OpenMobs: Mobile Broadband Internet Connection Sharing

Nicolae-Valentin Ciobanu[1], Dragos-George Comaneci[1], Ciprian Dobre[1], Constandinos X. Mavromoustakis[2], and George Mastorakis[3]

[1] University Politehnica of Bucharest, Faculty of Automatic Control and Computers, 313, Splaiul Independentei, 060042, Bucharest, Romania
nicolae.ciobanu1411@cti.pub.ro, dragos.comaneci@cti.pub.ro
ciprian.dobre@cs.pub.ro

[2] University of Nicosia, Department of Computer Science, 46 Makedonitissas Avenue, 1700 Nicosia, Cyprus
mavromoustakis.c@unic.ac.cy

[3] Technological Educational Institute of Crete
Estavromenos 71500, Heraklion, Crete, Greece
gmastorakis@staff.teicrete.gr

Abstract. We witness an explosion in the number of applications being developed for mobile devices. Many such applications are in need or generate a lot of Internet traffic, and as such mobile devices are today equipped with more networking capabilities, from mobile broadband (3G/4G) to WiFi, Bluetooth, and others. However, when it comes to mobile broadband Internet access, for economic reasons, today mobile providers tend to switch from unlimited mobile data plans to tiered data pricing models, putting pressure on mobile data subscribers to be more careful how they consume their subscribed traffic. In this paper, we propose OpenMobs, a mean to reduce the costs associated with mobile broadband access to Internet, by sharing under-utilized networking resources among co-located users through free wireless connections. When two or more handsets are in the vicinity of each other, OpenMobs forms an ad hoc mesh network to redirect traffic between mobile data plan subscribers, in the most economic and viable way. We present studies on the feasibility of such a system to minimize the costs users pay monthly to their mobile providers, and even financially compensates users' willingness to participate in the collaboration.

Key words: Internet, mobile broadband, sharing, traffic allocation

1 Introduction

Today mobile handsets and devices have come to outnumber traditional PCs several times, and applications for mobile devices exploded in number. But many such applications are in need or generate a lot of Internet traffic. For example, to address problems associated with the limited amount of resources (computation, storage, power) available within the mobile device, or to provide richer experience to their clients, many application developers appeal to resource providers (the 'Cloud') other than the mobile device. This is why today we do have various mobile applications connected to either Apple iCloud, Google's Gmail for Mobile, or Google Goggles. Of course, for this to happen, application developers rely on good networking connections with the Cloud.

The networking capabilities offered by mobile devices have become very diverse lately. Internet access options range from using free Wi-Fi at a hotspot, to having a mobile broadband (e.g., 3G) or a mobile hotspot access (the "anywhere, anytime" Internet access offered over cellular networks). Among these, mobile broadband access is still widely used, since it allows the user to go online anywhere there is a cellular signal.

For mobile broadband access, a mobile data plan from a cell phone provider allows a client to access the 3G or 4G data network, to send and receive emails, surf the Internet, use IM, and so on from his mobile device. Mobile broadband devices such as mobile hotspots and USB mobile broadband modems also require a data plan from a wireless provider.

Unlimited data plans for cell phones (including smartphones) have been the norm most recently (sometimes folded in with other wireless services in a one-price subscription plan for voice, data, and texting). Still, today most providers, following the example set by AT&T in 2010 [11], use tiered data pricing, thus eliminating unlimited data access on cell phones. Tiered data plans charge different rates based on how much data the client uses each month. The benefit is that such metered plans discourage heavy data usage that could slow down a cellular network. Thus, it is no wonder today that most mobile broadband plans for data access on laptops and tablets or via mobile hotspots are typically tiered [10]. The downside is that users have to be more vigilant about how much data they are using, and for heavy users, tiered data plans are more expensive.

When it comes to choosing a suitable tiered mobile data plan [4], clients generally tend to go for oversized mobile data plans. For choosing, clients estimate their peak monthly traffic needs, which is natural considering that mobile operators charge the extra traffic above the data plan limits. In our work, we started by analysing this fact, through interviews and questionnaires, and found out that *today most clients do tend to pay for a lot of mobile broadband traffic, but most of the time they never use their entire payed data plan traffic.*

On the other hand, when users (accidentally) exceed their mobile plan rates, they are generally charged extra by the mobile provider. Also, in roaming, the extra costs for connectivity can be occasionally quite prohibitive. So, the research question we are addressing in this article is: *Can we come up with a solution that mediates opportunistic sharing of networking resources, when needed, between users?*. For the sharing of WiFi Access Point traffic, opportunistic networking today provides an answer [2]. For mobile broadband access, we want to let users share parts of their unused mobile broadband traffic with others. But users pay a monthly fee to their mobile providers, so they might be reluctant in 'giving away' traffic to others, for free. Thus, we propose letting the user become a *re-seller of broadband traffic that he gets from his mobile provider*. This means that a user can sell part of his under-used traffic, and sell it to clients in need, making a small profit in doing this (such that, at the end of the month, some of his monthly mobile data plan fee gets payed by others). For a buyer, it is attractive to have other users let him use their mobile data broadband access, if he ends up paying less compared to the fees charged by the mobile provider.

