Internet Access for All: Assessing a Crowdsourced Web Proxy Service in a Community Network

Emmanouil Dimogerontakis, Roc Meseguer, Leandro Navarro

Universitat Politècnica de Catalunya
{edimoger,meseguer,leandro}@ac.upc.edu

Abstract. Global access to the Internet for all requires a dramatic reduction in Internet access costs particularly in developing areas. This access is often achieved through several proxy gateways shared across local or regional access networks. These proxies allow individuals or organisations to share the capacity of their Internet connection with other users. We present a measurement study of a crowdsourced Internet proxy service in the guifi.net community network that provides free Web access to a large community with many small proxy servers spread over the network. The dataset consists of Squid proxy logs for one month, combined with network topology and traffic data. Our study focuses on a representative subset of the whole network with about 900 nodes and roughly 470 users of the web proxy service. We analyse the service from three viewpoints: Web content traffic from users, performance of proxies and influence of the access network. We find clear daily patters of usage, excess capacity and little reuse of content which makes caching almost unnecessary. We also find variations and small inefficiencies in the distribution of traffic load across proxies and the access network, related to the locality and manual proxy choice. Finally, users experience an overall usable Internet access with good throughput for a free crowdsourced service.

Keywords: Community network, guifi.net, User experience, Proxy service

1 Introduction

The majority of the world's population does not have any or an adequate Internet access [12], implying that the Internet cannot provide service and reach everyone without discrimination. Global access to the Internet for all requires a dramatic reduction in Internet access costs especially in geographies and populations with low penetration [9]. Community Networks (WMNs) [17] allow local communities to build their own network infrastructures and provide affordable inter-networking with the Internet including the deepest rural communities worldwide [15]. Internet companies have also tried to address the digital divide with initiatives such as Facebook's FreeBasics [16] or the Google Global Cache. Sharing resources, such as local access infrastructure or global Internet transport, is encouraged at all levels [11,7] to lower the cost of network infrastructures and Internet services.

Among many other community networks, guifi.net exemplifies how regional communities can develop their own network infrastructures governed as a commons [2], using wired and wireless links to create a regional IP network, and sharing several Internet gateways among all their participants. These gateways are usually web proxies for Web access, the most popular traffic, but can accommodate other traffic through HTTP CONNECT, SOCKS or tunneling. Proxies, not exempt from the drawbacks of middleboxes, have also additional advantages: some content and DNS resolution can be shared in caches, and most important, proxies can protect the privacy of end users if they trust the proxy provider. Access to the Internet through Web proxy gateways relies on individuals or organisations sharing the full or spare capacity of its Internet connection with other guifi.net users. However, these crowdsourced gateway nodes have limited processing and Internet transfer capacity and might be overloaded by the demand.

In this paper we contribute an analysis of a large crowdsourced proxy service in a regional community network. A large population of C clients can browse the Web taking advantage of the aggregated capacity of a pool of P contributed web proxies, with $C > P$, spread over a regional network infrastructure, at a fraction of the cost of C Internet connections.

We first describe the guifi.net network, its proxy service and the collected datasets in Section 2. Then we analyse the service from three viewpoints: 1) service usage by end-users: patterns of usage and content in Section 3, 2) the proxy, Section 4, in terms of caching, users, performance and variability, and 3) the local network, Section 5, in terms of topology and usage. Our measurements describe the effectiveness of a simple setup of a regional network sharing a set of Web proxies in delivering free basic Web access to a large population.

2 The guifi.net Proxy Service

Guifi.net is an open, free, and neutral network built by citizens and organisations pooling their resources to build and operate a local network infrastructure, governed as a common pool resource [2]. The network infrastructure is mostly wireless [17] with a fiber backbone. Participants can extend the network to reach new locations and use it to access intranet services like the web proxy service.

