Coupons: Wide Scale Information Distribution for Wireless Ad Hoc Networks

Kevin Almeroth[1] and Anargyros Garyfalos[2]

[1] Department of Computer Science, University of California, Santa Barbara, CA 93106-5110, USA
[2] British Telecommunications Plc, Intelligent Systems Lab, Adastral Park, Ipswich IP5 3RE,UK

Abstract. Integrating ad hoc networks into the Internet requires overcoming a number of difficult technical challenges. In particular, ad hoc networks must not only overcome intermittent connectivity, but they also need a strong incentive mechanism to encourage users to participate in the cooperative relay of data traffic. We believe that inherent in solving these problems is the development of new applications that might, in fact, be more easily deployed in an ad hoc environment than in a traditional fixed network infrastructure. To this end, we develop and evaluate the idea of "coupons" for wide-scale information distribution in ad hoc networks. "Coupons" provides a simple incentive to nodes for relaying a piece of information. By using mechanisms on top of basic flooding to efficiently control distribution, it provides an elegant solution for scalable data dissemination with reduced network costs.

1 Introduction

We are evolving into a society with nearly constant access to the Internet and information. Currently access at the fringes of the network is provided using ad hoc networking and novel devices like PDAs, advanced cell phones, etc. Originally conceptualized for military networks, ad hoc networking allows a connection to the Internet even when not in range of a wireless access point. This is achieved by connecting through other users who relay transmissions to those who are connected.

However, deploying a wireless ad hoc network in a commercially arena raises a number of issues. Intermittent service along with a wide range of node capability may result in a frequently changing environment. Since users are not yet accustomed to such unpredictable behavior, emphasis is given on providing as efficient connectivity as possible. Much of the work to date attempts to address this problem by leveraging the resources of other nodes, such as cooperative nearby users. However it is often neglected that in a commercial setting this may not always be the case since operational costs prohibit users from participating in any scheme without appropriate compensation. Such issues have a significant impact on how existing Internet applications are ported to ad hoc networks and on the kinds of new applications that might be created.

The goal of this paper is then to develop applications and services that are uniquely suited to ad hoc environments. "Coupons" is a scheme that expands on the basic idea of information dissemination in ad hoc networks by investigating incentives for collaboration. Cooperation is stimulated by adopting a basic pyramid scheme where users are awarded credits as they pass a received piece of information (or a *coupon*) to other people. Using a simple flooding mechanism, as one user forwards the coupon to another, a pyramid is built with the users residing at the top receiving more benefit than those at the bottom.

The basic concept of a coupon has been proposed elsewhere[1, 2], but that work only contemplates the idea and neither fully defines how coupons would work or simulates their effectiveness in a realistic ad hoc environment. We are particularly interested in examining the network costs that are usually associated with flooding schemes. Although flooding can be effective in terms of distributing data to a large set of nodes, it is known to suffer from several problems such as high message redundancy.

The contribution of this paper is therefore twofold; first to perform a thorough evaluation of the behavior of "coupons" based on various parameters. Secondly to understand under which conditions, if any, "coupons" should be used and what tradeoffs should be expected. Our results show that "coupons" is a simple and elegant solution to the problems associated with network flooding. By deploying a simple scheme, the behavior of the users themselves can provide the answer to finding the desired compromise between scalable data dissemination and network efficiency.

The remaining of this paper is organized as follows. Section 2 presents the "coupons" idea in more detail and explains our motivation. Section 3 gives an overview of existing work in this area and Section 4 offers a discussion of the metrics and parameters that can measure the effectiveness of the scheme. Section 5 presents the simulation results and the paper is concluded in Section 6.

2 System Description and Motivation

The aim of coupons is to allow information distribution to as many users as possible and as efficiently as possible. Since we consider a wireless ad hoc network, there are two main challenges. First, the constantly changing network topology complicates the development of an efficient plan for wide scale data distribution. Second, nodes need to be encouraged through some sort of incentive to relay the information that needs to be broadcasted.