In the present work we present OpenMobs, a system designed to support the sharing of under-utilized resources available on mobile handsets in a distributed and opportunistic way. When two or more handsets are in wireless proximity, OpenMobs tries to forward part of one user's traffic through the mobile data plan of the other. To incentivize the payments between users, and motive them share resources in an accountable manner, a digital currency such as BitCoin [9] can be used as the form of payment for used resources. Here, we present extensive studies on the feasibility of such a system to minimize the costs users pay to their mobile providers at the end of the month.

Internet connection sharing has existed as an idea for many years, and every modern operating system has implemented its fair share of services in order to address it [12]. For smartphones, the latest venture that comes close to what the current work is trying to solve is Open Garden [6]. Open Garden leverages crowdsourcing to create seamless connectivity across 3G, 4G, Wi-Fi and Bluetooth. It enables users to create their own ad-hoc mesh networks with other Open Garden enabled devices (i.e., smartphones, tablets and PCs). Unlike the centralized idea proposed by Open Garden, OpenMobs allows users to share networking resources with minimal interaction with a centralized entity. OpenMobs tackles the problem of automatic sharing based on context, where the user node automatically chooses to use the opportunistic shared connection.

Our work can also be compared with the idea of offloading cellular networks through ad hoc vehicular wireless networks [8]. However, we do not rely only on the existence of wireless routers (it is even better when such devices exist), and optimize traffic consumption particularly considering the wireless mobile broadband charging fees. To the best of our knowledge, this is the first work to propose such a decentralized traffic sharing approach.

The rest of the paper is structured as follows. In Section 2 we first introduce the theoretical optimization problem linked to the allocation of networking resources, and propose a heuristic allocation approach for the maximization of compensation costs. Our approach is further evaluated in extensive simulation experiments in Section 3. Finally, in Section 4 we present the conclusions.

2 Traffic Allocation Model

2.1 Research Problem

The problem of sharing traffic can be modelled as follows: Given a set of users, each having a mobile data plan (defined by a data amount for which the user pays a monthly fee, and a cost model for computing an extra fee associated with the traffic consumption exceeding the data plan threshold set by the cell phone operator), and each consuming a certain amount of traffic each month, we want to find an equilibrium price auction model for bidding traffic between users, such that by re-routing traffic through other mobile phones and paying a fee for the temporary use of their data plans, users gain profits and/or pay less, compared to the case when each user acts selfish sticking to only the local data plan costs negotiated by each with their cell phone operators.

In this problem, any user can become a traffic provider for other users (*seller*). The traffic is auctioned, and any user is interested to buy traffic from other users (*buyer*) directly located in his wireless communication range (WiFi, Bluetooth or ZigBee could be employed at no extra costs), if the price is lower than the cost associated with sending the traffic over the 3G or 4G data network. As mentioned, such a situation appears, for example, when a buyer already consumed its entire monthly data plan traffic, and any extra traffic might be charged by the mobile operator at considerable higher fees. Or, when a buyer is in roaming, and the cost of transferring data over the mobile operator can be considerable higher compared to the costs negotiated with the mobile operator by another user (i.e., local to the mobile network).

This situation is illustrated in Figure 1. In the example, $user_B$ needs to transfer some data (send and receive emails, surf the Internet, use IM, and so on) from its mobile device. For this, he can send data over 3G, using the cost associated with the data plan negotiated with the mobile provider ($cost_B$). Luckly, in his wireless communication range, $user_A$ is offering to transfer this data, over a WiFi connection existing between these two users, at a cost (bid_{AB}) lower than $cost_B$ (so $user_B$ actually pays less, the difference being his 'gain'). For $user_A$, this situation also brings a small profit ($gain_A$), since the offered cost bid_{AB} is higher than the actual cost ($cost_A$) negotiated by $user_A$ with her mobile provider for transferring this data.

