Data Dissemination in Opportunistic Networks using Cognitive Heuristics

Marco Conti, Matteo Mordacchini, Andrea Passarella
IIT-CNR, Pisa , Italy
Email: {m.conti,m.mordacchini,a.passarella}@iit.cnr.it

***Abstract*—It is often argued that the Future Internet will be a very large scale content-centric network. Scalability issues will stem even more from the amount of content nodes will generate, share and consume. In order to let users become aware and retrieve the content they really need, thesenodes will be required to swiftly react to stimuli and assert the relevance of discovered data under uncertainty and only partial information. The human brain performs the task of information filtering and selection using the so-called cognitive heuristics, i.e. simple, rapid, low-resource demanding, yet very effective schemes that can be modeled using a functional approach. In this paper we propose a solution based on one such heuristics, namely the *recognition* heuristic, for dealing with data dissemination in opportunistic networks. We show how to model an algorithm that exploits the environmental information in order to implement an effective dissemination of data based on the recognition heuristic, and provide a performance evaluation of such a solution via simulation.**

I. Introduction

In the Future Internet scenario, mobile devices will be part of a crowded information landscape. Data generation and consumption patterns will be very dynamic, with a large fraction of data being produced and stored on the user devices themselves. A considerable part of these data will also be very contextualized, i.e. relevant only at specific times and/or geographic areas, and of interest only for specific groups of users. Solutions like web servers, CDNs will not be sufficient efficiently manage and disseminate data in such dynamic conditions, while users collaboration will be a must. Conventional P2P approaches will also not be enough, due to the presence of mobile nodes and challenged networking conditions, as in opportunistic networking scenarios [12]. According to this view, data dissemination paradigms are required, in which individual (mobile) nodes contribute a fraction of their resources to circulate data, selecting those data items to store based on their utility for the overall dissemination process. To optimize this process, each node should react quickly to the discovery of new data items and it should consider which are the most useful items to take among at a given point in time. This selection should be performed swiftly, since the contextualized nature of information could make it aged before a complex evaluation process has ended. Moreover, not all the variables required to perform a complete evaluation may be known. Finally, nodes - in general - will contribute limited resources to the dissemination process (e.g. in terms of computing and storage capabilities). Thus, the data selection process must be very lightweight and able to perform a sharp distinction between data items, since only a very limited part of them could be stored.

One way to overcome these problems is to add autonomic capabilities to those devices. In this paper, we explore a new (to the best of our knowledge) direction in the autonomic networking field, i.e., we exploit results coming from the cognitive psychology area, by using models of how the human brain assesses the relevance of information under partial knowledge. With respect to conventional artificial intelligence approaches, we do not seek to reproduce the physiology of cognition, but exploit *functional* description of cognitive processes, termed in the literature *cognitive heuristics* (e.g. [6]). In computer science, heuristics are computational methods that try to optimize a problem by producing stochastically good results. They are obtained by pruning the search space through an iterative improvement of a candidate solution, with regard to a given measure of quality. On the other hand, cognitive heuristics are fast, frugal and adaptive strategies of the brain that allow humans to face complex situations by addressing simpler problems. Cognitive heuristics are characterized by the fact that they are effective, simple rules, requiring little estimation time and working under incomplete knowledge of the problem space. They may act by ignoring the quality criterion used to evaluate the goodness of the final results, yet producing very effective results. Due to these characteristics, they are a very good candidate for translating a mental process of information selection and acquisition to devices with limited resources, dealing with a very dynamic and crowded information context.

Among these heuristics, Goldstein and Gigerenzer [6], [5] have studied and modelled one of the simplest and more effective of them: the recognition heuristic. This heuristic assumes that merely recognizing an object is sufficient to take decisions that would theoretically require much more information about the object's properties. A detailed description of the recognition heuristic is provided in Section III. This kind of heuristic has proved to be not only fast and frugal, but it is also *ecologically rational*, in the sense that it exploits structures of information coming from the environment in order to work.

