

Exploiting super peers for large-scale peer-to-peer Wi-Fi roaming

Efstratios G. Dimopoulos, Pantelis A. Frangoudis, and George C. Polyzos

Mobile Multimedia Laboratory, Department of Informatics
Athens University of Economics and Business

Email: stratos.dimopoulos@gmail.com, {pfrag,polyzos}@aueb.gr

Abstract—With the low installation and maintenance cost of IEEE 802.11-based equipment, dense Wi-Fi deployments are a reality, especially in today's urban areas. This vast number of WLANs can be exploited to achieve low-cost ubiquitous wireless Internet access, which is also demonstrated by the emergence of community-based wireless access schemes. In our prior work we have developed a reciprocity-based peer-to-peer architecture for Wi-Fi sharing, where peers provide free Wi-Fi access to others in order to enjoy the same benefit when they are away from their own Wi-Fi network. Our system tries to match peer consumption with contribution and we have shown it to work well for city-scale Wi-Fi sharing communities. However, when attempting to roam outside the city boundaries, the statistics are such that there is typically a lack of consumption-contribution information between consuming and providing members, which hinders the system's scalability. In this work, we extend our architecture with global-scale roaming capabilities by relaxing the requirement for full decentralization. In particular, we exploit special trusted *super-peers* which act as representatives of different Wi-Fi sharing communities (e.g., communities of different geographical regions) and which mediate transactions when there is insufficient information about peer contribution history. Extensive simulations show that this super-peer-assisted approach can significantly enhance the system's performance in terms of roaming coverage.

I. INTRODUCTION

Driven by the low cost and ease of deployment of IEEE 802.11-based equipment, as well as its operation in unlicensed spectrum, Wi-Fi has become the predominant technology for local wireless connectivity. In modern densely populated urban areas, Wi-Fi Access Point (AP) density is very high.

Considering these Wi-Fi jungles, a question that naturally emerges is whether low-cost alternatives to traditional mobile services offered by 2.5G and 3G systems can be achieved, at least for low-mobility *nomadic* users. Community-based approaches aiming to exploit the vast volume of user-provided infrastructure (e.g., residential Wi-Fi APs) are a promising answer to this question. Indeed, in recent years, there have been numerous efforts to this direction. In our prior work, we have adopted a fully-decentralized and self-organizing approach to this problem by proposing a Wi-Fi sharing scheme based on (indirect) service reciprocity.

Its fully distributed nature makes it more suitable for citywide roaming. In this paper, we focus on extending it so that it will more naturally allow for low-cost Wi-Fi roaming at a larger scale. The challenge lies in the inherent locality

properties of wireless communities; users tend to visit Wi-Fi networks in their vicinity (e.g. their city); a reciprocity-based scheme can work well in such cases. But what happens when a roaming user crosses his city boundaries?

In order to achieve the goal of larger-scale roaming, we add some degree of centralization to our design, attempting at the same time to protect its self-organizing and distributed character. In particular we introduce *super peers*, which can mediate interactions among peers when necessary. We provide incentives to users to cooperate with super peers and propose a lightweight algorithm that super peers should follow, keeping system implementation simple.

II. STATE-OF-THE-ART

A well-known Wi-Fi sharing system is FON [1]. FON provides mediation services for the development of a world-wide WLAN sharing community, where registered users can roam around FON-affiliated Wi-Fi networks. WLAN owners can either share their WLANs for a small monetary compensation or in exchange for similar service when they are away from their home WLAN. FON takes care of user registration and authentication and withholds a fraction of the money paid to the hotspot micro-operator for the provided service.

Biczók et al. [2] model the role of mediators as community providers. They analyze their interactions with users and ISPs in global-scale wireless community networks such as FON and explore the space of available parameters to determine the benefits of each player when joining the community.

Ben Salem et al. [3] study an environment where wireless ISPs have multilateral roaming contracts and register with a central authority that maintains reputation records, which are updated with QoS reports submitted by roaming users.