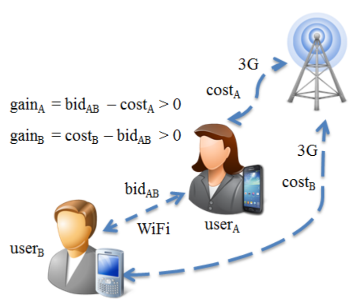

Fig. 1: Example scenario.

When applying equilibrium price auctions for the allocation of traffic, the two roles, buyer and seller, face distinct yet linked challenges. The buyer is interested to transfer traffic at the minimum cost possible, while the seller wants to maximize his profit (he will commonly pursue the objective of maximizing profit). Thus, we need to come up with specific equilibrium prices each time a user is interested to transfer some data, such that to avoid situations where a seller loses money at the end of the month, by selling traffic too cheap, compared to the cost he has to pay for his own traffic transfer needs. In this case, the profit is given by the difference between the revenue from the served bids, and the costs associated with all transfers over the mobile data network.

For the rest of this article, we assume that each user has a mobile data plan, negotiated with a local mobile provider/operator. Also, each user consumes a certain amount of traffic, each month, for his own personal needs (transfers generated from the local mobile phone, for emails, web, and others). We further assume that each user can participate in any auction, with any other users having different data plans and traffic needs.

2.2 Formal Notations

As a basis for the optimization approach presented in the following, we introduce a formal notation. First, we define the basic entities:

- $U \subset \mathbb{N}$ Set of participating users.
- $B \subset \mathbb{N}$ Set of buyers (willing to buy traffic from other users, where $B \subseteq U$).
- $S \subset \mathbb{N}$ Set of sellers (wanting to sell traffic to other users, where $S \subseteq U$).

According to the mobile data plan of user $u \in U$, he can monthly transfer traffic up to a specific amount (DP_u), at a constant fee (CF_u - generally negotiated with the mobile provider). If the user exceeds the DP_u limit, the extra traffic is charged separately by the mobile operator, according to a cost algorithm $CO_u(T)$, that depends on the actual amount of traffic T being transferred.

Generally, the data plans differ between users. Thus, for any 2 users u_1 and u_2, it can happen that $DP_{u_1} \neq DP_{u_2}$, $CF_{u_1} \neq CF_{u_2}$ and so on (but, it can also happen that two users can get similar data plans, especially when they are subscribed to the same mobile operator).

Also, we assume that each user transfers a certain amount of traffic, T_u each month. Naturally, there is no problem when all users manage to stay within their mobile data plans ($T_u \leq DP_u$) - but this is not always the case because of at least two reasons: (1) even if the user manages to provision a data plan satisfactory to his need, unscheduled events might actually lead to more traffic being generated than usual (*unforeseen networking needs*), and (2) whenever in roaming, or for different services provided by the mobile operator at 'extra costs', there are supplementary fees for the generated traffic, independently of the traffic included within the data plan (*extra costs*). We assume that other users are willing to sell traffic (actually, accept connections and transfer data for other users, through their local mobile data plan), because they do not manage each month to consume all the traffic they generate anyway (which is generally the case, as users tend to negotiate mobile data plans with mobile operators that are above their average traffic needs - the old saying 'better safe than sorrow' generally applies when we think what kind of mobile data plan is suitable to our monthly traffic needs).

Whenever two users, a buyer ($b \in B$) and a seller ($s \in S$) are in contact (they can exchange traffic at no costs over short- or medium- range wireless protocols, such as WiFi, Bluetooth or ZigBee), they can negotiate a price (CS_{bs}) for transferring data generated by b, over the wireless link to s, and from there over the mobile data network, using the seller's data plan. This price also depends on the amount of traffic T_{bs} that is transferred between these two users (otherwise, the buyers might transfer well above the data plan limits of the seller, which might reduce the seller's profit).

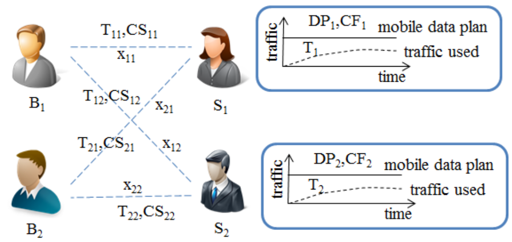

Fig. 2: Schematic overview of the optimization model, depicting the decision variables, and most relevant entities.

Finally, we assume users are mobile, such that with a high probability, over a longer period of time, any two users can meet at least once (such that $\forall b \in B, s \in S \quad T_{bs} \geq 0$ stands).

2.3 Optimal Allocation Approach

To compute an optimal solution to the indicated problem, we transfer the problem definition into a mathematical optimization model. The result is given in Model 1, and will be explained in detail in the following.