The basic concept of coupons is that it distributes a given piece of information through an ad hoc network using a controlled broadcast. Nodes are encouraged to control the broadcast and to relay the coupon to other nodes using an appropriate motivation. The incentive we propose is an ordered list of unique IDs appended to the end of the message. The idea is that once the information is eventually used, users contained in the ordered list receive some sort of benefit. The benefit can be uniform for all users on the list or can be variable, e.g. the higher a user is on the list the more the value.

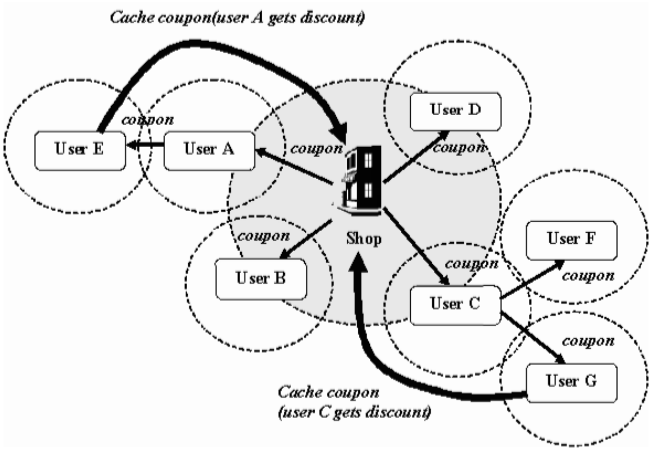

Fig. 1: Coupons operation

The key application scenario we envision is a true coupon where the information distributor is a store trying to advertise its service. A common mechanism for doing this is through discounts for merchandise. Normally users would have little incentive to relay this kind of information, but if they could attach their ID to the coupon then they would receive benefit if someone used the coupon with their ID in the list. Figure 1 shows a typical scenario with a store that is interested in advertising a coupon. As the coupon is received by a number of in-store nodes, users A and C decide to keep it and users B and D to discard it. The coupon is forwarded to other users who come into contact with A and C. When a user visits the store and uses the coupon, discount points will be awarded to all users listed in the ID list. For example if user E goes to the store, then user A could get a 10% discount for up to the value of what E spent.

While one application of coupons is that of stores advertising their products, the basic concept can be more broadly applied. For example, another application might be information dissemination in event-specific scenarios such as the Olympic Games. The coupon could instead be results of the various events. Then the use of signatures would allow a person that helped in the distribution of the information to get "credits". These credits could be used for various rewards such as free tickets, ticket upgrades or special event invitations. Conferences, trade shows, or shopping malls could benefit by adopting the coupon approach to content distribution.

The motivation for developing coupons is that it solves the key problem of *efficient* wide-scale information dissemination. Although basic flooding schemes offer simplicity and high effectiveness in terms of spreading the information to a wide range of participants, there are two main concerns: security and efficiency.

Security in ad hoc networks is still an open area of research[6, 8]. However, for the coupon scheme we have developed a simple but effective security mechanism. Basically, the store transmits its public key, *PK*, and a unique nonce along with the coupon. Then each user in the propagation chain will use PK to encrypt the received coupon along with his ID, the received nonce and a nonce generated by itself. When at a later stage the initiator gets a coupon back, he can retrieve the order of users by using his private key to sequentially decrypt the coupon. This solution protects the integrity of the IDs and their order in the coupon list.

In terms of efficiency, simple flooding is usually associated with the *broadcast storm* problem[3]. Broadcast storms cause three problems: a high rate of redundant messages, an increased probability of collisions, and increased congestion. It is therefore sensible to question whether flooding should even be deployed by an adhoc network, especially in densely populated areas. In order to answer this question as well as quantify the benefits of coupons, the rest of this paper examines how coupons fit with related work and then compares performance via simulation.

3 Related Work

There are three dimensions behind the concept of coupons and therefore three related research areas. First is the notion of providing an incentive in order to stimulate user collaboration. Second, there is work related to the performance of epidemic protocols. Finally, existing literature describing the idea of coupons.

Providing incentives to users has recently attracted much attention[4]. As an example, the Terminodes project[5] specifies *nuglets* as a special type of virtual currency in order to stimulate collaborative packet forwarding. Other approaches, like Sprite[6], adopt a closer link to economical factors by relying on the accessibility of banker nodes to convert from digital to real world money. Finally there are also attempts like the Confidant protocol[7] to address more specific problems such as the detection and isolation of malicious behavior. Our coupons is different in the sense that it targets more lightweight applications with a focus on minimizing complexity.