The most popular application in community networks is web access and guifi.net is no exception. Web proxy nodes connected both to guifi.net and an ISP act as free gateways to the Internet to the community network users. Proxies run on simple servers and take advantage of individuals or organisations (like libraries or municipalities) offering their Internet access to other guifi.net users. Using web proxies, public entities can provide free Internet access without infringing telecom market competence regulations. While some of the web proxies are kept as a private service, 356 out of the 477 registered web proxy servers in the network (May 2016) are shared with all the network registered participants (12,500). A registered member is allowed to use any proxy of their convenience, although recommended to use one nearby. Users can select or change its choice based on quality of experience. Therefore, while some proxies may become popular and highly used, others may remain underused.

Data Collection: For our analysis we chose to study the Llucanes guifi.net zone, a region in the Osona county of Catalunya, Spain. As explained in [6], this zone is representative of other rural guifi.net networks. Furthermore, Llucanes is the only guifi.net zone with published anonymized logs for all (four) involved operational prox-

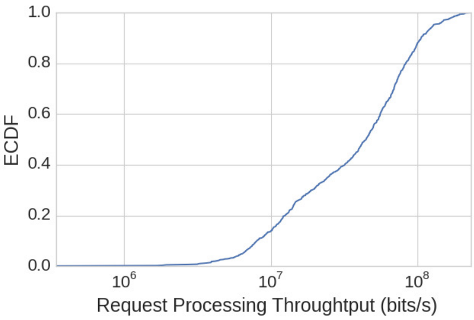

Fig. 1: Processing rate per request

Table 1: Top domains by traffic

Domain	Traffic Fraction
googlevideo	27.85%
mega	16.73%
fbcdn	5.40%
rncdn3	2.80%
nflxvideo	2.70%
xvideos	2.60%
tv3	2.54%
level3	2.51%
google	1.96%
apple	1.78%

ies. Even-day proxy log entries anoymise the client IP address and show information about the requested URLs, while odd-day proxy logs show the opposite. We assisted in the preparation and publication of these logs. The logs combined with other openly accessible information about network topology, network links and network traffic information, provide a consistent and complete view of this regional network.

3 Service Usage Viewpoint

The behavior of the users and the service can be described at macro-level as a set of time series concerning metrics that can be extracted from the monthly logs, namely bytes per request, number of requests and number of users.

The traffic time series for the aggregate set of proxies shows a daily repetitive pattern, but also strong aperiodic negative spikes, which were statistically verified as a dominant period of 1 day, and the second largest peak at 12 hours.

Service Usage: The majority of the traffic is due to a relatively small number of large requests (20% of the requests produce 97% of the traffic), while the rest of the requests present little variation in size. Additionally, as expected, the majority of the traffic (90%) is created by 15% of the users. However, in contrast to the distribution of request size, the distribution of traffic and number of requests per user varies exponentially across users. For the analysis of the service processing rate we calculated the **request processing throughput** as the bits per time elapsed for each request, depicted in figure 1, ranging from less than 10^7 for the worst 10% to at least 10^8 for more than 80% of the requests.

Content analysis: Using the even-day proxy logs we looked at the request types and target URL of the users' requests. The majority of the traffic, almost 50%, consists of HTTP CONNECT requests, which is the method to establish TCP tunnels over HTTP, mostly all HTTPS which is indisputably the main usage appearing in the logs. While for HTTP CONNECT we cannot know the corresponding content type, the most common type for the rest of the requests is the generic *application/** with 23%, followed by *video/** (19%) and *image/** (5.5%).

Table 2: Characteristics of the proxies

Id	CPU	RAM	Max Throughput
3982	Intel amd64 2-core 2.6GHz	2GB	80Mbps
10473	Intel x86 2-core 2.6GHz	0.5GB	6Mbps
11252	AMD Athlon(tm) XP 1700+	0.5GB	4Mbps
18202	Intel amd64 2-core 2.7	2GB	8MBps

Table 3: Average volume of data in all proxies and ratios in a month of logs

	Different Data (MB)			Data transferred (MB)				Ratio (/All transfrd)		
Proxy	All	Repetd	Cached	All	Repetd	Cached	Connect	Repetd	Cached	Connect
10473	606	37	9.2	1481	95	14.3	943	6.4%	0.9%	63.7%
11252	3572	1234	28	15352	5512	99	7578	35.9%	0.6%	49.4%
18202	6384	1498	151	15963	3039	253	9274	19.0%	1.6%	58.1%
3982	2542	435	55	6019	855	96	3128	14.2%	1.6%	52.0%
Avg	3276	801	61	9704	2376	115	5231	18.9%	1.2%	55.8%