In Equation 6, x is defined as a binary decision variable. Specifically, x_{bs} indicates whether the user b is willing to buy traffic from another user s. More specifically, when two users meet, each one presents his offer (the price he is willing to accept for the other to transfer traffic through his mobile data plan). Of course, this auction can be accepted or not by each user (which is interpreted by this binary decision variable). An overview of the optimization model, which highlights the relations between the decision variables and the most important entities, is depicted in Figure 2.

Equation 1 specifies the objective of the optimization model, namely the minimization of the costs implied for transferring cell phone generated data over the mobile networks operated by different mobile providers (but doing it such that in the end no user loses money). This is similar to the concept of profit, as the difference between the fees payed to the mobile operator without using any optimization, versus the smaller fees payed to the mobile operator when users collaborate and re-sell traffic in their data plan, is a cost which is shown in their own pockets. Thus, the two components shown in equation 1.

Equation 2 specifies the cost for transferring data for all users, as a sum between the data plan costs (CF, fix costs payed by all user at the end of the month; this generally includes the traffic D, corresponding to the data plan negotiated with the mobile operator), and costs associated with extra traffic (CO is a cost model specified by the mobile operator; according to this, generally the user pays proportionally with the amount of traffic transferred over the data plan limits $T^{'}$). The cost CO can be quite large, and in this case, the $T^{'}$ traffic could be redirected through the unused data plan offered by another user.

Equation 3 specified exactly this optimization. In this case, some traffic is transferred through other users (so we have sellers S, and buyers B). In this case, the cost is a sum between (1) the cost for transferring the local traffic by the buyer, through his local mobile provider, (2) the cost negotiated for traffic auctioned between the buyer and seller, and (3) the cost for transferring the local traffic plus the negotiated traffic, through the local mobile provider of the seller. In this equation, the cost of transferring data over the mobile operator depends, again, whether the traffic is included in the monthly data plan traffic, or exceeds the data plan traffic (which is presented in equation 4).

Finally, equation 5 links everything together, and presents the conditions needing to happen in order for a user to maximize his profit by selling data plan traffic (respectively, optimize the cost by buying traffic from a seller, at a cost lower than the one offered by the mobile operator). For users to gain from this collaboration (see Figure 1), two conditions must simultaneous stand: (1) the buyer must buy at a cost smaller than the cost of transferring the same amount of data through the mobile operator (right inequality), and (2) the seller must sell at a cost that covers at least the cost necessary for him to transfer the sold data over the local mobile operator (left inequality).

For the seller, the cost could be 0 if he manages to sell traffic included in his mobile data plan, negotiated with the mobile operator. This condition is also captured in equation 5; in this case, his operating/transferring cost C_s is kept to 0 if the traffic generated by the seller and transferred over the mobile operator (T_s), plus all traffic that he manages to sell, is still less than the data plan traffic negotiated with the operator (DP_s). Actually, the seller still has to pay at the end of the month the fee associated with his data plan, but this is independent of the amount of traffic sold (and is included in $Cost_{initial}$). In other words, C_s establishes a hint over the *profit* the seller manages to accomplish.

In equation 5, the selling price needs a predictor on the amount of traffic the seller expects to deliver through his mobile network for the current month. A fog-of-war probabilistic model, similar to the one proposed in [8], can be used to deal with this uncertainty, considering the history of traffic associated with the user monthly, on a historical base. With this construction, Model 1 still constitutes a Linear Program (LP), or more specifically, Binary Integer Program (BIP). This class of optimization problems can be solved using well-known methods from the field of Operations Research, most notably, the Branch and Bound (B&B) algorithm [7]. While the B&B algorithm can be very efficient in some cases, it is still based on the principle of enumeration, i. e., in the worst case, all potential solutions have to be examined [5]. Specifically, for a BIP, the solution space grows exponentially with the number of decision variables. As can be observed from Model 1, the number of decision variables increases quadratically with the number of traffic auctions (equation 6), and linearly with the number of mobile data plan types. Accordingly, the computational complexity of the optimal allocation approach is exponential and corresponds to $O(2^{\|U\|^2 * \|DP\|})$, where $\|DP\|$ is the number of different tiered mobile data plans in use.1