In this paper, we want to exploit the fast and frugal recognition heuristic in an opportunistic networking scenario. In this scenario, nodes carry some data, are interested in acquiring specific types of content and have the possibility to store some of the data encountered when moving in the environment. We propose an exploitation of the recognition heuristic to let each node rapidly decide which is the utility of taking one data item instead of another upon making direct (i.e. one-hop) contact with other nodes. First of all we define the requisites needed to implement the recognition

978-1-4577-0351-5/11/$26.00 ©2011 IEEE

heuristic in an opportunistic environment by defining the main variables involved in this process. Then, we propose an algorithm inspired by the model of Goldstein and Gigerenzer that exploits the recognition heuristic in order to simplify and limit the complexity of the data selection task. Finally, we evaluate by simulation the data diffusion process when nodes exploit the proposed solution.

The rest of this paper is organized as follows. In Section II we briefly survey the state of the art on data dissemination in opportunistic networks. In Section III we give a more precise description of the recognition heuristic. In Section IV we introduce how the recognition heuristic can be implemented by mobile devices, while in Section V we define an algorithm that exploits it for the purpose of data dissemination in an opportunistic network. Section VI presents the experimental results obtained via simulation. Finally, Section VII concludes the paper.

II. Related work

Data dissemination algorithms have been proposed for diverse families of mobile networks. The work in [14] is representative of a body of work focused on caching strategies for well-connected MANETs. In this paper we focus on more challenged networking environments, where such policies cannot be applied. To the best of our knowledge, the most advanced approaches for data dissemination in opportunistic networks exploit information about users social relationships to drive the dissemination process [15], [4], [2]. Specifically, the work in [15] defines a pub/sub overlay in which brokers are the most "socially-connected" nodes, i.e., those nodes that are expected to be most available and easily reachable in the network. SocialCast [4] proposes a first attempt to exploit social information in dissemination processes. This is also the goal of the work in [2], where, however, a more refined and complete approach is used. Specifically, dissemination is driven by the social structure of the network users, such that nodes store data items that are likely of interest to users they have social relationships with (and who, therefore, are expected to be in touch in the near future). Other data dissemination schemes for opportunistic networks include those defined in the PodNet project [8] which, however, do not exploit social information, but incorporate well-known caching policies such as uniform and greedy selection.

With respect to these approaches, in this paper we take a completely new direction, by borrowing models of human cognitive processes coming from the cognitive psychology domain. As this approach is still totally unexplored, in this paper we limit the set of contextual information that we use to the very minimum, and, for example, we do not exploit information about users social structures. This allows us to obtain initial exploratory results about the feasibility of this novel approach.

III. The Recognition Heuristic

Heuristics are simple and adaptive strategies used by living species (including humans) to perform specific tasks in face of limited time, knowledge and computational capabilities.The main characteristic of heuristics is that they do not make use of all the information required - in principle - to find an exact solution to a given problem. They reach a decision using only limited knowledge and using a process simple enough to produce a decision in short time. One of the simplest of such heuristics is the recognition heuristic. The recognition heuristic works by inferring that a recognized object has an higher value than an unrecognized one, with respect to a given criterion. It works even when the real criterion value is not available, not known or requires further, more complex (and longer and expensive) reasoning to be computed. It exploits the presence in the environment of some *mediators* that carry information (coded in variables) used by the heuristic to approximate the value of objects with respect to the criterion. Mediators spread these variables in the environment, thus determining which objects are recognized. As shown in Fig. 1, the correlation that binds the criterion and the mediator is called the *ecological correlation*. The relationship between the mediator and the contents of recognition memory is called the *surrogate correlation*, since the mediator is a surrogate for the inaccessible criterion. Finally, the accuracy of the recognition heuristic is called the *recognition validity*.

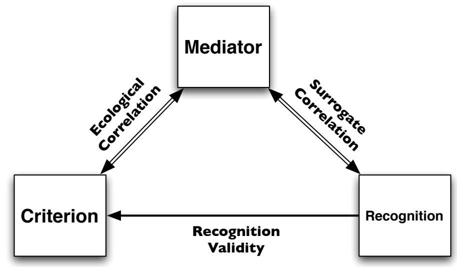

Figure 1. Ecological Rationality of the recognition heuristic

In order to better understand how the recognition heuristic works, Gigerenzer and Goldstein [5] use, as an example, the estimation of a university endowment. In this case the criterion is the value of the endowment. This is generally not publicly available. Anyway, newspapers could act as mediators, since they periodically publish news related to the biggest universities. The number of times a university institution appears on newspapers could be a strong indicator that it has larger endowments than institutions that do not or rarely appear on the media. In this case, newspapers play the role of mediators and the mediator variable related to the criterion is the number of citations. When a person has to choose which university has the biggest endowments between a couple of institution names, she uses the recognition heuristic and chooses a recognized name against an unknown one. Clearly, in this case, newspapers influence the recognition, since the more they cite an institution, the more probably it will be recognized. This example highlights the so-called *ecological rationality* of the recognition heuristic, i.e. the fact that the recognition ability is reinforced by the stimuli received from the environment.

