

Incentives and Algorithms for Broadband Access Sharing

Lefteris Mamatras, Ioannis Psaras, George Pavlou
Dept. of Electronic & Electrical Engineering
University College London
WC1E 7JE, Torrington Place, London, UK
{l.mamatras, i.pсарas, g.pavlou}@ee.ucl.ac.uk

ABSTRACT

The unprecedented growth of the mobile (Smart)phone industry that comes together with the corresponding application development market has made apparent that mobile networking through 3G links is just about to reach an unbreakable limit, in terms of network capacity. The networking research community has recently started considering alternative connectivity approaches to support and boost the performance of mobile networking. In particular, researchers have identified a big amount of "power", hidden at the edges of the network, which remains there unexploited and is no other than the WiFi technology deployed in home-networks. We explore incentives and algorithms for Broadband Access Sharing to support nomadic users and show that ubiquitous connectivity in densely populated areas is already possible, since the infrastructure is already there, waiting to be used.

Categories and Subject Descriptors

C.2.0.c [Communication/Networking and Information Technology]: Emerging technologies; C.2.2.a [Network Protocols]: Applications

General Terms

Algorithms; Design

Keywords

Broadband Access Sharing, User-Provided Networks (UPNs), UPN Incentives, Load-Balancing

1. INTRODUCTION

"In a commercial network, new technologies essentially get deployed for reasons of fear or greed." [6] "Fear" is the danger of a component of the network collapsing, or not being able to keep up with the demand placed on it and therefore, being in a commercial environment, a player will lose money.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

HomeNets 2010, September 3, 2010, New Delhi, India.

Copyright 2010 ACM 978-1-4503-0198-5/10/09 ...\$10.00.

"Greed" is the natural motive of a player getting involved in the deployment of the network and means that someone can make (more) money out of its deployment. Given the above, what is the status of the Mobile 3G Internet technology? Is it about to face a big problem soon? Is it able to keep up with current trends in the mobile phone industry?

On the one hand, 3G connectivity was deployed for reasons of greed, but admittedly gives the opportunity of some sort of connectivity to mobile users. On the other hand though, the users' experience is far from satisfactory, since in many (if not most of the) cases, 3G connectivity barely works: the explosion of Smartphone Applications makes the 10-year ago dial-up connection seem high-speed. Indeed, most applications require WiFi connectivity, while the rest are struggling to download even simple text. The worrying point about this situation is that the mobile phone application market has just only now taken off and nobody knows what the (very near) future is going to bring.

With these concerns in mind, researchers have recently investigated the broad field of *Broadband Access Sharing*, mainly in the context of *Citywide Ubiquitous WiFi Access* [13]. The main motive here is to find a balance between 3G and WiFi connectivity in densely populated areas. That is, to support the growing need for mobile connectivity, the operators have to massively overprovision their 3G networks, or otherwise offload some of the traffic to WiFi networks, saving this way money from 3G network extensions and implicitly increasing their revenue. Of course, all this is based on the fact that almost every house has a broadband Internet connection that terminates in a wireless Access Point (AP)¹.

In our opinion, realization of the *connectivity sharing* concept depends mainly on the following two factors: *i) the wireless interface's capability to share resources and ii) the resource sharing scheme between home- and roaming-users.*

Point i) above has been mainly studied with respect to the current wireless technologies, namely IEEE 802.11 and its performance when one of the wirelessly connected users is far away from an AP (e.g., [11], [2], [12]).

Authors in [14] have built a planning tool, called *Met-roSim*, to choose in the most cost-effective and performance-efficient way the home-APs that the operators should "leave open" in order to provide almost full coverage to roaming-users. This brings us to our second point above, namely, *the terms and conditions under which connectivity is shared.*

¹E.g., statistics indicate that 70% of the houses in the UK have broadband Internet connection: <http://statistics.gov.uk/cci/nugget.asp?ID=8>

We argue that an essential part of the broadband access sharing paradigm is the actual framework, in terms of rules and regulations, under which resources are shared. In particular, we contend that the *sharing scheme* or *sharing framework* is mainly formed of: i) the *Incentives* given to the home-user in order to share his connection with strangers, ii) the actual rules and regulations (i.e., *Algorithms*) for bandwidth sharing and iii) the *Security* guarantees given to both the home- and the guest-user for safe transactions.