Model 1 Optimal Allocation Model

$$\text{Maximize} \quad Profit_x = \{Cost_{initial} - Cost_{optim}(x)\} \geq 0 \tag{1}$$

$$Cost_{initial} = \sum_{u \in U} CF_u + \sum_{u \in U} CO_u \left(T_u^{'}\right) \tag{2}$$

$$\begin{aligned} Cost_{optim}(x) &= \sum_{b \in B} C_b \left(T_b - \sum_{s \in S} x_{bs} * T_{bs}\right) + \\ \sum_{b \in B, s \in S} & x_{bs} * CS_{bs}\left(T_{bs}\right) + \sum_{s \in S} C_s \left(T_s + \sum_{b \in B} x_{bs} * T_{bs}\right) \end{aligned} \tag{3}$$

where,

$$C_u(T) = \begin{cases} CF_u & \text{if} \quad T < DP_u. \\ CF_u + CO_u\left(T - DP\right), & \text{otherwise.} \end{cases} \tag{4}$$

$$\begin{aligned} &\forall b \in B, s \in S, \quad x_{bs} = 1 \iff C_s\left(T\right) < CS_{bs}\left(T\right) < CO_b\left(T\right) \\ &with, \\ &C_s(T) = \begin{cases} 0, \quad \text{if} \quad T_s + \sum_{\forall u, u \neq s} x_{us} * T_{us} + T < DP_s. \\ CO_s\left(T\right), \quad \text{otherwise.} \end{cases} \end{aligned} \tag{5}$$

$$x_{bs} \in \{0, 1\} \quad \forall b \in B, S \in S \tag{6}$$

2.4 An Heuristic Allocation Approach

For real-life application scenarios involving thousands of users, the optimal allocation approach may be problematic due to its exponential growth in computational complexity. Thus, we have developed a heuristic approach that trades reductions in computation time against potentially sub-optimal solutions. The idea is to determine an equilibrium price auctioned between any two users.

In our approach, whenever two users, A and B, meet, each presents to the other a price he is willing to accept for traffic forwarding. This means that A computes a price, CS_{AB}, he is willing to accept from B (per data unit). If user B needs to transfer data (for email, or others), he decides whether is cheaper to transfer it through the mobile network (3G), or send it through A (over WiFi or other 'cost-free' wireless communication protocol). In this case, A gains a small fee, which is still larger than what it costs him to actually send the data coming from B, over A's mobile network. If this is true, than user A becomes the 'seller', and user B the 'buyer'.

The CS_{bs} price depends on several parameters (as described in Model 1): the mobile data plan of the seller (DP_s), the traffic already used from this data plan by the seller from the beginning of the current month[1] (P_s), and the amount of traffic the buyer is interested to transfer (XP_b). The idea is to sell cheap when the user has plenty of traffic left from the mobile data plan (such that to guarantee that at least someone buys - the seller wants to maximize his profit, and use the traffic remaining in the data plan that otherwise would be waste), and sell at a high rate if the seller does not have much traffic left in the mobile data plan (such that, in the unfortunate event that in the future he will also want to use his data plan for own traffic needs, the higher fees operated by the mobile provider for any extra traffic are still covered by the fees he gains from his traffic buyers - the seller wants to stay 'in profit', and not lose money at the end of the month).

Thus, the heuristic formula we propose for computing the cost is:

$$CS_{bs} = e^{\min\{th, \frac{P_s + XP_b - DP_s}{coef}\}} \tag{7}$$

where th is a high upper-limit threshold (that ensures the negotiated fee does not grow indefinitely), and $coef$ is a coefficient that reflects the mobility environment. This means that for $coef$, we start with a predefined traffic value (the predefined preference of the user to sell traffic). If, at the end of the month, the user losses money (because his preference in selling made him sell cheaper that the is charged by the mobile operator), the value of this coefficient doubles. After several iterations, as the experiments presented next show, the system actually reaches a state of ecquilibrium, and all $coef$ are stable and individually defined such that we have a positive profit.

3 Evaluation

3.1 Approach and Methodology

For testing the proposed heuristic allocation approach, we used three publicly-available mobility traces. UPB [2] is a trace taken in an academic environment at the University Politehnica of Bucharest, where the participants were students and teachers at the faculty. It includes Bluetooth and WiFi data collected for a period of 64 days, by 66 participants. St. Andrew [1] is a real-world mobility trace taken on the premises of the University of St. Andrews and around the surrounding town. It lasted for 79 days and involved 27 participants that used T-mote Invent devices with Bluetooth capabilities. Finally, MIT Reality [3] contains tracing data from 100 users from the University of Helsinki. The collected information includes call logs and Bluetooth devices in proximity, collected over the course of an academic year. Thus, each scenario emulates different running conditions: users meet scarcely, regularly, or frequently.

On top of these traces, we simulated the behavior of users wanting to transfer data, with and without a system serving the proposed optimized heuristics. This

[1] Through the paper, a month is the time period usually charged by the mobile operator. As such, a month can actually begin with any day of the montly calendar.

system detects opportunistic wireless connections and tries to optimize the cost when possible, by transferring some of the traffic exceeding the mobile data plan, through other proximity-located users.