The simple description of the recognition heuristic made it a powerful tool for making predictions about a given criterion. The recognition heuristic can be exploited as a support in decision-making processes. As such, it has been successfully used in various fields [9], like financial decision-making processes [11], forecasting future purchase activities [7] or even sport events results [13] or political election outcomes [10].

IV. The Recognition Heuristic for Data Dissemination in Opportunistic Networks

To define how to exploit the recognition heuristic, hereafter we consider a scenario in which each data item pertains to a specific *channel* (or interest). Each node generates and owns some items of possibly different channels. It is interested in retrieving and keeping items belonging to a specific channel only. In this scenario nodes collaboratively contribute to the diffusion of information by storing and exchanging some of the items they discover coming in contact with other devices, even if those items do not belong to the channel they are interested in. We assume that node contribute a limited shared storage space to the diffusion process (whose size is throughout denoted with S). The diffusion process happens through contacts between nodes. When two nodes come in contact, they exchange summaries of their data items. The node ranks data items stored by itself and by the encountered peer according to their utility, and updates its storage space by storing the most useful items only (until the storage space is full). Computing the utility of items for the diffusion process is a hard (or impossible) target criterion to evaluate for a single node, as it requires in general complete knowledge of the status of the network. Thus, the application of a simple, fast and effective strategy such as the recognition heuristic can significantly reduce the complexity of evaluating this criterion.

In order to exploit the recognition heuristic, the first step is to define the elements upon which recognition will be made. After that, we need to design a proper algorithm that, starting from the recognition heuristic, effectively filters the information, with the aim of maximizing the utility of the exchange of objects among nodes. To this end, in this section we characterize the elements that allow to use the recognition heuristic in this environment. In the next section, we define an algorithm that relies on this heuristic, a modified version from one present in the cognitive science literature. The aim of this algorithm is to limit as much as possible the use of precise attributes and complex operations by effectively filtering out most of the information using the recognition heuristic only.

In order to use the recognition heuristic in our scenario, some steps must be followed. Specifically, we have to identify:

- the features (like the name of cities or universities in the examples of Goldstein and Gigerenzer) that are highly correlated with the selection criterion and that are thus spread by the mediators;
- the environmental mediators;
- the way by which nodes implement the heuristic based on the information collected from mediators

As for the first point, we consider two simple factors that determine the utility of a data item, i.e., the popularity of its channel, and its availability (these factors have always been considered as fundamental in the data management literature, starting from the area of web caching [1]. Specifically, the utility of a data items is positively correlated with the popularity of its channel (how many users are interested in that item), and negatively correlated with its availability (how many times that item is already replicated).

As for the second point, we use nodes themselves as mediators, while the variables they spread are, respectively, the channel they are interested into, and the set of items they are currently storing in their shared storage space.

As for the third point, the bottomline idea is to use two recognition heuristics to separately recognize channels and data items. Intuitively, a node recognizes a channel as soon as it becomes "enough popular", i.e., as soon as the node encounters enough nodes that are interested in that channel. Furthermore, a node recognizes a data item it is "spread enough", i.e., as soon as it is encountered on at least a given number of other nodes. This approach is very similar to what is referred to as *inference-from-memory* in [5]. More specifically, each node maintains a separate *recognition cache* for channels and data items. Entries of the cache correspond to channels of interest for or data items carried by encountered nodes, respectively. Each entry contains a counter and a TTL associated with the channel or data item. Whenever a node interested in a channel (or storing a data item) is encountered, the associated counter is incremented and the TTL reset. When the counter reaches a certain threshold, the corresponding channel or data item is deemed as recognized. Furthermore, the TTL is incremented at each time slot. When the cache becomes full and replacement must occur, the entry with the highest TTL is selected for replacement. If it corresponds to a recognized channel (or data item), this entry is stored in a Bloom filter. Otherwise, it is dropped. The complete recognition algorithm is shown in Algorithm 1.