Co-existence of home- and guest-user traffic in the same AP clearly raises security issues for both of them. Authors in [9] propose that tunneling the guest-user traffic through a "trusted" point elsewhere in the Internet (e.g., the guest's actual home-AP) can alleviate most of the security issues that can be possibly raised.

To the best of our knowledge there is no publicly available study to investigate the incentives that need to be given to home-users in order to share their connection with strangers, or to issues related to bandwidth sharing and load-balancing between home- and guest-traffic. Our view of the general *Broadband Access Sharing* research framework together with the already investigated topics and the missing bits is illustrated in Figure 1. Here, we elaborate on both the *incentives* and the corresponding *algorithms* that would guarantee smooth co-existence of home- and guest-users. We study the above issues in the context of *User-Provided Networks* [10]. We summarize the contributions of the present study in the next Section.

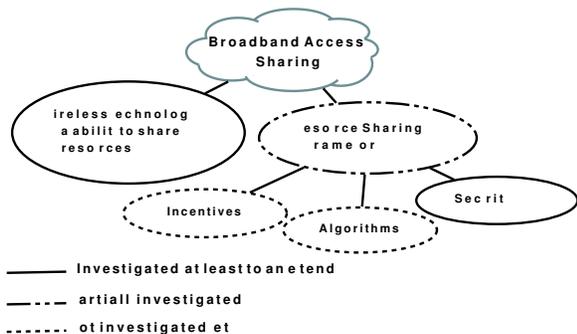


Figure 1: Broadband Access Sharing Research

2. SUMMARY OF CONTRIBUTIONS

We show that sharing his connection, the end-user can gain unlimited connectivity, when mobile, without any impact on the performance of his own broadband Internet connection, when at home. Throughout the paper, we use the terms *home-user* and *micro-provider* interchangeably; the home-user's AP is herein denoted as UPN-AP.

We realize the above by an *active queue management* service differentiation mechanism, which prioritizes home-against guest-traffic and a *load-balancing* technique that routes guest-traffic through the least-utilized access point within range in order to boost the performance of nomadic users.

Therefore, the incentive for the home-user to share his broadband Internet connection is simple: he shares a portion of his bandwidth *seamlessly*, when at home and gains unlimited connectivity, when mobile. We call this scheme *Offer Nothing - Gain Something*. Of course, this "Something" depends on the density of the available access points as well as

on the bandwidth utilization of the respective home-users. That is, connectivity may be poor and intermittent, which makes DTNs, or connectivity through 3G links an essential complementary part for UPNs. We argue that *connectivity is for mobile computing what bandwidth is for the wired core of the Internet*. Therefore, instead of high-speed core Internet links, *the backbone for a mobile environment is the area of dense connectivity*.

We note that service differentiation through active queue management has received a lot of attention in the past few years (e.g., [3]), mainly in order to prioritize high-paying customers' traffic against low-paying ones', or to boost the performance of non-elastic, or non-congestive applications against bursty TCP transfers. Load-balancing, on the other hand, has been studied mainly in the context of server farms (e.g. [1]), peer-to-peer networks (e.g., [5]) and multipath routing mechanisms (e.g., [4]). The novelty introduced herein is the actual marriage and application of these techniques in the context of UPNs in order to realize new communication technologies and connectivity paradigms that explicitly support mobile computing. To the best of our knowledge, this is the first study that puts these two research fields under a different context, that of connectivity sharing and elaborates on their potential to set the foundations for the realization of "Broadband Access Sharing through User-Provided Networks".

This work has been undertaken in the context of content-based networks. In particular, we are currently investigating how mobile users can identify content and download it to their mobile devices in the fastest possible manner. This has to be studied mainly with regard to two issues: i) how connectivity sharing affects the respective home-user and ii) how content is identified and routed back to the mobile-user. Here, we investigate point i), while point ii) constitutes our immediate next step. On that direction, we plan to define content-identifiers, investigate content-caching techniques and routing approaches that guarantee smooth content fetching to the mobile device.

3. PROTOCOL DESIGN FOR UPNS

We realize the above-mentioned *Offer Nothing - Gain Something* scheme through a *Service Differentiation* technique, which is further complemented by a *Load-Balancing* mechanism. We briefly elaborate on the proposed schemes below. The interested reader is referred to [8] for a complete version of the proposed algorithms and detailed discussions on their functionalities.

