to exchange. In order to carry out this operation, we exploit another simple decision rule derived from the cognitive science field: the *tallying* heuristic. For this step, we refer to the pseudo-code given in Alg. 3. In our scenario, each data item is associated to a set of tags, that we term as its semantic description (*semanticDesc(i)* in Alg. 3, line 6). We consider the vertices selected for inclusion in the *contributed* network as the m “cues” (out from all the M nodes in the donor network) needed by the *tallying* heuristic. Counting the number of favourable cues simply corresponds to counting the cardinality of the intersection between each data item semantic description and the set of nodes included in the *contributed* network (lines 5–7). We consider that the other party sends to the node the list of the IDs of the items it already owns. Thus, those data items can be directly pruned out from the

selection process (line 4). Moreover, like for the exchange of semantic concepts, we assume that there is a maximum number of exchangeable data items *data_limit*. Thus, once the data items have been ranked according to *tallying*, the first *data_limit* ones are selected to be passed to the other interacting peer (lines 8–9).

Algorithm 3 Tallying algorithm

- 1: Let \bar{V} be the nodes of the contributed network;
 - 2: Let I be the set of data items of the node
 - 3: Let J be the data items owned by the other peer
 - 4: Let $I' = I - J$
 - 5: **for** each $i \in I'$ **do**
 - 6: Let $tall(i) = |semanticDesc(i) \cap \bar{V}|$
 - 7: **end for**
 - 8: Rank I' in descending order according to the *tall* values
 - 9: Send the first *data_limit* items of I' to the other node
-

V. PERFORMANCE EVALUATION

A. Simulated Environment

In this section, we show the results of the proposed solution obtained in a simulated environment. The scenario consists of 99 mobile nodes moving in a 1000²m area. Peers’ movement patterns are taken according to the HCMM model [3]. This is a mobility model that integrates temporal, social and spatial notions in order to obtain an accurate representation of real user movements. Specifically, in HCMM the simulation space is divided in cells representing different social communities. In this preliminary study, we consider that there exists only one social community. Data assigned to these nodes is selected from the CoPhIR dataset [14]. This dataset is made up of more than 100M images coming from Flickr. For each user’s image, it is possible to know the list of associated tags. In order to create a useful dataset to test our solution we proceeded in the following way. Images were selected in order to have the initial users’ semantic knowledge clustered around three main concepts. Figure 2 represents the entire knowledge present in the network at the beginning of each simulation, i.e. the graph of the union of all nodes’ initial semantic networks. We pointed our attention on such initial configuration in order to both study the ability of each user to retrieve the information semantically related to its initial interests and to analyze the overall permeation of data in the network. We observed and studied the evolution of the knowledge and content acquisition processes generated by the interactions between users. Here, we defined two different measures of knowledge and content dissemination. We measure knowledge dissemination (KD) by computing how much of the starting global knowledge reaches nodes at the end of the simulation. 100% would mean that all semantic data reach all nodes. Given the clustered structure of the initial overall semantic network

(i.e. the fact that users’ interests are clustered) we don’t expect to reach 100% for this index. We also define a second measure, the *coverage*, as the portion of items owned by a node over all the items that contain a tag matching one of the concepts in its semantic network. For this index we expect to reach 100%, as this means that users receive all the data items related to concepts they have in their semantic network. Results reported in the following simulations are the statistics collected on 10 different tests obtained from 10 different mobility traces of the HCMM model and averaged across all nodes. Each simulation experiment ran for 25000 sec (required for our indices to reach stationarity).

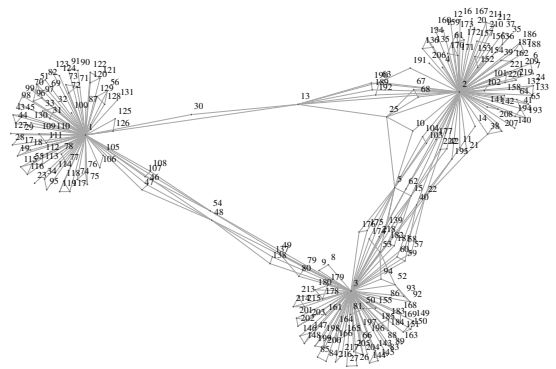


Figure 2. Initial knowledge graph defined as the union of all nodes’ semantic networks at the beginning of the simulation

B. Experimental results

In order to simplify the analysis, in the following we set $\theta_{rec} = 5$. As stated in Section IV, the remembering process plays (both in human brain and in our model) an important role in knowledge acquisition. Figure 3a shows that different settings of *remembering threshold* f_{min} produce largely different results in knowledge acquisition. As one could expect, the longer the tags lifetime in the semantic network, the higher the diffusion they reach. Interestingly, there exist a threshold (50s) under which nodes’ semantic networks cannot grow, triggering an information loss process that leads to a decrease of the owned knowledge. In Figure 3b we can notice that the coverage follows, even with a certain delay, the tag dissemination trend reaching the convergence after 1000s. This delay is a peculiar behaviour of our approach and it is due to some kind of inertia related to the spreading of tags and data items (tags spread first, data items are ”pulled” as a side effect). Indeed, there is a transient moment (from 200s to 800s) where nodes start to increase their knowledge faster than they exchange correlated data items. When the semantic networks growth starts to slow down the data dissemination process reaches the convergence.