Intuitively the algorithm keeps track of encountered channels or data items until they are recognized. When the cache is empty, the entry corresponding to the least recently seen object is selected. Bloom filters allow nodes to distinguish, among entries that are not in the cache, those that correspond to recognized items (stored in the Bloom filter), and not recognized items. This is important in case such items are encountered again, as, if they are in the Bloom filter, they can be immediately recognized again. Note that this algorithm mimics the way in which the human brain refreshes, flushes and recalls "items" in memory.

V. A Modified *Take The Best* Algorithm for Opportunistic Networks

Having described how to implement heuristics, we now present an algorithm that exploit them in the data dissemination process. Also in this case, we take inspiration from the cognitive psychology literature [6]. The algorithm defined in [6] (named *Take the Best*) mimics a fast and frugal way of reasoning for choosing among two alternatives. The goal of the algorithm is comparing two objects. To this end, objects are tested against an ordered set of cues, stopping at the first (best) cue that discriminates among them. When none of the cues can discriminate, the algorithm chooses by some additional discriminating criterion, which usually requires much more complex information to be evaluated

Algorithm 1 Recognition algorithm

Let i be an observed channel/item;
Let H be a hashed index of removed channels/items
Let R_θ be the recognition threshold
if Cache.contains(i) **then**
  Increment i counter
  reset i.TTL
else
  if Cache is full **then**
    Select the item o with the oldest TTL
    if o.counter $\geq R_\theta$ **then**
      Move o to H
    end if
    Drop o
  end if
  Put i in the Cache
  i.counter = 1
  Set i.TTL
end if

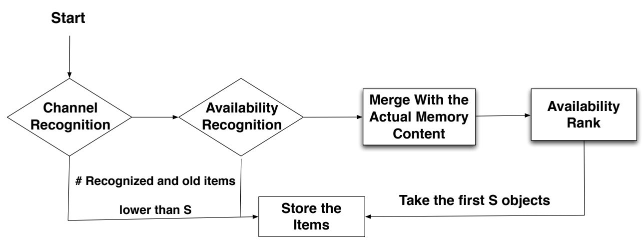

Figure 2. Modified Take The Best Algorithm

with respect to the cues. We propose to adapt this algorithm in the scenario we are considering. When a node meets another peer, it ranks the objects of the other node using an adaptation of the *Take The Best* algorithm, depicted in Fig. 2 and Algorithm 2. The first two cues we considered consist of the recognition of channels and items, using the algorithm presented in Section IV. The first cue is the channel recognition: items of recognized channels are ranked higher than the others and selected for the next steps. If the total size of remaining items (considering both the node and the peer shared storage spaces) is greater than S (the size of the node's shared storage space), items are further discriminated using the second cue, i.e. the recognition of items. In this case the recognition assumes a *negative* meaning, as recognized items (already very spread) are ranked lower than the others and they are not considered anymore. If further discrimination have to be carried out to fill the node's shared storage space, the precise value of estimated availability of items is considered, and less available items are ranked higher. As for the original *Take The Best* Algorithm, not all the steps are required, and the last (and more costly) one is run only on a subset of the items.

A. The Less–is–More Effect

Goldstein and Gigerenzer show that the recognition heuristic [5] and the *Take The Best* algorithm are subject to the so-called *less–is–more* effect [6]: When increasing the number of recognized items, the number of accurate inferences will increase up to a certain point, but thereafter decrease. This is due to the diminished discrimination power of the recognition heuristic, since too many items are recognized at the same time. In order to take advantage of the

Algorithm 2 Modified Take The Best Algorithm

Let R be a set of items received from another node;
M be the actual storage content
Let S be the storage capacity limit
Let C be the channel chache and O be the item cache
Let H_C and H_I be the hashes of old recognized channels and items
Let RC_θ be the recognition threshold for channels
Let RI_θ be the recognition threshold for items
Let $I = R - M$
Let $recChannels = \emptyset$
for each $i \in I$ **do**
  if C.contains(i.channel) & i.channel.counter $\geq$ RC_θ OR H_C.contains(r) **then**
    $recChannles \cup = i$
  end if
end for
Let $recItems = \emptyset$
if $recChannels$.size + M.size $> S$ **then**
  for each $r \in recChannels$ **do**
    if (O.contains(r) & NOT r.counter $\geq$ RI_θ) OR NOT H_I.contains(r) **then**
      $recItems \cup = r$
    end if
  end for
  if $recItems$.size + M.size $> S$ **then**
    Let $M' = M \cup recItems$
    Rank M' in ascending order according to the counters of its items
    Select and keep in M the first S objects of M'
  else
    $M \cup = recItems$
  end if
else
  $M \cup = recChannels$
end if

less–is–more effect, the storage capacity of the recognition caches should be adapted in order to achieve the maximum accuracy. Identifying autonomic algorithms to adjust this capacity is one of the key directions of future work.