3.1 Service Differentiation for UPNs

If the user connects to his own access point, he is identified as the home-user or the *micro-provider*, through a pair of credentials, while if he is mobile and connects to an unknown access point, he is identified as the guest-user. In turn, whatever packets are sent or received by that user are classified accordingly. This way, we implement a packet-classification algorithm to (i) identify *home-* and *guest-packets* and (ii) apply the corresponding service differentiation. We call this algorithm *User-Provided Network Queuing* (UPNQ).

Service differentiation is applied on a per-packet basis and depends mainly on the queue utilization at the time a packet arrives at the *micro-provider's* AP. In particular, we implemented a packet scheduling mechanism based on the non-preemptive priority queuing scheme. In our case, the home-

packets receive higher priority at all times, which gives a constant performance advantage to the home Internet connection owner. Since the non-preemptive nature of the mechanism may indeed cause a slight impact on the performance of the home-user, we complement the mechanism with one extra capability. Whenever the percentage of the guest-packets in the queue exceeds a certain threshold, which we call *upnthresh*, the newly-arriving guest-packets are forcefully dropped. In our results, a value between 15-20% produces the expected outcome; as shown in [8] through extensive simulations, this is the worse-case scenario for UPNQ. We have shown in [7] that a small rate of packets can be serviced with zero impact on the performance of the flows sharing the same channel.

The above algorithm guarantees that the guest-traffic is not noticeable from the home-user, while maintaining, in the worst case, connectivity to a mobile device that may access services that are not demanding (e.g., Web or e-mail). We note that in case the home-network is underutilized, the offered service may reach the level of a paying customer.

In [8], through queuing theoretical analysis we have shown that if:

- the arrival rate of the home- and the guest-traffic at the UPN-AP is λ_1 and λ_2 , respectively,
- the average service time for the home- and the guest-traffic at the UPN-AP is T_{s1} and T_{s2} , respectively, and
- k is the impact of the guest's presence to the home-user's performance in terms of queuing delay,

then the guest's arrival rate can be regulated according to Equation 1:

$$\lambda_2 \leq \frac{k(2 - \lambda_1 T_{s1})}{T_{s2}} \quad (1)$$

In [7], we have shown that when the impact of the guest's queuing delay to the home-user (i.e., k) is below 5%, this is, in general, unnoticeable by the home-user's applications. That said, setting $k \approx 0.01 - 0.05$ (i.e., 1% - 5%) and regulating the rate of the guest's traffic (λ_2) according to Equation 1 will have statistically no impact on the performance of the home-traffic (λ_1). We refer the interested reader to [8] for the full theoretical analysis as well as for performance evaluation results that prove its validity.

Our proposal inherently assumes that the bottleneck is at the micro-provider's AP, which is a reasonable argument for today's networks. That is, the DSL connection to the ISP is usually slower than the wired or wireless LAN that spreads further away from the AP towards the edge of the home-network. Therefore, in all cases the guest will not cause network overload further away from what the micro-provider's contract allows. In case the bottleneck is at the ISP link, then adjustment of the incoming/outgoing AP's bandwidth can move the congestion bottleneck back to the UPN-AP.

3.2 Load-Balancing for UPNs

To guarantee the best possible service for the guest-user, the service differentiation algorithm presented above is complemented with a simple Load-Balancing mechanism, which we call *User-Provided Network Load-Balancing* (UPNLB). It is possible that a mobile-user will be within range of more

than one *micro-providers*, at least in densely populated, urban areas. Clearly, the guest-user can make the most out of his UPN experience, if he connects to the AP with the lowest home-user traffic/usage. For that reason, each AP broadcasts the home-user's queue-usage in regular time intervals (every two seconds for our experimental study included herein) to the guests within range².

We consider the topology illustrated in Figure 2: three home-users (H_1 , H_2 and H_3) provide connectivity to three roaming, guest-users (G_1 , G_2 and G_3). We assume that the least-utilized UPN-AP is that of home-user H_2 . We investigate whether and in which cases it is efficient for both guests G_1 and G_3 to change their point of attachment to H_2 . We identify three main approaches:

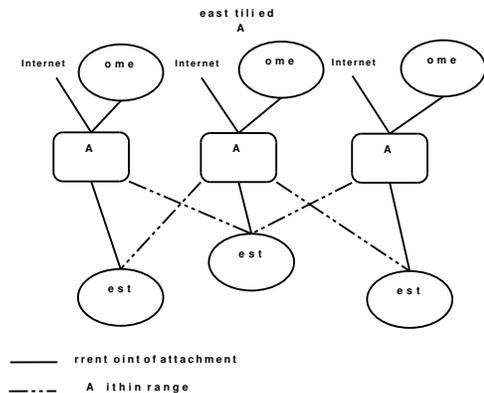


Figure 2: Load-Balancing Topology

1. **Guests move to the least-utilized UPN-AP within range with probability $p_m = 1$.**
2. **Guests move to the least-utilized UPN-AP within range with probability $p_m < 1$.**
3. **Guests move to the least-utilized UPN-AP with probability equal to the queue-usage at their current point of attachment.**

Modeling the act of *switching* or *not switching* as a *Binomial Random Variable*, X , with parameters (n, p) , where n is the number of users that consider switching (i.e., G_1 and G_3 , therefore two in our case) and p is the probability of *switching* or *not switching* (hence, $p = 0.5$ in our case), we find the following. According to the first approach above, all guest-users will switch to the same UPN-AP, i.e., H_2 . In that case, H_2 may become congested because of the large number of guests (i.e., the approach would still scale for a small number of guest-users). For the second approach, we find that the probability that both G_1 and G_3 decide to switch to H_2 is given by $p_m(i) = \binom{n}{i} p^i (1-p)^{n-i} = \frac{n!}{(n-i)! i!} p^i (1-p)^{n-i}$, where $i = 2$. Based on the above, we get that $P\{X = 2\} = 0.5^2 \cdot 0.5^0 = 0.25$. Although this approach has better scalability properties than the previous one, it is still "blind" with regard to the AP's queue-utilization.

Proceeding to the third approach, we reason as follows: given that none of the guest-users knows how many others

²The broadcast notification messages consist of a few bytes and cross one hop only, hence, the associated communication overhead is insignificant.

are attached to the UPN-APs within common range, a guest-user should only change his point of attachment only if there is a real need to do so. That is, he should switch to a less-congested AP, only if his current point of attachment is really busy. For example, if the queue utilization at H_1 is 75%, at H_2 is 25% and at H_3 is 50%, then the probability that all three guest-users move to H_2 is 0.375, according to the third approach. In particular, chances are that G_1 will move to H_2 , while G_3 will remain attached to H_3 . We consider this a successful outcome indeed, since it manages to load-balance guests among the available UPN-APs. In [8], we include a comprehensive discussion on the load-balancing mechanism, as well as results for its switching properties, based on the simple probabilistic analysis introduced earlier on in this Section.

4. EXPERIMENTAL ANALYSIS

4.1 Alternative Sharing Schemes

Broadband access sharing schemes have already been commercially deployed and enjoy wide popularity among millions of Internet users around the world. We list the main commercial players in this market below, together with their main operational properties.

- **FON Community** (<http://www.btfon.com>): FON is the biggest connectivity sharing community around the world. According to the FON scheme, home-users allocate explicitly and at all times a portion of their bandwidth to mobile users. Here, we focus on the FON package provided by BT: the maximum speed that the guest-user can get is 512kbps, regardless of the home-user's network usage. In Section 4, we compare our proposed algorithms with the BT FON scheme and discuss its advantages and disadvantages.
- **OpenSpark Community** (<http://open.sparknet.fi>): According to OpenSpark, there are no limits on the resources that the guest-user can exploit. However, according to this scheme, the guest-users have to be added manually. Clearly, this scheme is not so flexible in giving access to mobile, roaming users.
- **Wifi.com** (<http://www.wifi.com/>): This scheme follows a social-networking approach: each member of the community grants access to friends and colleagues at will and free of charge, realizing a so-called *Shared Hotspot* framework. Again, there are no bandwidth limitations for the guest-users.

The success and growth of the above initiatives clearly shows that users are keen on giving away a portion of their bandwidth in order to have access when out of home or office. In our opinion, whatever the *Broadband Access Sharing* scheme, it has to balance between two fragile service-points: (i) the guarantee to the home-user for seamless sharing, performance-wise and (ii) the mobile-user's quality of experience. Requirement (i) is met only by the FON scheme, while requirement (ii) is only partially met by the OpenSpark and the Wifi.com schemes. That is, we argue that the mobile-user should be able to attach to different APs *on-the-fly* in order to benefit from mobile Internet services. In that sense, our proposal comprises an extension/enhancement to the above-mentioned schemes. We base our design on the

fact that underutilization of resources in home networks is the rule and not the exception. That given, we highlight that *the proposed scheme does not upper-bound the portion of bandwidth that the guest-user can exploit, provided that sharing is seamless to the home-user at all times.*