As this simplicity comes from the adoption of a simple flooding protocol, we also need to examine the implications of the broadcasting problem. Related work in this area refers to *efficient flooding*; mechanisms that aim to control the number of unnecessary retransmissions by having each node make a decision on whether a message should be retransmitted or not. This decision can be based on a number of different heuristics, for example, a randomly chosen probability or retransmitting only when the number of received duplicate packets is below a certain threshold[3]. Although such schemes have been shown to perform more efficiently than simple flooding, they impose a number of deployment requirements. In addition to it's relative complexity, a heuristic works only if there is an incentive for users to participate, and equally importantly if all users are aware of the type of the heuristic and its value. The remainder of this paper shows that coupons can avoid these concerns while at the same time it achieves similar levels of performance.

Finally, the basic idea of coupons has been previously proposed in the research literature in two places: in iClouds and in eNcentive. iClouds[1] is an architecture to enable transparent data exchange through mobile node collaboration. Aiming at scalable and dynamic service advertisement, a similar coupon scheme is described with a coupon being an ID list where nodes append their unique signatures. eNcentive[2] is an agent based framework that provides a communication layer for ad hoc wireless networks. The authors of both papers give good incentives for the use of coupons and they present a middleware architecture to enable this facility. However, although the basic concept is similar, the focus is not on the network aspects of disseminating coupons. No evaluation of the scheme is available. By assuming an agent framework similar to the ones presented in iClouds and eNcentive, this paper describes the idea more thoroughly and discusses its network implications.

4 Simulation Environment

The goal of our simulations is to evaluate specific metrics and examine how well coupons work in different circumstances. Before explaining the results, we first describe our simulation environment.

We have performed our simulations using the GloMoSim network simulator[9]. The coupons scheme was implemented on top of a simple flooding protocol where nodes periodically broadcast the coupons to their surrounding neighbors. Initially the coupon is transmitted by a single stationary node that takes the role of the "store". The parameters used in our simulations are described below and summarized in Table 1:

- **Number of nodes**: number of nodes in the network.
- **Node degree**: the average node connectivity degree. Node degree is calculated as a combination of three parameters: number of nodes, terrain size, and radio range. The association between the degree and the corresponding set of parameters has been calculated through a number of simulations where each node was constantly broadcasting a message to its neighbors. Values range from 0.02 to 8.
- **Coupon levels**: the maximum number of ID signatures each coupon can carry. We assume that a node may broadcast a received coupon only if there is at least one free slot in the ID list. Otherwise there is no incentive to re-distribute the information since there is no potential gain. By experimenting with different levels we evaluate whether there are any major effects of spreading value among fewer or more nodes.
- **Mobility model**: the patterns in which nodes move. All models have been generated by the Bonnmotion mobility model generator [1].
- **Duration**: length of the simulation.
- **Broadcasting rate**: the rate at which the store and participating nodes broadcast the coupon.
- **User behavior**: the way that users behave when receiving a coupon. Users might only cache the coupon or might also start re-broadcasting it.

Parameter	Descrption	Value Range	Nominal Value
N	Number of nodes	100	100
node-degree	Node connectivity degree	[0.02...8]	0.3 and 5
coupon-levels	Length of ID list	2,5,8,10	5
mob-model	Mobility model	RandomWaypoit, ManhattanGrid, Attractions	RandomWaypoint
duration	Duration of simulation	3600 sec	3600 sec
br	Shop broadcasting rate	1 coupon/sec	1 coupon/sec
behavior	User behavior profile	Greedy, Casual, Prob-Greedy, Passive	Greedy

Table 1: Simulation parameters.

Based on this setup we performed our simulations according to the following three metrics:

- **Application efficiency**: measured by both the percentage of nodes that get the coupon and the average time it takes for nodes to first receive the coupon.
- **Network efficiency**: measured by both the total number of transmissions and the average number of times a node gets the same coupon.
- **Incentive efficiency**: measured as the number of nodes that actually contribute to the scheme and more importantly how much benefit each receives. Since the incentive for a user is to have other people cache a coupon with its signature ID in it, effective distribution is largely dependent on the incentive mechanism.