VI. Experimental Results

We evaluate the proposed algorithm by simulating the following scenario. We consider 45 nodes, divided into three different groups. In order to simulate real user movement patterns, nodes move in a 4 x 4 grid (1000 m wide), according to the HCMM model [3]. The HCMM model is a mobility model that integrates temporal, social and spatial notions in order to obtain an accurate representation of real user movements. In the simulation scenario, groups represent set of users that have social and spatial relationships. Groups are initially assigned to a home cell and any physical contact among groups is avoided. Thus, the only way to exchange and obtain data among groups is through node mobility. Nodes can move in the cell of their group only. A few nodes in each group (named *travellers*) bridge between communities by visiting more than one group. This model well represents social communities, in which people typically stay, with a few people commuting between different communities due to different social relationships. In the simulation settings, each group has two travellers, one for each of the other groups. The objects available in the network are assigned to channels and there are as many channels (n_c) as groups. Each channel has a total of 99 objects and each group originates $1/n_c$ items per channel, i.e. there are 33 items of each channel per group. All the objects are generated at the start of the simulation. Each node subscribes to one channel only. Within each group, node interests follow a Zipf law with parameter 1. Interests are rotated, so that the most popular channel in a group is the second in another and the third one in the other, and so on. Note that in this scenario data items can reach interested users in communities other than those where they are generated only through nodes mobility. Therefore, it allows us to highlight the effectiveness of the data dissemination algorithm. The simulation runs for 50,000 seconds. At the start of the simulation each node subscribes to a given channel. The performance figure is the hit rate, computed at various time instants after the simulation starts. It is defined as the ratio between the number of retrieved objects of the subscribed channel and the total amount of objects of the channel. By default, on each node, the channel recognition cache size is 3 (i.e. all channels can be recognized), the items recognition cache is 10, and the shared storage space has 10 slots (items are assumed to be of equal size). Simulation have been replicated in independent conditions 10 times, and confidence intervals (with 95% confidence level) have been computed. As they are very narrow, in the plot we only show the central values of the confidence intervals.

Simulations were run changing the values of the thresholds for both the channel and the items recognitions and by varying size of the shared storage space, S. The first set of experiments (Figs. 3–4) shows, first of all, that by using the recognition heuristic the system is able, after some time, to reach 100% hit rate. This is a very important result. Using the proposed recognition heuristic result in an extremely lightweight data dissemination system, which nevertheless proves to be very effective. Another aspect highlighted by this set of results is the effect of varying the item recognition threshold with different values of the channel recognition threshold. The main result is that for very low value of the item threshold ($RI_\theta = 2$) the information diffusion algorithm is not able to obtain a 100% hit rate. For the other values, changes in the item threshold implies little or almost no differences in the algorithm performance. This results (particularly evident in Fig. 4) are a sign of the $less-is-more$ effect. In order to have a in-depth analyses of this case, Fig. 5 shows the results obtained by fixing $RI_\theta = 2$ and varying the channel threshold. Such a low value of the item threshold implies a higher probability of recognizing items as widespread. As a consequence, at a given time *all* the items are recognized as too diffused by all the nodes. In this case, items are not exchanged anymore, since all of them are not recognized as useful. In the other cases in Fig. 3 and 4, not all the items were recognized as too diffused, thus allowing to reach a 100% hit rate. $RI_\theta = 2$, as in Fig. 5, always leads to lower performances. Moreover, this figure highlights that different values of the channel threshold imply different information diffusion speeds. The lower RC_θ, the more rapid the convergence of the algorithm. The longer it takes to converge, the higher is the probability for a node to see more copies of a given data before starting to take it. Thus, with a slower convergence rate, is more probable to have all the items recognized as widespread at increasingly further points from 100% of hit rate. The