The main benefit of such a system is that it creates a virtual market of Internet traffic. Basically, it addresses the problem of rigid mobile subscription plans by providing opportunistic ad hoc sharing of traffic, below the normal rates. Those with excess traffic included in their mobile subscription plan can offer it to others for a much lower price than they would pay the network provider.

There is also a hidden benefit for the mobile network provider: the network usage can be more predictable since users will tend to use up their mobile subscription plan completely; at lower costs, users will be stimulated to use broadband communication more, which will result in profit for the provider.

Another benefit is present when travelling abroad since the rates offered by the local peers will be much less than the roaming rates offered by the mobile network provider.

Table 1: Data Service Plan parameters used in simulations.

Price/MB (EUR) for traffic included in plan	Price/MB (EUR) for extra traffic outside the data plan	Amount of traffic included in data plan	Assignment probability
0.02	0.01	90 MB	0.05
0.012	0.01	200 MB	0.05
0.008	0.01	350 MB	0.10
0.007	0.01	450 MB	0.15
0.006	0.01	500 MB	0.30
0.004	0.01	1 GB	0.25
0.0037	0.01	1.5 GB	0.05
0.0032	0.01	2.0 GB	0.05

In these experiments, we were particularly interested whether there are benefits (profit) for the users, and whether users interact frequently enough in real-world so that the system is useful. We want to measure the benefits previously mentioned, and see if the degree of interaction that occurs between peers increases the benefits provided by the system.

Each of the cases was run several times on each trace, with varying random seed values, for a confidence level of 95%.

In order for the simulations to resemble real usage as close as possible, we have chosen to allocate each mobile device from the trace a specific service subscription data plan, which has three associated parameters: price/MB when the user still has traffic in his normal data plan, price/MB when the user has exhausted the traffic in his data plan and the amount of included data traffic. Each data service plan also has an associated probability that is taken into account when generating the data plan associations for the mobile devices.

After associating a data subscription plan, the next step was to designate a level of data usage and traffic pattern per device. A traffic pattern has two associated parameters: the average amount of traffic used per month and a propensity to use that traffic when in a social context (when the device is in contact with other devices). The propensity parameter is represented as the probability that a device will consume traffic when in the presence of another device.

Table 2: Traffic usage pattern parameters used in the simulation.

Average Data Traffic per Month (MB)	Propensity to use traffic while in the presence of other mobile devices	Traffic Pattern Distribution Probability
200	0.5	0.25
400	0.5	0.25
600	0.5	0.25
800	0.5	0.25

The data service plans parameters used in our simulation are presented in Table 1. These plans and their distribution have been empirically determined based on the real data plan offered in Romania, by the Orange mobile operator. The Price/MB for traffic included in the data plan has been determined by factoring out the included data traffic as representing one fifth of the value of the mobile subscription plan. The 1/5 factor has been selected such that the price for traffic included in the plan will be less that the price for extra traffic for most of the service plans. Also, for each service plan there tends to be up to 5 components included in the offer (data traffic, internal voice traffic, external voice traffic, SMS, MMS).

The traffic usage pattern parameters used for the simulations are presented in Table 2. These parameters have been selected based on usual data traffic consumption patterns.

3.2 Results and Discussion

We used four metrics to evaluate OpenMobs's ability to support traffic sharing (all results below are averaged over all simulated months). The first one is *greediness*, which is the percent of monthly traffic a user uses, versus the traffic offered implicitly by his mobile data plan. In simulations, we have (a) users consuming less traffic than they are offered by the mobile provider, and (b) others consuming up to 10x times the traffic they could otherwise use (which happens because now users are motivated to use more, at lower offered prices). For the users in the first category, the gain is triggered by the traffic they sell. For the others, the gain is computed as the difference between what they would have paid without OpenMobs, and what they actually paid when using OpenMobs. As results in Figures 3a, 5a, and 4a show, because costs in OpenMobs are optimized, users selling traffic gain a profit maintained within such limits that users buying would not lose money at the end of the month, compared to what they would pay if buying traffic directly from the mobile operator.

