

3GOL: Power-boosting ADSL using 3G OnLoading

Claudio Rossi¹, Narseo Vallina-Rodriguez², Vijay Erramilli³, Yan Grunenberger⁴
Laszlo Gyarmati⁵, Nikolaos Laoutaris⁶, Rade Stanojevic⁷, Konstantina Papagiannaki⁸
Pablo Rodriguez⁹
Politecnico di Torino, Italy¹
University of Cambridge, Cambridge, UK²
Telefonica Research, Barcelona (Spain)³

ABSTRACT

The co-existence of cellular and wired networks has been exploited almost exclusively in the direction of OffLoading traffic from the former onto the latter. In this paper we claim that there exist cases that call for the exact opposite, *i.e.*, use the cellular network to assist a fixed wired network. In particular, we show that by “OnLoading” traffic from the wired broadband network onto the cellular network we can usefully speedup wired connections, on the downlink or the uplink. We consider the technological challenges pertaining to this idea and implement a prototype 3G OnLoading service that we call *3GOL*, that can be deployed by an operator providing both the wired and cellular network services. By strategically OnLoading a fraction of the data transfers to the 3G network, one can significantly enhance the performance of particular applications. In particular we demonstrate non-trivial performance benefits of *3GOL* to two widely used applications: video-on-demand and multimedia upload. We also consider the case when the operator that provides wired and cellular services is different, adding the analysis on economic constraints and volume cap on cellular data plans that need to be respected. Simulating *3GOL* over a DSLAM trace we show that *3GOL* can reduce video pre-buffering time by at least 20% for 50% of the users while respecting data caps and we design a simple estimator to compute the daily allowance that can be used towards *3GOL* while respecting caps. Our prototype is currently being piloted in 30 households in a large European city by a large network provider.

Categories and Subject Descriptors

C.0 [Computer Systems Organization]: System architectures

General Terms

Design

Keywords

Cellular, Wired, Networks, 4G, 3G, Onloading, DSL

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CoNEXT'13, December 9–12, 2013, Santa Barbara, California, USA.

Copyright 2013 ACM 978-1-4503-2101-3/13/12 ...\$15.00.

<http://dx.doi.org/10.1145/2535372.2535400>.

1. INTRODUCTION

The adoption of bandwidth hungry¹ mobile devices in the last few years has posed serious provisioning challenges for carriers worldwide. In response, a fair amount of research has gone into devising mobile data offloading [20, 14] solutions that offload data *from* resource constrained cellular networks *onto* (relatively) resource rich wired networks using Wi-Fi and Femtocells.

Such cellular network offloading strategies assume that the cellular network is constantly under strain while the wired network has abundant resources to assist. The reality, however, is that wired networks themselves have bottlenecks, first and foremost of which is the access link [21]. For example, ADSL is often constrained by the distance between the customers and the telephone exchange and has much less bandwidth on the uplink than on the downlink, thus effectively limiting the adoption of applications, specially the ones that source content from within the home.

A simple back of the envelope calculation (§2) shows that a cellular network has smaller total aggregate capacity by roughly 1-2 orders of magnitude than the corresponding wired network that covers the same geographic area. Still, at a specific location (*e.g.*, a household) and when there is free capacity, a cellular network can provide a comparable, if not greater, amount of bandwidth than the overlapping wired network. This points to opportunities when the resource-constrained cellular network can augment the typically assumed resource-rich wired network.

In this paper we present a system that “OnLoads” data from a wired network onto a cellular network for residential users with access bottlenecks, using existing mobile devices with cellular data connectivity. We refer to this as 3G OnLoading (3GOL). We first present our case for one application, *i.e.*, augmenting video-on-demand in the initial payout phase (also known as ‘Powerboosting’) and initially propose a network-integrated solution, where there is a single operator that provides both wired and cellular broadband connectivity. While such a solution simplifies economic constraints towards deployment, we still need to quantify if there is enough free capacity in the cellular network *during certain times* to offer opportunities for traffic OnLoading. We perform active measurements of available capacity over different locations and over time in a large European city to answer these questions (§3). We develop an architecture for 3GOL that can be easily deployed today (§2). We de-

¹Monthly global mobile data traffic will surpass 10 exabytes in 2016 [8].

scribe our implementation of 3GOL that does not require modification to either the residential gateway or the cellular infrastructure (§4) and evaluate its performances in the wild (§5) including the case of full video downloads as well as of picture uploads.

We then consider the case when the operator providing the wired and the cellular network service can be different. In this case, economic constraints and volume caps on cellular data plans make the problem of OnLoading harder: a cell tower might have capacity to offer, but the volume caps on individual 3G data contracts might not permit end users to use 3GOL. We use a large trace to understand typical usage patterns of users with monthly cellular data plans, giving us an idea on the amount of leftover volume that can be used for 3GOL. We show that users typically use less than 20% of their monthly allowed volume; the other 80% is *already-paid-for* volume that we can exploit with 3GOL (§6). In order to avoid cap overruns, we develop a simple estimator that given the past data usage estimates the daily safe volume that can be devoted to 3GOL (§6). We also study the amount of traffic onloaded onto the cellular network (§6) using our traces. A prototype of 3GOL is being tested in 30 households of a large European city, with the intention of a larger scale deployment later.