5 Simulation Results

This section presents a representative set of our simulation results. Through our results we hope to achieve three goals. First, we want to understand how the system behaves and the extent of the associated costs. Second, we want to analyze how the first set of results may be affected by the movement pattern of the participating nodes. Finally, we want to examine whether different user profiles affect network costs and application efficiency.

5.1 System Behavior

In this set of simulations we attempt to understand the general behavior of the system. The results for each metric are described below:

Application efficiency. We ran experiments to measure both the total number of nodes that receive the coupon (spreading rate) and the time that was

[1] http://web.informatik.uni-bonn.de/IV/Mitarbeiter/dewaal/BonnMotion/

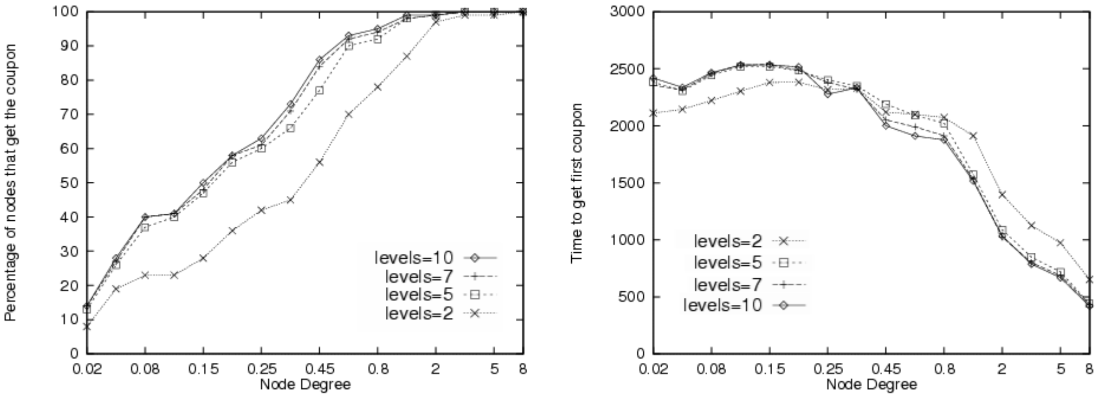

(a) % of nodes that get the coupon.

(b) Average time to receive coupon.

Fig. 2: Application Efficiency

required. A spreading rate of 100% would mean that at some point during the simulation each node received a copy of the coupon. Multiple coupon levels were simulated in order to examine whether more levels would give better results.

Figure 2(a) shows the spreading rate and Figure 2(b) the average time to first receive the coupon. In both diagrams the x-axis shows the node degree, a combination of the number of nodes, terrain size and radio range.

Overall, Figure 2 shows that high node degree results in faster and more thorough coupon spreading. With a node degree of 0.45 and above we can achieve a spreading percentage of 80% to 100%. Concerning the coupon levels, the more levels, the faster the coupon distribution. Moreover, for the cases where node degree is between 0.1 and 2, more levels also results in wider distribution.

Network efficiency. In this set of experiments we measured the network and node load. Figure 3(a) shows the network load as the count of the total number of coupon broadcasts. Node load is shown in Figure 3(b)as the average number of times each node gets the coupon. Both diagrams have the same x-axis (node degree) and both show results for multiple coupon levels.

Overall the cost of distribution is closely related to the node degree. In terms of coupon levels, as we move to a more densely populated environment, higher coupon levels are associated with much higher cost. Our experiments have also shown that there seems to be no relationship between the number of times a coupon is seen and when the node sees the coupon for the first time. The reason is because when a node sees a coupon for the first time and how many times it sees it thereafter is almost exclusively affected by the mobility pattern, i.e. who the node comes into contact with during the simulation period.