Another metric is the *usedConections*, which shows the link between the number of times a user buys or sells traffic, versus the gain OpenMobs brings at the end of the month. As seen in Figures 3b, 5b, and 4b, users manage to gain at the end of the month proportionally to the number of times they are able to sell

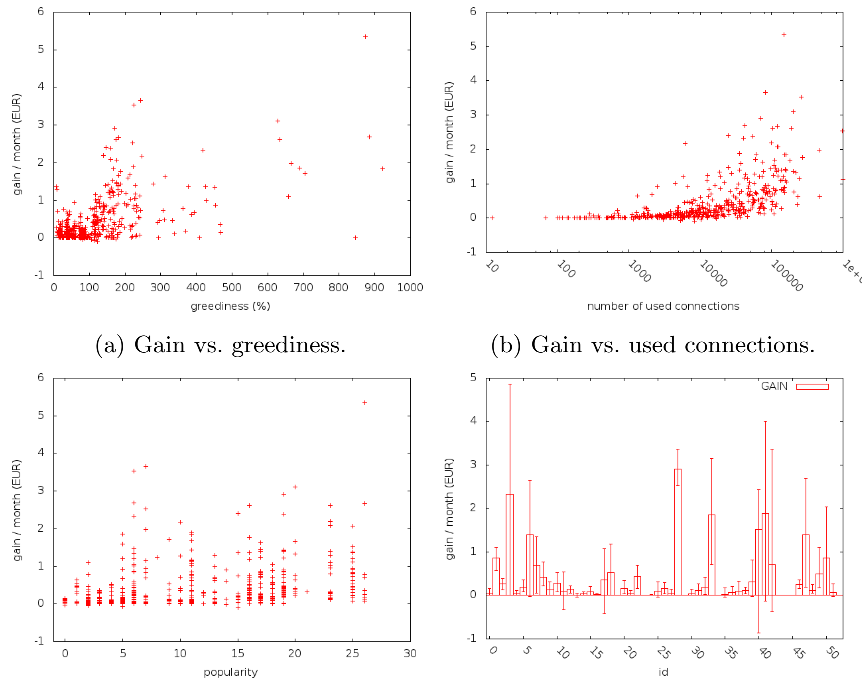

(a) Gain vs. greediness.

(b) Gain vs. used connections.

(c) Gain vs. popularity.

(d) Gain per node.

Fig. 3: Results for the experiments on UPB.

or buy traffic - thus, OpenMobs can actually incentivitize users to participate in the traffic sharing collaboration, just by the fact that the more clients use the system, the more profit they manage to gain.

The *popularity* shows the relation between the homophily of an user, and his monthly gain when using OpenMobs. Two users are considered to have a connection if they spend enough time in contact and share a number of friends in common [2]. The results in Figures 3c and 4c show that a more popular user has a higher probability of making a certain profit from using OpenMobs – in the left part, the users with relatively few friends are clustered near the bottom, while in the right part of the plots, clients with more friends are scattered and tend to gain more from using OpenMobs.

Finally, Figures 3d and 4d show the relation between the average gain for each node (user), and the actual variance of the gain during the experiments. We recall that each experiment lasts for several month, and in the beginning users can actually lose (the lower values for gain). But, because OpenMobs adapts the strategy and corrects the coefficient used in the auctioned prices for selling traffic, in the end all users start gaining. This means that, in real life, when OpenMobs is used for even more consecutive months, it can actually start bringing more profit to each user – which we believe can act as another incentive for users to use and participate in the collaboration.

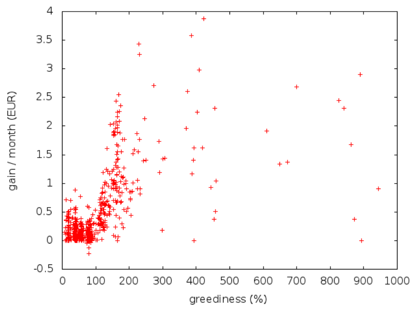

(a) Gain vs. greediness.

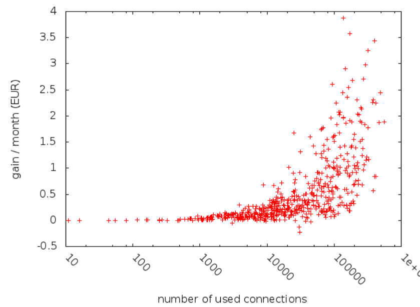

(b) Gain vs. used connections.

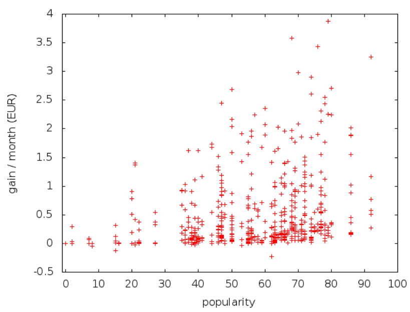

(c) Gain vs. popularity.