The traffic for all analysed proxies in Table 1, including HTTP CONNECT, shows that the top video portal traffic occupies 36% of the traffic, which is an impressive large amount. For completeness, we mention that this is not reflected in the number of requests, therefore it is attributed on the size of the objects requested. Since video is by far the HTTP type with most traffic, it is not surprising to find that 4 out of 10 top domains are video portals. We also found that the distribution of web traffic per URL can roughly approximate a Zipf distribution, equivalent to results in [14] with domestic Internet connections.

4 The Proxy Viewpoint

In this section we investigate the capabilities and influence of the proxy servers involved. Our dataset concerns the only 4 proxies operating in the Llucanes zone. Table 2 shows the CPU and RAM characteristics of the proxy servers, as well as the nominal maximum throughput of the Internet connection they offer. They are very diverse, with great differences in Internet throughput (4–80Mbps). We also observe that proxy 11252 has the slowest combined characteristics. Despite that these servers provide other services, e.g. SNMP, the interference caused by other services is expected to be negligible.

The analysis of logs for the four proxies is summarized in Table 3. The values are averages for each proxy over a month of daily logs. The first group of columns (Different data) shows a data object storage perspective, with the amount of different data objects requested (disregarding the number of requests for each). The second group (Data transferred) shows a data transfer perspective, with the amount of traffic in each category. The third group shows data transfer ratios to the total transferred. We distinguish between “All” content, seen or transferred by the proxy, content requested repeatedly (same URL, cacheable or not), content served from

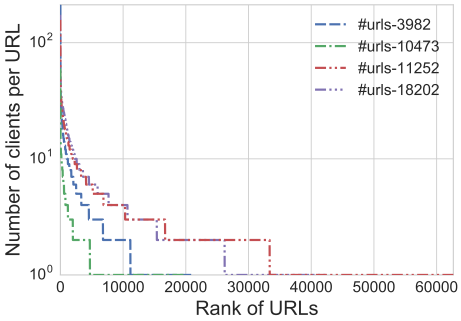

Fig. 2: Rank of URLs by number of clients requesting them per proxy

the cache (checked or not against the server), and content that is invisible (Connect method, typically HTTPS, passed through blindly).

Cache effectiveness: As introduced before, the passed-through content (HTTPS) represents the majority of the proxy traffic (49.4–64%). Although URLs repeat significantly (6.4–36% of proxy traffic), the content successfully served from the cache (after validation or not) only represents a negligible amount (1–1.6%). Considering the number of requests instead of the amount of data, despite URLs repeat often (20–41%), the content does not seem cache friendly, as cache hits only represent a very small portion (3–10%). The analysis in number of requests compared to byte count indicates that cached content usually corresponds to small objects. Bad cache performance can be attributed to characteristics of the proxy service, such as small cache size, small number of concurrent users per proxy, or to increasingly non-cacheable served content. We next look at how these apply to our scenario, claiming that non-cacheable content is the main factor affecting cache performance.

Cache size: As far as the cache size, the default allocated cache size in guifi.net proxy settings is 10GB of secondary storage, while in some proxies caching is not even enabled. However, we discovered that cached content that results in cache hits only accounts for a maximum of 151MB (if all repeated URLs were cacheable) and an average of 61MB (based on cache HITs) of data per day. In the extreme case where all content were cacheable and discounting the transparent CONNECT/HTTPS data, the amount of daily data seen (i.e. all content for all URL seen) accounts for a maximum of 1.5GB or 801MB on average, easily achievable with RAM-based caches.

Sharing across clients: Proxies can provide the benefit of sharing network resources, reusing not only HTTP content, but also reusing DNS resolution data as client web browsers delegate, or even reusing (pooling) established TCP connections among different clients. Figure 2 shows the popularity of URLs across different clients in each proxy over a month, with top values between 60 to 212 different clients accessing each given URL. The number is related to the structure of the service, with many decentralized proxies with few users each and no inter-cache cooperation, which limits the potential of sharing cached content across more users.