To summarize, the basic behavior of coupons is analogous to that of any normal epidemic network: high spreading efficiency is contrasted by high network

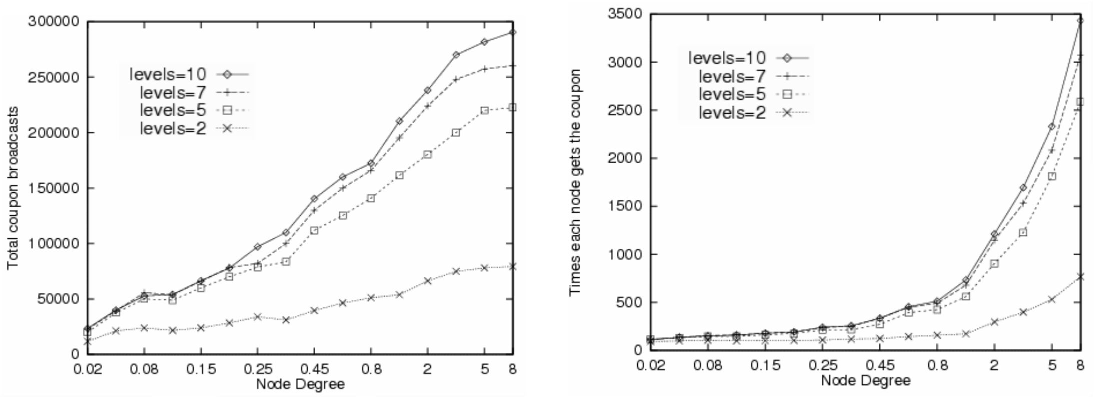

(a) Total broadcasts

(b) Average times each node gets the coupon

Fig. 3: Network Efficiency

costs. Our experiments show that one way to reduce the network overhead is by using fewer *coupon levels*.

Incentive efficiency. From the previous graphs it seems that the only benefit of having more coupon levels is that they achieve higher spreading rates. And in densely populated environments (i.e. high node degree), spreading is already quite good. However, there is one more important gain and this is related to the efficiency of the incentive.

With the *Incentive efficiency* we try to see how many nodes actually contribute to the scheme. We consider coupons to be more efficient when it is not one node that simply visits all others and passes the coupon but when the distribution is more evenly balanced.

Figure 4(a) shows a log scale of the awarded points (y-axis) over the population of nodes (x-axis). A highly populated area is assumed (node degree=5) since it gives almost identical spreading efficiency for all coupon levels. As can be observed, only a small percentage of nodes actually benefit when we use only two coupon levels. It is also interesting to note that the more coupon levels we use, the higher the awards for the benefited nodes are. So with more levels we get both higher participation rate and higher awards for those who collaborate, thus higher incentive.

Figure 4(b) highlights another important aspect of the system. By mapping the awarded points (y-axis) to the time the first coupon was received (x-axis), we see the nodes with the highest probability of receiving many reward points are among those that are first to receive the coupon from the store.

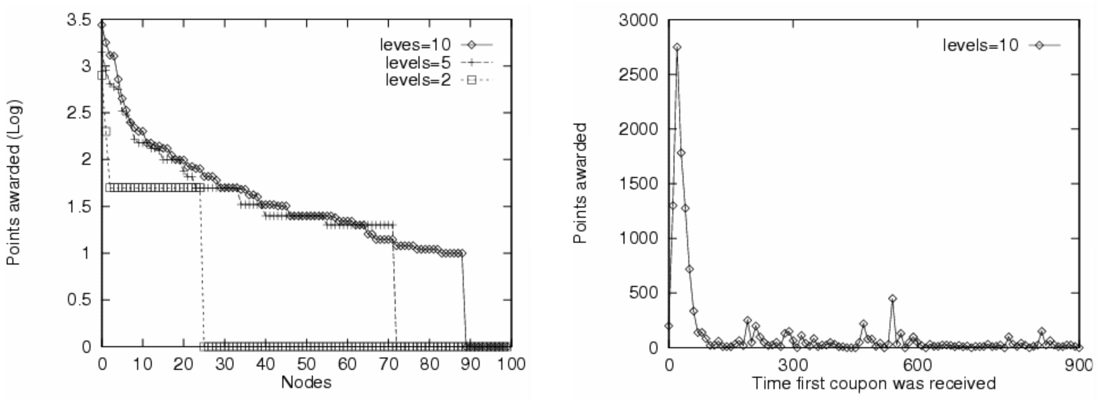

(a) Credit distribution.