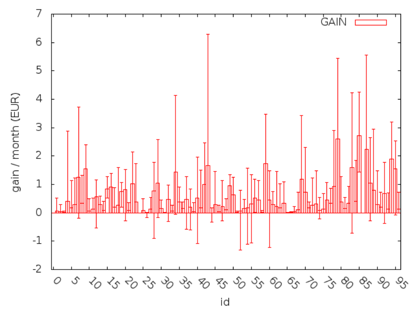

(d) Gain per node.

Fig. 4: Results for the experiments on MITReality.

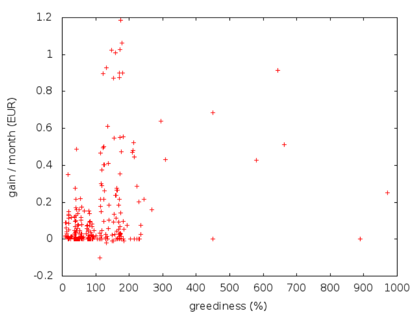

(a) Gain vs. greediness.

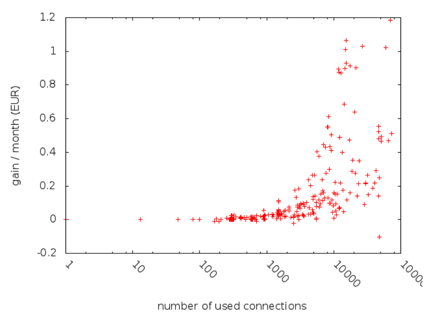

(b) Gain vs. used connections.

Fig. 5: Results for the experiments on StAndrews.

4 Conclusions

OpenMobs is a system designed to optimize the economical costs in accessing mobile broadband Internet through mobile handset devices. In this paper, we presented our approach to share under-utilized networking resources among co-located users through free wireless access. Whenever two or more users are in the vicinity of each other, OpenMobs forms an ad hoc mesh network and redirects traffic in the most economic and viable way. We presented extensive studies on the feasibility of such a system to minimize the costs users pay to their mobile providers at the end of the month, and even financially compensates users willingness to participate in the collaboration. We are currently well-underway with a real-world implementation of OpenMobs, on Android-operated devices. Also, ink the future, we aim to address also energy consumption as another parameter for our cost model.

Acknowledgment

The research is partially supported by COST Action IC1303 AAPELE, and by the national project MobiWay, Project PN-II-PT-PCCA-2013-4-0321.

References

1. Greg Bigwood, Devan Rehunathan, and all. Exploiting self-reported social networks for routing in ubiquitous computing environments. In *Networking and Communications, 2008. WIMOB'08. IEEE International Conference on Wireless and Mobile Computing,*, pages 484–489. IEEE, 2008.
2. Radu Ioan Ciobanu, Ciprian Dobre, and Valentin Cristea. Sprint: social prediction-based opportunistic routing. In *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2013 IEEE 14th Int. Symp. and Works. on a*, pages 1–7, 2013.
3. Nathan Eagle and Alex Pentland. Reality mining: sensing complex social systems. *Personal and ubiquitous computing*, 10(4):255–268, 2006.
4. Joseph Hanlon. Choosing the best broadband plan. http://www.whistleout.com.au/Broadband/Guides/choosing-the-best-bradband-plan, August 2013. [Accessed April 10th, 2014].
5. Frederick S Hillier and Gerald J Lieberman. *Introduction to operations research*. Tata McGraw-Hill Education, 2001.
6. Frederic Laridinios. Open Garden 2.0 makes sharing your wifi and mobile connections easier and faster. http://techcrunch.com/.
7. Eugene L Lawler and David E Wood. Branch-and-bound methods: A survey. *Operations research*, 14(4):699–719, 1966.
8. Francesco Malandrino, Claudio Casetti, C Chiasserini, and Marco Fiore. Content download in vehicular networks in presence of noisy mobility prediction. *IEEE Transactions on Mobile Computing*, 13(5):1007–1021, 2014.
9. Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Consulted*, 1:2012, 2008.
10. Melanie Pinola. Overview of Wi-Fi, 3G and 4G Data Plans. http://mobileoffice.about.com/. [Accessed April 10th, 2014].
11. Melanie Pinola. AT&T's new data tiers: cost less, less value? http://mobileoffice.about.com/, June 2010. [Accessed April 10th, 2014].
12. Christos Tsiaras, Samuel Liniger, and Burkhard Stiller. An automatic and on-demand mno selection mechanism. In *IEEE/IFIP Network Operations and Management Symposium (NOMS 2014), Management in a Software Defined World, Krakow, Polland*, 2014.