Two significant issues pertinent to wireless communities are studied by Manshaei et al. [4]. First, they study how initial community network coverage and user payoffs and fees affect the evolution of the community. Second, they focus on the competition between licensed wireless access providers and community-based ones, which is an important step towards answering whether wireless communities can be a viable alternative (or complement) to licensed cellular networks.

In our prior work, we have proposed that Wi-Fi be shared in a fully decentralized peer-to-peer manner [5]. A detailed description follows in Section III.

III. A PEER-TO-PEER Wi-Fi SHARING SCHEME

Our system [5], which we extend in this paper, is based on the concept of sharing wireless Internet bandwidth in a reciprocal manner; one shares his Internet connection with anonymous passers-by over Wi-Fi with the anticipation that he will enjoy the same benefit from another peer when mobile. Private WLAN owners have an incentive to contribute Internet bandwidth, given that they value much the mobile network access that they will enjoy as good contributors.

No registration with central authorities is required, nor any strong user identification scheme. Participants are identified by self-issued, uncertified public-private key pairs. To join the community, users simply configure their APs for open access and install the necessary software.

Accounting is based on digital proofs of service (*receipts*) that mobile users sign with their private keys and provide to visited APs. Receipts represent “debt” among peers, which we assume to be transitive; if peer A has provided service to peer B and the latter has served peer C, then C indirectly “owes” some service to A. A’s contribution can be reciprocated when he visits one of B’s or C’s APs. Receipts are stored in repositories distributed across the system; each peer maintains his own repository, which represents his (partial) view of the system’s history of service provisions. With each receipt encoding an “I-owe-you” relationship, a receipt repository (RR) can be viewed as a logical directed graph, whose vertices correspond to peers and its edges represent receipts and point from a service consumer to a service provider.

Such a graph is the input to the *reciprocity algorithm*, which is invoked by a prospective provider each time a user requests to be served. The internals of the reciprocity algorithm are not within the scope of this paper. Its output is a value expressing a user’s “reputation” in the eyes of a prospective service provider. This value, which we have termed *Indirect Normalized Debt (IND)*, ranges from 0 to 1, inclusive. The closer IND is to 1, the more a prospective provider “owes” to a peer requesting service. The reciprocity algorithm should be intelligent enough to detect free riders and prevent from attacks to the accounting mechanism. Our design and evaluation [5] have shown that our algorithm has these desired properties.

Receipts are generated according to a *receipt generation protocol*. Following the invocation of the reciprocity algorithm, a receipt request-response cycle begins, with the AP periodically requesting an acknowledgement for the volume of Internet traffic forwarded on behalf of the visitor thus far and the latter responding with a cumulative “fresh” receipt. The last receipt of a session contains information about the session’s aggregate amount of traffic and is forwarded to the provider’s RR. The receipt generation protocol ensures that a visitor’s session has not been *hijacked* by an unauthorized party; in order to sign a receipt and maintain the session, the hijacker would need access to the visitor’s private key.

To assist in giving potential service providers a better view of their overall contribution and have better chances of getting access, visitors also supply parts of their own RRs via a

gossiping protocol. Gossiping takes place at the beginning of a session, before the reciprocity algorithm has been executed; the visitor presents the AP with a subset of his own RR carried in his mobile device. These receipts are then merged with the visited peer’s RR and reveal debt directly or indirectly owed to the visitor.

IV. GLOBAL Wi-Fi ROAMING EXTENSIONS

A. Motivation and requirements

When users roam close to their home area, they have better chances of consuming service, because they can prove their contribution (direct or indirect) to the visited AP more easily. This is based on the assumption that users tend to roam around their city more often. On the other hand, when community scale grows and users roam far away from their home area, it is rare to find an AP that owes them directly or indirectly.