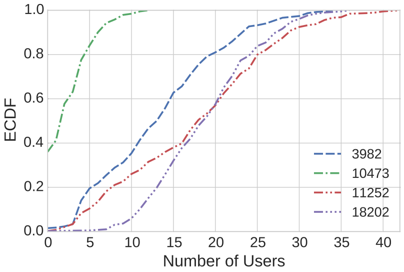

Fig. 3: Hourly average users per proxy

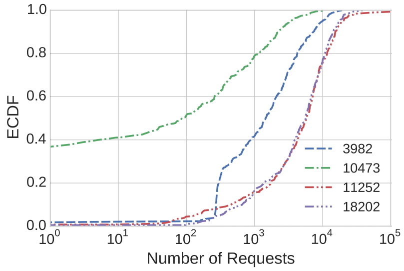

Fig. 4: Hourly average requests per proxy

Proxy selection: Users are instructed to check the public list of nearby proxies (in their network zone) in the public network management directory of the community network with shows a list of nearby proxies, including status and availability ratio, or follow the advice of trusted neighbors with previous usage experience. Therefore the choice is influenced by social factors and the reputation of each proxy, but in most cases the first choice is the nearest operational proxy with acceptable availability or reputation. Typically several nearby Web proxy services are configured in client Web browsers. As all federated proxies use the same authentication service, users are free to choose whatever proxy they prefer. The choice of proxy is rather fixed and prioritized, only switching to lower choice proxies when the first fails to reply.

Users and proxies: Figure 3 presents the distribution of the average number of users per hour. The different proxies show similar distributions, though we observe that proxy 10473 has a differentiated demand, with 40% of time without any user and a maximum of 10 users per hour. The rest of proxies, the majority of time (60%) have an almost linear distribution between 5 and 25 users, with near equally distributed values, and an average of around 17 users per hour for proxies 11252 and 18202, and an average of 12 users for proxy 3982. The different distributions among proxies is a result of preference for proximity and manual selection. To complete the picture, we found an average of 10 users in periods of 10 secs, an average of 76 different users per proxy and day, and a maximum of 254 in a month. The user's distribution among proxies has a clear impact in the distribution of the number of requests in figure 4. The ordering of proxies according to the number of users remains visible in the distribution of requests. Also, there is near-linear behavior between 20% and 60% for all proxies except 10473. For proxies 11272 and 18202 the number of requests per hour is typically between 1K and 10K requests, with a mean of 8187 and 6716 respectively. In proxy 3982 typical values are between 500 and 1K requests per hour.

Regarding the number of clients seen by a proxy every day, the values (min, average, max) range from the lowest in proxy 10473 (14, 20, 27) to the highest in proxy 3982 (59, 82, 101). These numbers reflect the spirit of a highly decentralized service with many small capacity local proxies.

Internet connection and processing performance: Figure 5 provides the distribution of the Internet connection usage per proxy, calculated as the approximate instant connection throughput of each proxy normalized by its maximum Internet