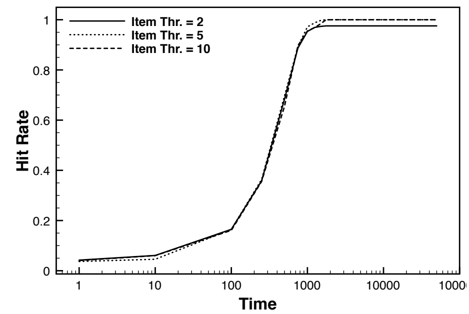

Figure 3. Hit Ratio with a Channel threshold = 3

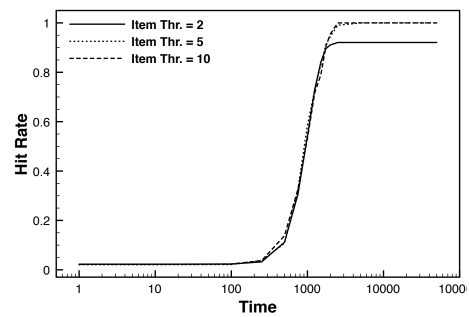

Figure 4. Hit Ratio with a Channel threshold = 10

effect of RC_θ values on the convergence speed is visible also with other values of RI_θ. Fig. 6 reports the results achieved with $RI_\theta = 10$. In this case, the hit rate always converges to 100%, but , as in the previous figure , the convergence velocity is strongly determined by the different values of RC_θ. Fig. 7 highlights the effects of the shared storage space to the algorithm effectiveness. The results are obtained with $RC_\theta = 2$, $RI_\theta = 5$. The number of slots of the shared memory is fixed to 2 and 50, respectively. A larger amount of shared memory allows a greater replication factor, since is more probable that multiple copies of the same data can be stored by different nodes at the same time. This fact easy the diffusion of items, thus increasing the converge speed. Nonetheless, even a very low number of slots does not affect the final performance (100% hit rate). On the other hand, they result in a slower convergence rate, due to the increased difficulty to share and spread items in the network.

VII. Conclusions

Cognitive heuristics are models of how the human brain assess the relevance of information using only partial knowl-

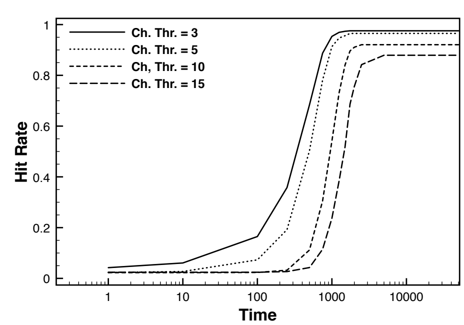

Figure 5. Hit Ratio with an Item threshold = 2

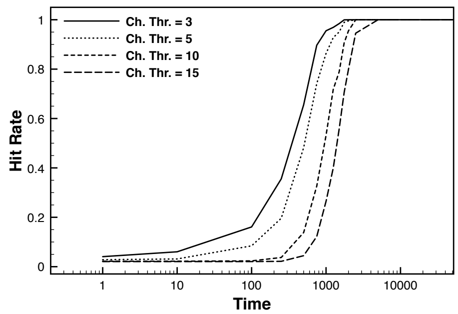

Figure 6. Hit Ratio with an Item threshold = 10

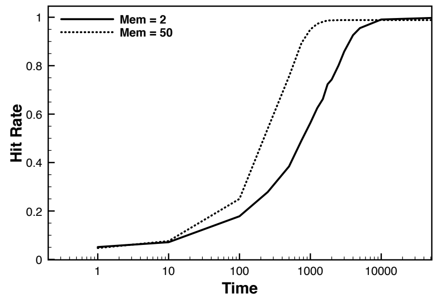

Figure 7. Hit Ratio with Variable Memory Sizes

edge of the problem space and very limited resources. In this paper we present an initial attempt to exploit these models (already established and coded in the cognitive psychology field) to drive data dissemination processes in opportunistic networking environments. Specifically, we show how a fast and frugal cognitive heuristic, i.e. the *recognition* heuristic, can be used. The recognition heuristics discriminates objects with respect to a given criterion, without requiring to collect all the information needed to exactly compute the criterion. It assumes that recognized objects have higher value (with respect to the criterion) than non recognized objects, and discriminates among them accordingly. Several instances of recognition can be chained together, in order to model how the human brain discriminates in complex scenarios, solving efficiently complex decision making problems. In this paper we adapt these models to address the data dissemination problem in content-centric mobile networks. We first define an algorithm by way of which the recognition heuristic can be implemented by the nodes of an opportunistic network. Then, we show how nodes can efficiently combine multiple instances of the recognition heuristic to assess the relevance of available data objects, thus deciding what to store and what to drop. Simulation results show the potential of such an approach and highlight how a correct tuning of the heuristic parameters leads to a fast and highly effective dissemination of data items. In particular, the emergence of the *less-is-more* effect (well understood in the cognitive psychology field) highlights that an optimal configuration exists for the recognition algorithm to obtain the best performance.