The above problem lies in the fact that cooperation decisions in this scheme are based on private information. Given this limited input, it is hard for the reciprocity algorithm to identify a visitor to a distant region as a good contributor. A solution to the above problem would be to relax one of the basic premises of our system, i.e., that of full decentralization. Facing the tradeoff between adding a central authority which would offer a global view of the system’s history of transactions or strictly adhering to the decentralization principle (with all its advantages, but with limited information), we chose to introduce *super peers (SPs)*, in order to keep the autonomy, freedom and sense of the fully distributed system. Before we present our extensions in detail, we summarize the requirements for a super-peer-based solution:

- Peers should have the right incentives to become SPs and mediate transactions; we can not rely on their altruism.
- Every peer could potentially become a SP. More computational capabilities are not a requirement.
- Peers should have the right incentives to accept roaming requests when roamers deserve service or a SP suggests they deserve it. Receipts from roamers are not as useful as other receipts, because the provider can use them only if he visits, in the near future, the area from which the consumer (signer of the receipt) came from.
- APs should contribute service to all SPs without executing the reciprocity algorithm. (Enjoying service always and everywhere is the only incentive for a SP.)

B. A mechanism for global scale roaming

Here we present a mechanism to enhance our scheme with global scale roaming capabilities. After a mobile user provides the visited AP with receipts stored in his mobile RR (gossiping), the visited peer invokes the reciprocity algorithm. If the outcome is negative, SPs need to mediate the transaction. First, the visitor informs the visited AP about the SPs of his home region. Then, both visited and home SPs run the reciprocity algorithm for the prospective consumer. Optionally, gossiping may first take place between the visitor and SPs. The visited AP has incentives to forward the mobile user’s receipts during this gossiping procedure; a successful transaction benefits

himself, too, and gossiping helps to this direction. If a region has more than one SP, all SPs run the reciprocity algorithm best of the results is kept for the prospective consumer. Two IND results thus arise; one computed from SPs of the visited peer (VSPs) and one from home SPs (HSPs).

The guarantor for the mobile user is the SP with the greater result, because this one seems to “owe” the user more. The QoS the consumer will enjoy is affected by the IND values of both the HSP and the VSP as follows:

$$IND_{TOTAL} = a \times IND_{HSP} + (1 - a) \times IND_{VSP}. \quad (1)$$

The role of a is evaluated in Section V-E.

Eventually, the AP grants access to the mobile user, if IND_{TOTAL} is over a desired threshold¹. The consumer signs receipts to the guarantor SP and the guarantor signs receipts to the AP. This way, the guarantor erases his debt to the consumer and the AP gains the valuable receipts of the SP. These receipts could be used in the near future by the peer in order to consume service when roaming in foreign areas.

Fig. 1 shows the interaction between entities during a roaming session. Note that this requires the extension of the original protocol [5] with new messages: **WARREQ** (WARranty REQuest) is used to request SPs to warrant for a roamer asking for service and contains the identity of the user. When its sender is the RR of the contributing AP, it also contains the identities of the SPs belonging to the user’s home region.

The response to a **WARREQ** message is **WARRES** (WARranty RESponse). It returns an error message if the SP cannot warrant about the user or an acknowledgement containing the corresponding IND value.

Finally, with a **WAR_NOTIFY** (WARranty NOTIFY), the visited area SP informs the visited peer and the SP with the greatest IND result (guarantor) about the identities of the mobile user and the guarantor and the calculated IND value for the mobile user.

C. Properties of our mechanism

The final cooperation decision is affected by two different SPs (HSP and VSP), because we want peers to offer service to SPs no matter if they owe them or not. If we considered only the response of, e.g., the HSP, users would have incentives to offer service only to them, because only HSPs could help them in future roaming transactions. Contribution to foreign SPs would not be of any help. In Section V-E we show that the relative weight of these two values in the decision (the a factor) is significant.

On the other hand, IND_{HSP} is typically higher than IND_{VSP} . For this reason we want to keep VSP contribution to the final result as low as possible, but high enough to keep the aforementioned incentive. A small value of a thus keeps the incentives strong but degrades service quality, since IND translates to QoS.