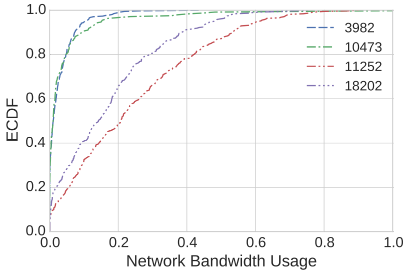

Fig. 5: Network usage per Proxy

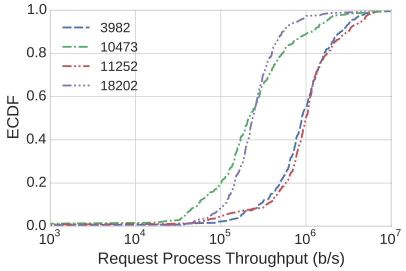

Fig. 6: Hourly average request processing throughput per Proxy

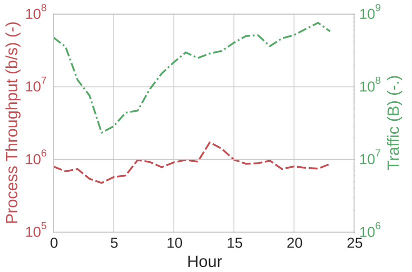

Fig. 7: Daily average request processing throughput compared to traffic

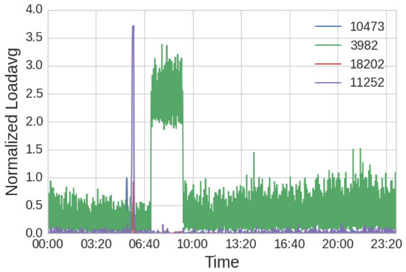

Fig. 8: Daily Median Loadavg per proxy normalized by #CPUs

throughput as provided in Table 2. All proxies show low utilization of their network resources, being approximately less than 0.3 (30%) for all the proxies for 80% of the time. Nevertheless, proxies 11252 and 18202 have significantly higher traffic. Moreover, figure 6 shows the distributions of the request processing throughput. We observe that all proxies have almost identical distribution but around different mean values, depending on the individual characteristics of the proxy. Moreover, we can see that a significant percentage (>60%) of the time proxies serve at a very narrow range of processing throughput, meaning they can offer a stable service. Even in the worst cases, the service does not suffer from extreme degradation, while remaining higher than 100Kbps 80% of the time. We also observe that for proxies 3982 and 11252, the processing throughput distribution resembles the number of requests distribution in Figure 4 possibly indicating, as before, that the proxies are not saturated.

To gain a more complete perspective we also studied the daily aggregates of the traffic, users and requests clearly observing not only the expected human daily pattern but also a clear effect of the different way each proxy receives and serves request as a result of the users' manual proxy selection. Moreover, studying the mean daily patterns, we noticed that, as seen in figure 7, the processing throughput presents very small variations implying a stable service behavior. Furthermore, the traffic volume varies more than 1.5 orders of magnitude. The fact that the processing throughput is not affected by the traffic size confirms our observation that the servers are not saturated. Addition-

Table 4: Summary of Llucanes network graph

graph	nodes	edges	degree max/mean/min	diameter
base-graph	902	914	98/2.04/1	11
proxy-clients-graph	463	472	60/2.04/1	10
backbone-graph	47	56	10/2.38/1	9

ally, in order to verify that the processing capabilities of proxies are not a bottleneck for the service, we monitored the proxies' CPU using the *loadvg* Linux metric. The results, presenting a strong daily cyclic pattern, are summarized in figure 8 that shows the daily median of the per-minute loadavg for each proxy normalized by the number of CPUs. Except from 3982, affected by other co-located network services, the proxies are not overloaded. The brief daily peak in each proxy is due to the daily restart of the proxy that includes a cache reindexing. Even at that small scale, we observed the daily cycle of human activity with preference for evenings and really reduced traffic during the first hours of the day. The pattern is visible in all the described metrics in different degrees.

From all the above we can conclude that the proxies are able to offer a stable service, with respect to the traffic load, allowing them to be used as an alternative domestic Internet connection. Moreover, in our concrete scenario, the network capacity of the proxies is underutilized assuming that no other services co-located in the host of the proxy are heavily using the Internet network capacity.

5 The Local Network Viewpoint

The local network infrastructure has also an influence in the final user experience. For the analysis we used information extracted from odd day logs that provide these details while hiding URL destinations.

Network structure: We refer to the entire zone network as the *base graph*. Moreover, we refer as *proxy-clients graph* to the part of the Llucanes network including only the nodes (clients, routers, proxies) that participate in the proxy service. Similarly to many rural community network deployments, the network consists of a small set of interconnected routers, the *backbone graph*, where each router is connected with a large number of end nodes, mostly 802.11b wireless links [17]. Users access the entire guifi.net network from the end nodes. Some of the routers act also as hosts for various guifi.net services, including the proxy service. More information concerning the network structure, hardware characteristics, and protocols used in guifi.net is available in [17].