Due to space limits we do not show the results concerning to the impact of the *retrieval value* on the knowledge dissemination. However, as one can expect, the knowledge




Figure 3. Knowledge diffusion (a) and Content coverage (b) over time for different remember thresholds

dissemination increases with both the maximum number of tags shared and the minimum retrieval value threshold. However, in contrast with the common sense, it is not true that the more tags nodes exchange the more data items they receive. As shown in Table I, we measured the dissemination performance for different values of tag and data limits noticing that even passing almost 10 tags and 5 contents for each encounter we reach a data dissemination greater than 70%. It is worth noting that, this value refers to the spreading of the all data items present in the network, that differs from the coverage measure which refers to percentage of items received by nodes w.r.t. their acquired knowledge. In the light of these results we can state that once set the tag limit, increasing the data limit does not affect the content dissemination performance, but has a strong impact on the convergence speed of this index. For example, we observed a difference of more that 10000s on the convergence time between the parameter pairs ($tag_limit = 10, data_limit = 2$) and ($tag_limit = 10, data_limit = 10$). This behaviour has a strong impact in terms of resource consumption, because our approach turns to be thrifty both in terms of memory and bandwidth usage.

Tag Lim.	Data Limit		
	2	5	10
5	0.67 ± 0.11	0.64 ± 0.12	0.64 ± 0.12
10	0.75 ± 0.10	0.71 ± 0.10	0.71 ± 0.10
25	0.79 ± 0.05	0.79 ± 0.05	0.79 ± 0.05

Table I
FINAL MEAN CONTENT DISSEMINATION WITH CONFIDENCE INTERVALS AT 95%

VI. CONCLUSION

In this paper we presented a model enabling nodes belonging to an opportunistic network to deal with semantic knowledge and content dissemination. According to the CPW convergence scenario where nodes are actual avatars of their users, we based our solution on well established models coming from the cognitive science field. Namely, we exploited the associative network memory model to craft a representation of the semantic knowledge on which we run two fast and frugal decision making strategies, i.e.

Fluency and Tallying Heuristics, in order to identify which semantic data and data items to exchange between nodes upon encounters. Preliminary results demonstrate the feasibility of our model to take advantage of the data semantic aspects in order to disseminate relevant information among nodes. Moreover, the initial sensitivity analysis conducted highlights the (sometimes not obvious) relations between the model parameters.

ACKNOWLEDGMENT

This work is funded by the EC under the FET AWARENESS RECOGNITION (FP7-257756), FIRE EINS (FP7-288021).

REFERENCES

- [1] M. Conti, et al *Looking Ahead in Pervasive Computing: Challenges and Opportunities in the Era of Cyber-Physical Convergence*, Elsevier Pervasive and Mobile Computing, 2011
- [2] C. Boldrini, M. Conti, A. Passarella, *Design and performance evaluation of contentplace, a social-aware data dissemination system for opportunistic networks*, Comput. Netw. 54, 589-604, 2010.
- [3] C. Boldrini and A. Passarella. *Hcmm: Modelling spatial and temporal properties of human mobility driven by users social relationships*, Comput. Commun. 33, 1056-1074. 2010.
- [4] C. Boldrini, A. Passarella, *Data Dissemination in Opportunistic Networks*, Ch. 12 of Mobile Ad hoc networking: the cutting edge directions, Eds. S. Basagni, M. Conti, S. Giordano, I. Stojmenovic, Wiley, 2012.
- [5] M. Conti, M. Mordacchini, and A. Passarella, *Data Dissemination in Opportunistic Networks using Cognitive Heuristics*, Proc. IEEE WoWMoM Workshop on Autonomic and Opportunistic Computing (AOC), 2011.
- [6] G. Gigerenzer, *Why heuristics work*. Perspectives on Psychological Science, 3(1), 20-29. 2008.
- [7] V. Lenders, M. May, G. Karlsson, C. Wacha, *Wireless ad hoc podcasting*, SIGMOBILE Mob. Comput. Commun. Rev. 12, 65-67, 2008.
- [8] E. Yoneki, P. Hui, S. Chan, J. Crowcroft, *A socio-aware overlay for publish/subscribe communication in delay tolerant networks*, Proc. MSWiM, 225-234, 2007.
- [9] R. Bruno, M. Conti, M. Mordacchini and A. Passarella. *An analytical model for content dissemination in opportunistic networks using cognitive heuristics*, Proc. MSWiM, 61-68, 2012.
- [10] B. Gawronski, B. K. Payne, *Handbook of implicit social cognition: Measurement, theory, and applications*. The Guilford Press, 2010.
- [11] L. J. Schooler, R. Hertwig. *How forgetting aids heuristic inference*. Psy. Rev. 112.3 (2005): 610.
- [12] R.M. Dawes, *The robust beauty of improper linear models in decision making*. American psychologist 34.7 (1979): 571.
- [13] H. Ebbinghaus.: *Memory: A contribution to experimental psychology*. Number 3. Teachers college, Columbia university (1913)
- [14] P. Bolettieri, et al. *CoPhIR: a Test Collection for Content-Based Image Retrieval*. CoRR abs/0905.4627: (2009)