4.2 Simulation Setup

We have gathered detailed information regarding the exact position and the number of BTFON users in the city center of London, which we parsed into the *ns-2* simulator. Based on this information, we simulate two different topologies (Figures 3 and 4), where mobile-users walk and occasionally attach to home-APs. In *Topology 1*, we simulate a 553-meter distance, from the Electrical Engineering Department of UCL to the British Museum (Figure 3), while in *Topology 2*, we simulate a roaming user walking in a random, central street of the city (Figure 4). The guest-users attach to home-networks, whenever within range and transfer a small file using FTP (i.e., a 10MB file, unless stated differently). The guest-users' default speed is 1m/sec (i.e., 3.6km/hour), unless explicitly stated otherwise. This setting could represent a guest-user attempting to attach a photo to an e-mail, or upload it to a social-networking site, or download a map or touristic info file to his device. In all cases, the guest-users connect through 11Mbps wireless links to home-APs; in turn home-APs, connect to the Internet through regular ADSL 8Mbps links.

In Figures 3 and 4, we show all the available BTFON APs. Our walking users do not connect to all APs shown, but only to those that are within range according to the data gathered from the BT FON web-site.

We compare the performance of the proposed UPN Protocols with DropTail (i.e., completely open APs) and with the BTFON scheme. To represent a realistic situation, we simulate home-users that transfer a 100MB file at random times and then suspend. We measure the time that guest-users need in order to complete their task; we call this time the *Task Completion Time* (TCT). Our target is to show that using the *UPN Protocols* (i.e., UPNQ and UPNLB), the guest-user gets acceptable service, connectivity-wise, while sharing is seamless to the home-user.

Recall that the BTFON home-user allocates 512kbps to the guest-user, which in turn means that the guest-user's bandwidth is bounded by that value. Furthermore, according to the OpenSpark and the Wifi.com connectivity sharing approaches, the guest-user is not bounded as for the bandwidth he can use. Therefore, in the following, the performance of DropTail can be assumed to represent these connectivity-sharing schemes. The *upnthresh* value for UPNQ is set to 20% for simulations included herein; as shown in [8] through extensive simulations, this is the worse-case scenario for UPNQ.

We note that the BTFON, OpenSpark and Wifi.com connectivity schemes were not originally designed for mobile environments. Here, we evaluate their performance comparatively to our design proposals in order to illustrate their potential in mobile environments.

4.3 Topology 1: UCL EE to British Museum

As a first step, we attempt to assess the performance of the home- and the guest-users for an increasing number of guest-users. The goal is to evaluate the scalability properties of the candidate connectivity sharing approaches under

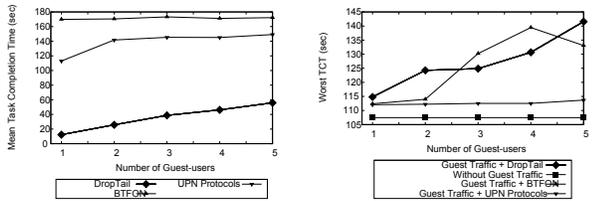


Figure 3: Topology 1: UCL EE to British Museum



Figure 4: Topology 2: Pentonville Road

realistic mobility conditions. In Figure 5(a), we observe (as expected) that the least aggressive algorithm with regard to the guest-user's traffic (i.e., DropTail) outperforms BTFON and the UPN Protocols. BTFON on the other hand, as the most aggressive approach against guest-traffic increases TCT for the guest-user. In Figure 5(b), we present the home-users' worst-case TCT; we see that both DropTail and BTFON impact the performance of the home-user, while UPN Protocols perform nearly optimal (i.e., the impact on the home-user, delay-wise, is less than 5 seconds, in a total of more than 100 seconds, in the worst-case, which we consider insignificant). The Load-Balancing mechanism of UPN Protocols will choose and transmit the guests' traffic through the least-utilized UPN-AP within range. This feature together with the UPN Queuing algorithm, which occasionally becomes aggressive against the guest-users if conditions indicate so, guarantees smooth performance for the home-user, regardless of the number of guests (Figure 5(b)).