Table 4 describes the main characteristics of the aforementioned graphs. We notice that the mean degree of the base graph and of the proxy-clients graph is very low since the end-nodes dominate the distribution of degrees. The low mean degree value in the backbone graph is more interesting though, since it implies that the majority of the routers have only two neighbors. Figures 9 and 10 provide a view of the proxy-clients graph and the backbone graph. The colors of the participating nodes and routers indicate that they are using the proxy with the same color. Moreover,

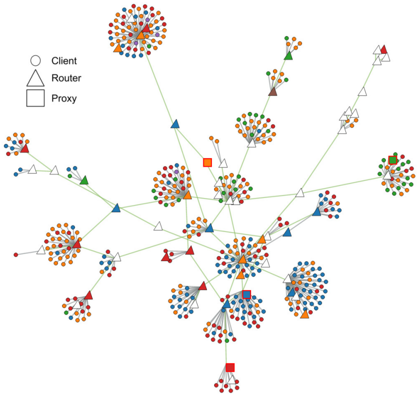

Fig. 9: Llucanes proxies and clients

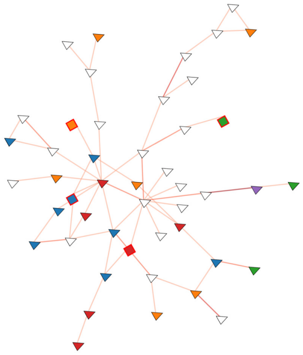

Fig. 10: Llucanes backbone

in figure 10 the darkness of the link color denotes the cost in latency for a byte to cross this link, therefore the darker the color the more expensive is the link to use.

Network usage: Due to the almost static (manual) proxy selection, the analysis of local network usage can show the effect of selection on local network usage and the perceived user experience. Towards that end, we first analyse metrics of distance between the users and the proxies. Figure 11 shows the distribution of the number of hops between users and their selected proxies. The distribution is almost uniform for 95% of the users with values from 1 to 6 hops. The remaining 5% is split between 7 and 8 hops. Nevertheless, we observe that manual choices result to a slight increase in number of hops, possibly introducing small unnecessary overheads. The latency involved, depicted in figure 12, shows a different behavior. Almost 80% of the users experience an average latency smaller than 15ms to reach their proxy. The remaining 20% lies between 20ms to 35ms. Despite the almost uniform distribution of hops, latency values vary much less, implying that during normal network conditions, the distance between the users and proxies is not significantly deteriorating the user experience for web services.

Download throughput: As described earlier, the request processing throughput is calculated using the request elapsed time, which includes the time the proxy requires until the last byte of the web object is sent to the client. Therefore, any significant local network deterioration affects the throughput behavior. Based on this observation we can utilize the request processing throughput metric for objects larger than 1MB, in order to estimate significant degradation on the user experience. Including smaller objects would give unreliable throughput results due to the noticeable influence of network buffering in the proxy, DNS caching and network latency variations for short connections. Figure 13 illustrates the individual user experience in throughput estimated for objects larger than 1MB, under the simplifying assumption that users focus on few or a single large object at a time. If so, our measures could be taken as a lower bound for the experienced individual download throughput. Median values of download throughput appear quite stable with median values ranging from 0.1Mbps to 10Mbps for different users. A quite good result for the many users of a free crowdsourced service.