Results presented in this paper are promising, and provide strong indications that using cognitive heuristics to cope with scalability issues in Future Internet environments is a sensible direction. Key topics for future research include a complete understanding of the heuristic parameters on the data dissemination efficiency. Specifically, analytical models are required to understand the impact and the interplay of the parameters. For example, results presented in the paper already highlight a non-trivial joint effect on the data dissemination efficiency played by the thresholds used by the different recognitions heuristics. Furthermore, understanding how data dissemination works when additional context information (such as social relationships between users) is exploited is another interesting topic. Finally, it will also be interesting to understand whether other heuristics (beyond recognition) can be effectively applied to the data dissemination or other related problems.

VIII. Acknowledgements

This work is funded by the EC under the FET-AWARENESS RECOGNITION Project, grant 257756.

References

[1] Abdullah Balamash and Marwan Krunz. An overview of web caching replacement algorithms. *IEEE Communications Surveys and Tutorials*, 6(1-4):44–56, 2004.

[2] Chiara Boldrini, Marco Conti, and Andrea Passarella. Design and performance evaluation of contentplace, a social-aware data dissemination system for opportunistic networks. *Comput. Netw.*, 54:589–604, March 2010.

[3] Chiara Boldrini and Andrea Passarella. Hcmm: Modelling spatial and temporal properties of human mobility driven by users' social relationships. *Comput. Commun.*, 33:1056–1074, June 2010.

[4] Paolo Costa, Cecilia Mascolo, Mirco Musolesi, and Gian Pietro Picco. Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks. *IEEE Journal on Selected Areas in Communications*, 26(5):748–760, 2008.

[5] Gerd Gigerenzer and Daniel G. Goldstein. Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1):75–90, 2002.

[6] Daniel G. Goldstein and Gerd Gigerenzer. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996.

[7] Daniel G. Goldstein and Gerd Gigerenzer. Fast and frugal forecasting. *Int. Journal of Forecasting*, 25:760–772, 2009.

[8] Vincent Lenders, Martin May, Gunnar Karlsson, and Clemens Wacha. Wireless ad hoc podcasting. *SIGMOBILE Mob. Comput. Commun. Rev.*, 12:65–67, January 2008.

[9] Julian N. Marewski, Wolfgang Gaissmaier, and Gerd Gigerenzer. Good judgments do not require complex cognition. *Cogn. Process*, 11:103–121, 2010.

[10] Julian N. Marewski, Wolfgang Gaissmaier, Lael J. Schooler, Daniel G. Goldstein, and Gerd Gigerenzer. From recognition to decisions: Extending and testing recognition-based models for multialternative inference. *Psychonomic Bulletin & Review*, 17(3):287–309, 2010.

[11] Marco Monti, Laura Martignon, Gerd Gigerenzer, and Nathan Berg. The impact of simplicity on financial decision-making. In *Proc. of CogSci 2009, July 29 - August 1 2009, Amsterdam, the Netherlands*, pages 1846–1851. The Cognitive Science Society, Inc., 2009.

[12] L. Pelusi, A. Passarella, and M. Conti. Opportunistic networking: data forwarding in disconnected mobile ad hoc networks. *Communications Magazine, IEEE*, 44(11):134 –141, 2006.

[13] Sascha Serwe and Christian Frings. Who will win wimbledon? the recognition heuristic in predicting sports events. *J. Behav. Dec. Making*, 19(4):321–332, 2006.

[14] Liangzhong Yin and Guohong Cao. Supporting cooperative caching in ad hoc networks. *IEEE Trans. Mob. Comput.*, 5(1):77–89, 2006.

[15] Eiko Yoneki, Pan Hui, ShuYan Chan, and Jon Crowcroft. A socio-aware overlay for publish/subscribe communication in delay tolerant networks. In *ACM MSWiM 2007*, pages 225–234, 2007.