A receipt signed by the roamer, as the usual receipt generation procedure implies, wouldn’t satisfy the visited AP. The

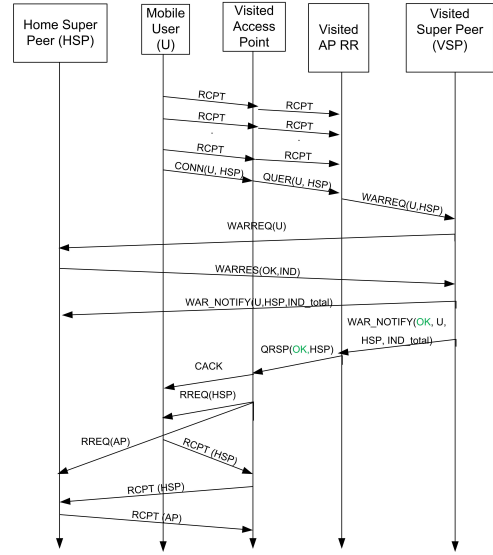


Fig. 1. Message exchange during global scale WLAN roaming.

benefit from this receipt exists only if the providing peer visits the home area of the roamer in the near future and more specifically an AP of the particular consumer. Because this situation is rare, we extended the receipt generation protocol to satisfy contributors of roaming transactions. As already mentioned, the receipt signer to the contributor is the super peer who guarantees for the consumer. This solution satisfies contributors because they can use super peer’s receipts on their future roaming or local transactions to increase their chances of getting high quality service.

V. EVALUATION

Here we present a simulation-based evaluation of our super-peer-assisted scheme by comparing it to a fully-decentralized Wi-Fi sharing approach in view of global-scale user mobility.

A. Evaluation metrics

Our metrics of interest are the *Hit Ratio (HR)* and the *Social Welfare (SW)*. The HR is the percentage of successful roaming access attempts, namely service requests by roaming users which have been accepted by the visited APs.

As to the Social Welfare, it represents the sum of the *scores* of all community members. The score of a peer is the benefit the user enjoys from the service received minus the cost associated with the times the peer has offered service to visitors. In our evaluation, the cost for a contributing peer per transaction is assumed constant, while the benefit for a successful service reception is a linear function of the IND value and ranges from 0 to b_{max} . In our simulations, the cost for providing service is 1 and $b_{max} = 10$. Also, note that in the following figures, instead of the (community-wide) SW, we plot the average SW per match (transaction).

B. Simulation environment

We have developed our own custom simulator in Java. Our simulated world is divided in *regions*. All regions have the

¹Recall that IND values are translated to QoS. For instance, a visited AP can prioritize the traffic it forwards on behalf of visitors with high IND values.

same number of *local* peers, who may either visit APs of peers within the same region or roam to other regions.

After a number of simulation rounds spent on the bootstrapping phase, when peers contribute unconditionally in order to build their reputation in the community, peers enter a roaming phase, when they “travel” to another region (with a fixed roaming probability) to request service. The visited region is picked uniformly at random. They stay there for a number of rounds (stop-over rounds) in which they still try to consume service and then return back to their home region. In all transactions they use their home (server) and mobile (client) repositories to store receipts and run the reciprocity algorithm. In the super-peer-based architecture, we put *at least* one super peer per region.

It should be noted that we only consider reciprocal users, contrary to our prior work [5], where we evaluated our algorithms in the presence of attackers. The reason is that here we wish to quantify the benefits of exploiting super-peers in view of global mobility and without the influence of different strategies selection. In most figures we compare the performance of a super-peer-assisted scheme *vs* a fully decentralized system (the one proposed in our prior work).