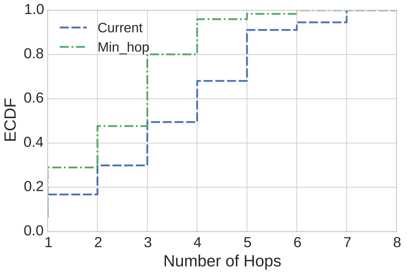

Fig. 11: Number of network hops between users and their selected proxies

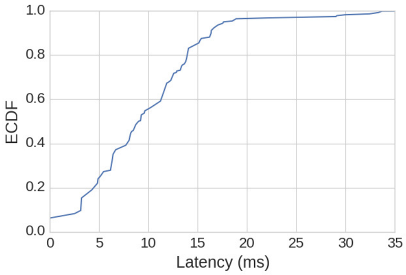

Fig. 12: Average latency between users and their selected proxies

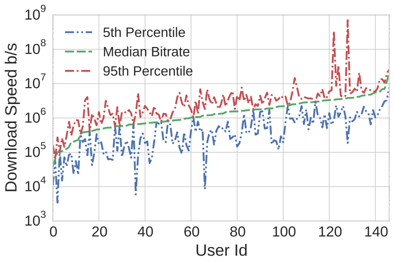

Fig. 13: Estimation of user experience throughput with objects >1MB

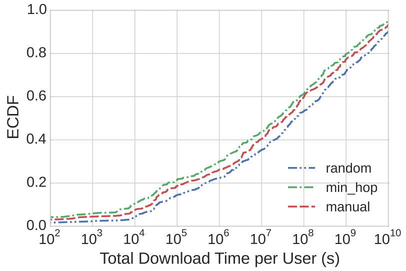

Fig. 14: User cost as sum of download times (1 month)

Furthermore, in order to show the margin for improvement in the user experience using other proxy selection strategies, we simulated the traffic of the users using a *min_ hop* and a random strategy considering local link latencies. As seen in Figure 14, the total download time of each user throughout the month in the manual selection is asymptotically better than the random selection while asymptotically worse than the min_hop selection. Considering that the proxies are not the bottleneck, this result shows that a proxy selection mechanism could improve the user experience of the proxy service. Nevertheless, we plan to extend our simulations taking into account the proxies processing and download speed.

6 Related Work

Most work on wireless networks focuses on usage traffic patterns, link level characteristics and topologies, but not user experience, e.g. MadMesh [4], Google WiFi [1] and Meraki [3] networks. In these studies, Internet access is direct instead of using proxies, and these wireless networks are homogeneous. Thus, measurement results cannot easily be compared with this. In the Google WiFi and MadMesh transfer rates are limited to 1Mbps, but 80% getting less than 80Kbps in Google WiFi. In MadMesh 80% get less than 1Mbps with 85% of the clients connected within 3 hops to Internet, comparable with our results that achieve higher speed but more hops to a web proxy.

Facebook's Free Basics [16] shows comparable performance (80-600Kbps for FB vs. 0.1-10Mbps median speeds) better in our case, despite significant differences: in clients (mobile devices vs. any device), access network (cellular mobile carrier vs. wireless fixed community network), web proxies (centralized large servers vs. distributed small servers with network locality), and web service and content providers (redesigned and optimized vs. unmodified content).

The web proxy business has changed significantly over the years. The percentage of cacheable content has been decreasing, coupled with a dramatic increase of HTTPS traffic. The performance of web proxies is not only about high-level metrics such as hit rates. Low-level details such as HTTP cookies, aborted connections, and persistent connections between clients and proxies as well as between proxies and servers have a visible impact on performance, particularly in heterogeneous bandwidth environments [8]. In [5], authors analyse a mobile network topology with a two level cache hierarchy. Their claim that a caching system can be efficient when 5.1% of traffic is suitable for caching, shows that the lower rates of caching in our case may not be that beneficial.

Wireless network user experience has been characterized previously. The first [13] focuses on web traffic and the use of proxies to access Internet content in rural areas. Five years ago, using a single high latency and slower VSAT Internet connection (64-128Kbps) obtained RTT values sometimes over 10 secs, closer to a DTN case, and cache hit rates of 43%. There are also complementary lessons, about security or that content from CDN is usually not cacheable, but the scenarios are too different. The second study [10] looks at web traffic patterns and content caching. They mention the decreasing cache hit rates over previous studies, even lower in our study 5 years later with a dramatic increase of HTTPS traffic.

7 Summary of Lessons Learned

The analysis of the guifi.net proxy service describes a crowdsourced, social solidarity driven, free basic Internet service built from many small proxy servers spread across a regional community network, contributed by locals for locals. These proxies act as gateways to Web content and DNS, that can be cached and shared among clients or act as middleboxes for HTTPS transfers, the majority of traffic. Being in the middle can also help protect the privacy of clients.