(b) Credit vs. when coupon received.

Fig. 4: Incentive Efficiency

5.2 Effect of Mobility

So far we have seen how the system is generally expected to behave from an abstract level. This section investigates whether different mobility patterns have an impact on the performance and whether the scheme should be adopted according to the different models. For this reason we have considered three different models:

- **RandomWaypoint**. This is the model that has been used so far. Each node chooses new destinations uniformly distributed from the simulation area and moves towards them.
- **Attractions**. A RandomWaypoint model but with a number of attraction areas (four have been used in this paper, located near the corners of the terrain). Instead of uniformly selecting a new destination from the simulation area, nodes have a tendency to move from one attraction area to the other.
- **ManhattanGrid**. In this model nodes move only on predefined paths. In our experiments 100 by 100 blocks have been specified.

We run similar experiments to those that have already been explained in the previous two sections. Figure 5(a) shows the spreading rate. The x-axis serves as an indicator for the different coupon levels while a single node degree equal to 2 is used. As ManhattanGrid indirectly affects the node degree, it seems to perform better. However there are no major differences between the models. The same observation can be made in Figure 5(b) as Manhattan Grid poses higher network cost in terms of the total messages. The only lesson is that when coupons are to be applied in a city environment (or any other similar scenario with high node degree), fewer coupon levels should be used as opposed to an open field case.

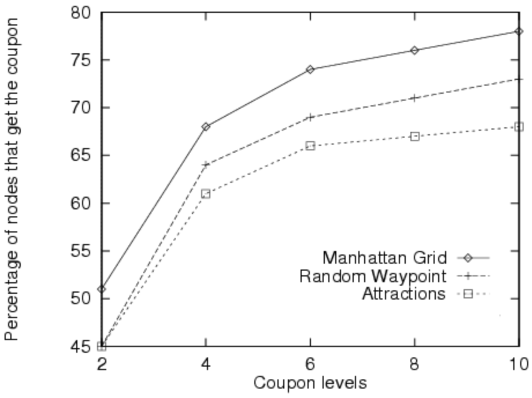

(a) Percentage of nodes that get the coupon

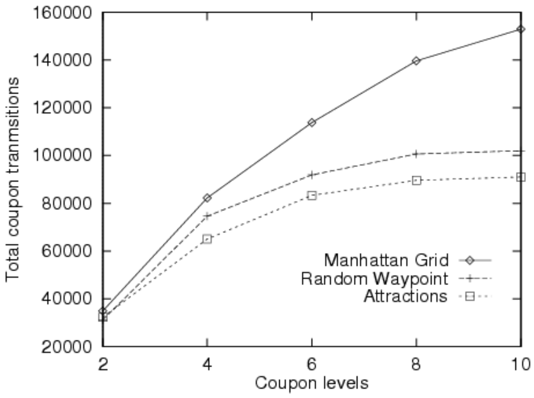

(b) Total transmissions

Fig. 5: Effect of Mobility Patterns

5.3 Considering User Profiles

So far we have seen that coupons can be very efficient in terms of propagating information, but at the same time it has significant overhead. There is a high number of redundant transmissions and nodes seem to receive a large number of duplicate messages. However, what we have not yet discussed is user behavior.

All experiments up to this point have made a simplistic assumption: all nodes forward coupons all the time (assuming they find an empty slot in the ID list). For the rest of this paper we will refer to this type of user behavior as *Greedy*. What is now important to see is how the system behaves when only a percentage of nodes are *Greedy*, and when different user behavior is possible. For this purpose we have implemented and tested the following user profiles:

- **Greedy**. Once a node gets a coupon, if there is at least one free level in the ID list, it broadcasts it once each second.
- **Casual**. This type of user wants to keep the coupon, but is not interested in distributing it.
- **Prob-Greedy**. Similar to *Greedy* but with one important distinction. There is now a probability function to decide whether a coupon is broadcast or not. This works by taking into consideration the number of free levels in the coupon and the time elapsed. These two metrics attempt to capture the essence of the user's incentive to redistribute the coupon (more free levels imply higher incentive) and the concern for battery consumption.
- **Passive**. A node broadcasts in a "passive" fashion. This implies that it distributes a coupon only if no other neighbor transmit at the same time. The motive for this model is to preserve energy and distribute without directly competing with others. We have implemented this type by recording when