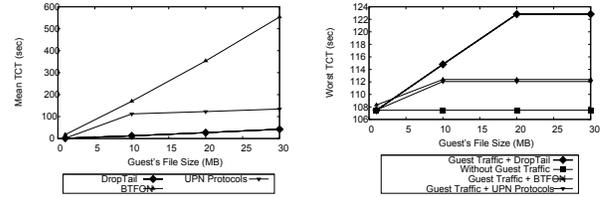


(a) Guest-User Mean TCT (b) Home-User Worst TCT

Figure 5: Topology 1 - Increasing Guests

We go one step further to investigate the scalability properties of the algorithms with regard to the guest-user's file size. In particular, we simulate one walking guest-user and we vary the amount of data to be transferred. The results are shown in Figure 6.

The performance problems of the inflexible and non-adjustable flat-rate of the BTFON scheme (i.e., 512Kbps) are made clear in this scenario. We see in Figure 6(a) that the BTFON guest-user may need up to 550 seconds to transfer a 30MByte file, while the proposed UPN Protocols complete the task in less than 150 seconds. As expected, the DropTail approach is even faster (i.e., DropTail needs 50 seconds at most to complete the task), but this comes at the cost of the home-user's performance. In Figure 6(b), for example, we see that this performance degradation for the home-user may reach up to 14%.



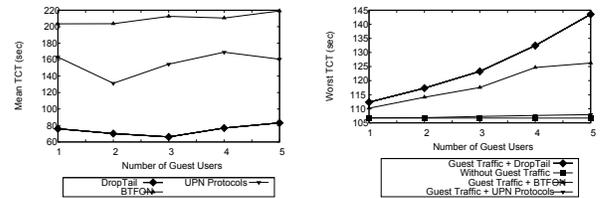
(a) Guest-User Mean TCT (b) Home-User Worst TCT

Figure 6: Topology 1 - Increasing File Size

4.4 Topology 2: Pentonville Road

In order to further evaluate the validity of our claims, we randomly pick a road in central London and repeat the previous simulations. The distance from the starting to the ending point is now 1,217 meters to allow for more diverse connectivity conditions; the topology is shown in Figure 4.

We observe that the performance difference between UPN Protocols and the alternatives now widens even more (Figures 7, 8 and 9). For example, in case of a group of mobile-users the average Task Completion Time is approximately 25% lower for UPN Protocols than for the BTFON user, as we see in Figure 7(a). At the same time, the impact to the home-user is negligible for UPN Protocols, even when considering the worst-case performances (i.e., Figure 7(b)). In contrast, in Figure 7(b), we observe that the home-users that share their connection with strangers under the BTFON or the simple DropTail schemes may occasionally experience severe performance degradation, which may be up to 30-40%, respectively (Figure 7(b)).



(a) Guest-User Mean TCT (b) Home-User Worst TCT

Figure 7: Topology 2: Increasing Guests

Increasing the size of the file under transmission, we observe in Figure 8(b) that although the BTFON scheme does not impact the performance of the home-user, it fails to provide acceptable services to the mobile, roaming guest. In Figure 8(a), for example, we see that the BTFON user

may need up to 10 minutes to download/upload a 30MByte file; UPN Protocols, on the other hand, exploit available resources in the most efficient manner and complete the transmission in less than three minutes. Our packet-level traces, which we do not present here due to space limitations, show that the UPNLB algorithm always switches the guest to empty UPN-APs, if any, or to the least-utilized ones, otherwise. In all cases, UPN Protocols preserve the quality of experience for the home-user (Figure 8(b)).

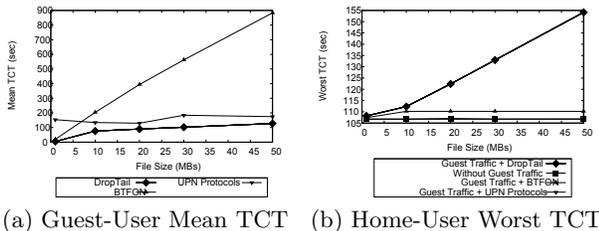


Figure 8: Topology 2: Increasing File Size

The performance gain of our proposed Load-Balancing mechanism is made clear in our final simulation scenario, where we vary the speed of the roaming users (Figure 9). The speed-range simulated here varies from 1m/s (i.e., 3.6km/h) for a slow-walking pedestrian to 20m/s (i.e., 72km/h) for a moving vehicle. We observe that in some cases (i.e., 10-20m/s, Figure 9) UPN Protocols complete their task even faster than DropTail, which explicitly owes to the UPNLB algorithm. That is, DropTail guests get hooked to the first AP within range and stay attached to that specific one till they complete their task or lose connectivity. In contrast, UPN Protocols constantly probe for less-utilized UPN-APs and switch once a better offer appears (Figure 9).