C. The effect of the number and size of regions

Here we study how the two systems behave as the number of regions grows, while keeping the number of total users constant. In all cases, the system based on SPs behaves almost perfect for both user categories, but the performance difference gets greater when we divide our system in many low populated regions. In this case, which is the most realistic one, users of foreign regions are very unlikely to trust each other and, as a result, the need for SPs is greater. Not only do roaming users have greater probability to consume service (Fig. 3) but also the QoS they will enjoy (expressed by the average SW) is much greater (Fig. 2).




Fig. 2. Number of regions effect on SW.

D. Server repository size effect

The interesting fact about this experiment (Fig. 4 and 5) is not that the system with SPs performs better but that, when the size of the server repositories grows, meaning that all peers have a better view of the system, then the SW per roaming transaction decreases. The explanation to this is given




Fig. 3. Number of regions effect on Hit Ratio.

by Fig. 6; when the server repository size is large, roamers have higher probability to consume service without the help of their SPs, based only on their reputation. Their reputation is enough in this case to let them consume but it is still not enough to let them consume service of good quality. In other words, some times it is better to have no reputation at all, than having little but enough reputation to consume service without the mediation of SPs.




Fig. 4. Server repository size effect on Hit Ratio




Fig. 5. Server repository size effect on SW.

E. Participation of SPs in the IND result

This simulation result (Fig. 7) proves what we have discussed in Section IV-C about home and visited SP participation in the calculation of IND for the prospective consumer. It is always better for roamers to use in the greatest possible




Fig. 6. Server repository size effect on super peers usage.

percentage the IND calculation of their home super peer, but we must be careful to keep this percentage as low as needed to keep strong user contribution incentives.




Fig. 7. Home super peer participation effect on SW

F. The effect of the number of SPs per region

Here we answer the question of whether more SPs per region can bring better results. As shown in Fig 8, the answer is positive. In densely populated areas, where only one SP may not be adequate to support all of its region's roamers, it seems that the addition of one more super peer can improve performance.




Fig. 8. Super peers per region effect on SW.

G. Scale effect

Here we study how the community performs when its size grows, while keeping all other parameters fixed. The results

are very promising (Fig. 9), since HR decreases a little but still remains in a good level. Furthermore, roaming users still receive the same a high QOS level(Fig. 10).




Fig. 9. Scale Effect on Hit Ratio




Fig. 10. Scale Effect on SW.

VI. CONCLUSION

We deal with a fundamental problem a decentralized Wi-Fi sharing scheme would have to face, i.e., the locality of movement that characterizes users; they typically interact with their local peers, which makes it harder to demonstrate their contribution and community reputation in the case of occasional visits far from their region. Since the root of this problem is the lack of a global view of peer contributions, we have extended the fully-decentralized Wi-Fi sharing scheme that we have developed in our prior work with the use of *super peers* mediating service requests when necessary. Our evaluation has shown that with this approach we can overcome this problem to a significant extent and move towards achieving global-scale Wi-Fi roaming.

REFERENCES

- [1] "Fon," <http://www.fon.com/en/>.
- [2] G. Biczók, L. Toka, A. Vidacs, and T. A. Trinh, "On incentives in global wireless communities," in *Proc. of the 1st ACM workshop on User-provided networking*, 2009, pp. 1–6.
- [3] N. Salem, J. Hubaux, and M. Jakobsson, "Reputation-based Wi-Fi deployment," *ACM Mobile Computing and Comm. Rev.*, vol. 9, no. 3, pp. 69–81, 2005.
- [4] M. H. Manshaei, J. Freudiger, M. Félegyházi, P. Marbach, and J.-P. Hubaux, "On wireless social community networks," in *Proc. IEEE INFOCOM*, 2008, pp. 1552–1560.
- [5] E. C. Efstathiou, P. A. Frangoudis, and G. C. Polyzos, "Controlled Wi-Fi sharing in cities: A decentralized approach relying on indirect reciprocity," *IEEE Trans. Mobile Computing*, vol. 9, no. 8, pp. 1147–1160, August 2010.