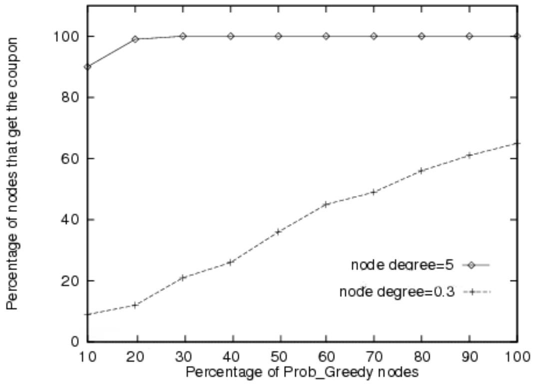

(a) % of nodes that get the coupon vs. profiled node density.

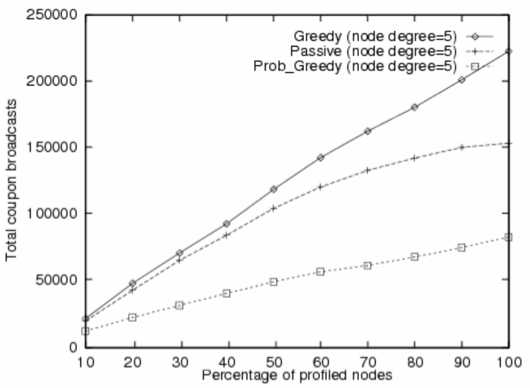

(b) Total broadcasts vs. the density of different profiled nodes.

Fig. 6: Different User Profiles

the last received coupon was received. A broadcast takes place only after not hearing the same coupon broadcast in the last 10 seconds.

By considering all these profiles, we evaluated each metric by assuming that all nodes are at least *Casual*. Then we simply varied the percentage of the other three models. The main results are available below:

Application efficiency. Two important observations can be made. First, in terms of spreading the coupon there is little difference between the models. This is an important issue since it indicates that application efficiency does not require participating nodes to always operate in the *Greedy* mode. The second observation concerns the importance of how many nodes are actually willing to participate, or in simulation terms, how many nodes have one of the three forwarding types (*Greedy*, *Passive* or *Prob_Greedy*). For the sake of clarity we call these three types of nodes "profiled nodes".

Figure 6(a) shows how different percentages of profiled nodes perform with regard to spreading rate. In this graph we show two population density scenarios: low density (node degree=0.3) and high density (node degree=5). In the low density scenario, the spreading rate increases proportionally with the increase in greedy nodes. In other words, only if there are enough Prob_Greedy nodes will there be significant coupon spreading. In the second case however, we observe that even with a 10% of Prob_Greedy nodes, we can achieve a 100% distribution of the content. This is an important point since it highlights the potential of coupons: in densely populated environments, only a few nodes are required to have an incentive.

Network efficiency. An obvious benefit of having only a small percentage of nodes forward the coupon is that network cost is reduced. This is shown in Figure

6(b) where with 10% of profiled nodes (either Greedy, Passive, or Prob_Greedy) there is a saving of approximately 90% compared to the case when all nodes are profiled nodes. From the same graph we can also see that even in the case where all nodes are "profiled", the more realistic profiles perform much better. Compared to Greedy, sProb_Greedy and Passive can reduce the cost by 64% and 32% respectively without significantly affecting the percentage of nodes that eventually receive the coupon.

Incentive efficiency. Incentive efficiency results so far have assumed that all nodes are of the same type. An interesting question is how the different profiles perform in terms of incentive efficiency. In order to evaluate this aspect, we performed another set of experiments with a mix of profiles: 10% *Greedy*, 10% *Passive*, 10% *Prob_Greedy* and the remaining 70% *Casual*. Our results showed that none of the user profiles had a clear advantage. This might be surprising at first since we might have expected *Greedy* to earn most of the credit. The reason is that what matters most is not how a user behaves but the ability to be among the first nodes that receive the coupon from the store and then to move frequently reaching new nodes more often.