Although the results of this particular scenario cannot be generalized, since clearly connectivity times/points depend on the specific setting and the availability of APs, we still claim that the UPN Protocols' increased performance is due to the efficient handling of the UPNLB load-balancing mechanism introduced earlier on.

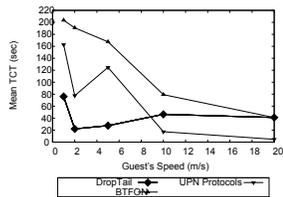


Figure 9: Topology 2: Increasing Speed

5. CONCLUSIONS & OPEN ISSUES

Our study targets the broad topic of *Broadband Access Sharing*. We have investigated related works and found that the picture is still vague as for the scheme under which resources are shared between the home- and the guest-users. Due to the limited amount of studies in this area, this work places more questions than the ones it answers. Some of our future plans include optimization of the proposed mechanisms and investigation of scalability issues with regard to

(i) the number of guest- and home-users and (ii) the corresponding applications and their requirements throughput- and delay-wise. In particular, we consider that the real gain of the load-balancing mechanism, has to be evaluated together with the nature of the requested content (e.g., real-time video vs static web-page). Furthermore, UPNLB can potentially get input from the moving patterns of the mobile users and manage connections accordingly.

6. ACKNOWLEDGMENTS

This work was partly funded by the EU FP7 COMET project - Content Mediator Architecture for Content-Aware Networks, grant agreement 248784.

7. REFERENCES

- [1] J. Brendel et al. World-wide-web server with delayed resource-binding for resource-based load balancing on a distributed resource multi-node network, June 1998. US Patent 5,774,660.
- [2] V. Bychkovsky et al. A Measurement Study of Vehicular Internet Access Using in Situ Wi-Fi Networks. In *Proceedings of MobiCom '06*, pages 50–61, New York, NY, USA, 2006. ACM.
- [3] M. Claypool, R. Kinicki, and A. Kumar. Traffic sensitive active queue management. In *Proceedings of the 8th IEEE Global Internet Symposium*, Miami, Florida, March 2005.
- [4] Y. Ganjali and A. Keshavarzian. Load balancing in ad hoc networks: single-path routing vs. multi-path routing. In *Proceedings of IEEE INFOCOM*, 2004.
- [5] B. Godfrey et al. Load balancing in dynamic structured P2P systems. In *IEEE INFOCOM*, volume 4, pages 2253–2262, 2004.
- [6] M. Handley. Why the Internet Only Just Works. *BT Technology Journal*, 24(3):119–129, 2006.
- [7] L. Mamas and V. Tsaoussidis. Differentiating Services with Noncongestive Queuing (NCQ). *IEEE Transactions on Computers*, 58(5):591–604, May 2009.
- [8] I. Psaras and L. Mamas. On Demand Connectivity Sharing: Queuing Management and Load-Balancing for User-Provided Networks. Technical Report, Available at: <http://www.ee.surrey.ac.uk/Personal/I.Psaras/ondemand-upn-ipsaras.pdf>.
- [9] N. Sastry, J. Crowcroft, and K. Sollins. Architecting Citywide Ubiquitous Wi-Fi Access. In *Proceedings of HotNets-VI*, 2007.
- [10] R. Sofia and P. Mendes. User-provided networks: consumer as provider. *IEEE Communication Magazine*, 46:86–91, December 2008.
- [11] M. Solarzski et al. An Experimental Evaluation of Urban Networking using IEEE 802.11 Technology. In *Wireless Mesh, 2006*, pages 1–10.
- [12] A. Subramanian et al. Drive-By Localization of Roadside WiFi Networks. In *Proceedings of Infocom '09*, pages 718–725, 2008.
- [13] C. Thraves et al. Driving the Deployment of Citywide Ubiquitous WiFi Access. In *SimulWorks*, 2008.
- [14] G. A. Urueta et al. Metrosim: a planning tool for metropolitan wifi networks. In *Proceedings of Simutools '09*, pages 1–8, 2009.