The analysis confirms the trend to non-cacheable content, small cacheable objects, and therefore small object caches that can even fit in RAM. Each proxy has a small number of clients, ranging from 14 to 101 per day. Moreover, there is a good balance of traffic and number of clients per proxy despite the manual proxy selection, driven by locality (same zone), end-user choice and advice from people living nearby. The system is simple and resilient since each proxy is independent and clients just switch to their next choice in case of failure of their proxy.

The service has satisfactory performance (0.1-10 Mbps, good client-proxy latency), without perceived Internet, access network or service congestion, despite the typical daily usage patterns. That can be attributed to the structure of the service with many small servers across a regional access network, close to end-users with locality preference. Nevertheless, scaling or coordination between services between different zones is not trivial. Future work lies in exploring optimization by automating service selection with a global perspective, under diverse cost and performance metrics.

In summary, a crowdsourced service that fulfills the goal of an usable, satisfactory and inexpensive free basic alternative Internet access service for many.

Acknowledgments

This work was partially supported by the EU funded Erasmus Mundus Joint Doctorate in Distributed Computing (EMJD-DC) (FPA 2012-0030), the EU Horizon 2020 project netCommons (H2020-688768), the Spanish government (TIN2016-77836-C2-2-R), and the Generalitat de Catalunya (2014-SGR-881). Our thanks to Davide Vega, Roger Baig and several Guifi.net members that have made this work possible.

References

1. Afanasyev, M., Chen, T., Voelker, G., Snoeren, A.: Usage patterns in an urban wifi network. IEEE/ACM Transactions on Networking 18(5), 1359–1372 (2010)
2. Baig, R., Roca, R., Freitag, F., Navarro, L.: guifi. net, a crowdsourced network infrastructure held in common. Computer Networks 90, 150–165 (2015)
3. Biswas, S., et al.: Large-scale measurements of wireless network behavior. ACM SIGCOMM Computer Communication Review 45(4), 153–165 (2015)
4. Brik, V., et al.: A measurement study of a commercial-grade urban wifi mesh. In: Internet measurement conference. pp. 111–124. (IMC) (2008)
5. Catrein, D., et al.: An analysis of web caching in current mobile broadband scenarios. In: New Technologies, Mobility and Security. pp. 1–5. (NTMS) (2011)
6. Cerdà, L.: On the topology characterization of guifi. net. In: Wireless and Mobile Computing, Networking and Communications. pp. 389–396. (WiMob) (2012)
7. European Parliament and Council: Directive 2014/61/EU on measures to reduce the cost of deploying high-speed electronic communications networks. (May 2014)
8. Feldmann, A., et al.: Performance of web proxy caching in heterogeneous bandwidth environments. In: INFOCOM. pp. 107–116 (1999)
9. GAIA WG: Global access to the internet for all research group (2016), `https://irtf.org/gaia`, [Online; accessed 14-September-2016]
10. Ihm, S., Pai, V.S.: Towards understanding modern web traffic. In: Internet measurement conference. pp. 295–312. (IMC) (2011)
11. International Telecommunication Union: Trends in telecommunication reform 2008: Six degrees of sharing (d-pref-ttr.10) (July 2009)
12. Internet Society: Global internet report 2015 (October 2015)
13. Johnson, D.L., Pejovic, V., Belding, E.M., van Stam, G.: Traffic characterization and internet usage in rural africa. In: World Wide web. pp. 493–502 (2011)
14. Maier, G., et al.: On dominant characteristics of residential broadband internet traffic. In: Internet Measurement Conference. pp. 90–102. (IMC) (2009)
15. Rey-Moreno, C., et al.: Experiences, challenges and lessons from rolling out a rural wifi mesh network. In: ACM Computing for Dev. p. 11. (ACM-DEV) (2013)
16. Sen, R., et al.: On the free bridge across the digital divide: Assessing the quality of facebook's free basics service. In: Proceedings of the 2016 ACM on Internet Measurement Conference. pp. 127–133. (IMC) (2016)
17. Vega, D., et al.: A technological overview of the guifi.net community network. Computer Networks 93, 260–278 (2015)