Our belief is that users will gradually learn that there is no need to be greedy in order to earn a large bonus. Given that battery power will be an important concern, most users will default to rational behavior like the Casual or Prob_Greedy profiles. In the end, the system will seek its own equilibrium and balance effectiveness and efficiency.

6 Conclusion

The goal of this paper has been to develop applications and services that are uniquely suited to ad hoc environments. "Coupons" is a scheme that expands on the basic idea of information dissemination in ad hoc networks by investigating incentives for collaboration. Cooperation is stimulated by adopting a basic pyramid scheme where users are awarded credits as they pass a received piece of information (a *coupon*) to other people.

Although the basic concept of a coupon has been previously proposed, that work only contemplated the idea and did not simulate its performance in an ad hoc environment. The contribution of this paper is therefore twofold. First, we have performed a thorough evaluation of the behavior of coupons based on various parameters. Second, we have examined under which conditions coupons could be used and what tradeoffs should be expected.

Our results have shown that our scheme has similar properties to epidemic protocols. High spreading efficiency can be achieved but this advantage is offset by the disadvantage of high inefficiency. The important distinction of our work comes from the observation that the use of a reasonable user profile helps control distribution. Spreading effectiveness is high while at the same time controlling distribution costs (efficiency). This is more evident in a densely populated environment. Our results have shown that with only 10% of nodes being Prob_Greedy, we achieve the same spreading rate with 90% fewer messages. Effectively we achieve the performance of an *efficient flooding* protocol, but without the complexity or need for a heuristic.

Coupons can be a simple and elegant solution to the problems associated with network flooding. By deploying a simple scheme that actually takes advantage of user behavior, we can create a desirable compromise between scalable data dissemination and network efficiency.

References

1. A.Heinemann and J.Kangashrju, "Ad Hoc Collaboration and Information Services Using Information Clouds", In Proceedings of the *International Workshop on Applications and Services in Wireless Networks (ASWN)*, Berne, Switzerland, July 2003
2. O.Ratsimor, T.Finin, A.Joshi and Y.Yesha, "eNcentive: A Framework for Intelligent Marketing in Mobile Peer-To-Peer Environments", In Proceedings of the *International Conference on Electronic Commerce (ICEC)*, Pittsburgh, PA, September 2003
3. S. Ni, Y. Tseng, Y. Chen, and J. Sheu, "The broadcast storm problem in a mobile ad hoc network", In Proceedings of the *ACM/IEEE International Conference of Mobile Computing and Networking (MOBICOM)*, Seattle, Washington, August 1999.
4. P.Orbreiter and J.Nimis, "A Taxonomy of Incentive Patterns - the Design Space of Incentives for Cooperation", Technical Report Nr 2003-9, University of Karlsruhe, Germany, May 2003
5. L.Butty and J.Hubaux, "Stimulating Cooperation in Self-Organizing Mobile Ad Hoc Networks", In *ACM/Kluwer Mobile Networks and Applications (MONET), Special Issue on Mobile Ad Hoc Networks*, Vol. 8 No. 5, October 2003.
6. S.Zhong, J.Chen, and Y.Richard Yang, "Sprite: A simple, cheat-proof, credit-based system for mobile ad-hoc networks", In Proceedings of the *IEEE Infocom*, San Francisco, CA, April 2003.
7. S. Buchegger and J. Le Boudec, "Performance Analysis of the CONFIDANT Protocol: Cooperation Of Nodes - Fairness In Dynamic Ad-hoc NeTworks", In Proceedings of the *IEEE/ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHOC)*, Lausanne, CH, June 2002.
8. L.Zhou and Z.J.Haas, "Securing Ad Hoc Networks", In *IEEE Network Magazine*, Vol. 13, No. 6 November 1999
9. X.Zeng, R.Bagrodia and M.Gerla, "GloMoSim: A Library for Parallel Simulation of Large-Scale Wireless Networks", In Proceedings of the *ACM Workshop on Parallel and Distributed Simulation (PADS)*, Banff, Canada, May 1998