Main Results

- Using active measurements we find that 3GOL can increase downlink throughput of ADSL connections by $\times 2.6$ and uplink capacity by $\times 12.9$, while using 3 devices for 3G OnLoading. We see variations across locations, where rural areas seem to experience greater speedup but urban areas also have non-negligible benefits (§3).
- We implement a 3GOL prototype and quantify performance gains in the wild for two applications: video-on-demand and uploading photos. These gains tend to be consistent across the day and lead to a maximum performance improvement of $\times 3.8$, $\times 4$ and $\times 6$ for pre-buffering, downlink and uplink respectively. We find that performance gains are present even when we use 3GOL to assist a 20 Mbps downlink residential connection (§5).
- Our trace driven analysis shows that we can decrease latency for video delivery by at least 20% for 50% of the users, while respecting volume caps (§6).
- We investigate the load on the cellular network when 3GOL operates under volume caps and we find that the cellular network can cope with the onloaded traffic. We further quantify the additional load in function of the percentage of users adopting 3GOL (§6).

2. CONTEXT

In this section, we present a case for the viability for 3GOL as well as the high level description of the architecture. We first perform a back-of-the-envelope calculation to understand the capacities of the different networks (ADSL, cellular), then we study the leftover bandwidth pattern of cellular networks as well as the opportunities to use the same. We finally end this section with a high-level description of the architecture for 3GOL.

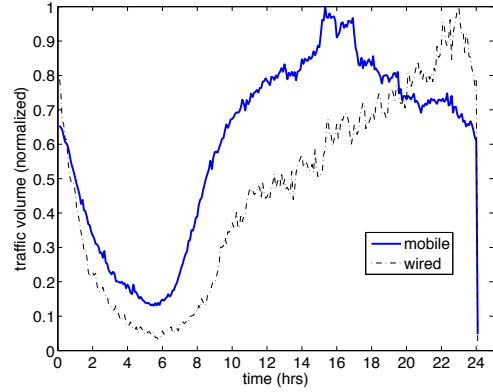


Figure 1: Traffic pattern over a day on cellular and wired network. Note the different peak timing on the two networks and the diurnal pattern in the cellular network

2.1 Comparison of capacity

If we assume that one cellular tower provides coverage to an area of 200 meters radius, and a typical population density of 35000 inhabitants per Km^2 , as in a typical downtown dense metropolitan area², then each cell offers services to 4375 subscribers. If we assume that each household has 4 people and that we have 80% penetration of ADSL connectivity, then each cell covers 875 ADSL connections. We do not consider VDSL or fiber relevant as VDSL benefits are limited in range (2km), whereas cable/fiber deployment is planned based on demographic parameters such as population density. Consequently, fiber is mainly deployed in urban areas. We focus therefore on generic ADSL; and with an average downlink speed of 6.7 Mbps³, the overall ADSL downlink capacity for the cell area would be 5.863 Gbps. Given that wired networks tend to be oversubscribed at the access, this number would be slightly lower but still on the order of a couple Gbps. The same area is covered by a cell tower with a typical 40 – 50 Mbps backhaul connectivity to the rest of the Internet. Therefore the cellular network is 1-2 orders of magnitude smaller in terms of capacity than its wired counterpart serving the same geographical region. In the uplink direction, the difference can be smaller since ADSL might involve 1/10 asymmetry in capacity in uplink/downlink.

2.2 Cellular has leftover bandwidth at certain times

From the above analysis, it results that we cannot use the cellular network at all times. We now investigate if there are some times that offer opportunities to onload traffic to the cellular network.

In Fig. 1, we plot normalized traffic from a DSLAM (Table 1) and a 3G network (data traffic, Table 1) from a single country. We do not have access to similar datasets from other countries but are in the process of acquiring them. Our first observation is that the cellular data follows a diurnal pattern as also reported by other studies [12]. This points to the fact that the cellular network is not constantly loaded,

²<http://www.demographia.com/db-citydenshist.htm>

³6.7 Mbps is the average downlink throughput reported by Netanalyzr [5].

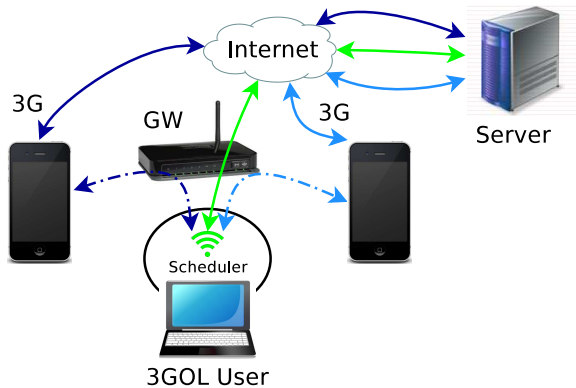


Figure 2: 3GOL architecture with one client (3GOL User), one residential gateway (GW) plus two 3G enabled devices.

and hence data from the wired network can be unloaded to the cellular. Our second observation from Fig. 1 is that the peaks for wired and the cellular are not perfectly aligned, which is confirmed in [12]. This is not too surprising – most users at home prefer to use their wired connection (via Wi-Fi or other technologies), as opposed to the cellular network. Furthermore, previous work [7] has shown that there is high spatial diversity in terms of load between cells, and some cells have left over capacity even during peak hours.

2.3 Using leftover capacity: Powerboosting

Although leftover capacity exists in cellular networks, we should carefully select the wired applications that would be allowed to use it. For example, we would not want to on-load long running bulk (*e.g.*, P2P) downloads. Applications that need onloading only for a short period of time seem ideal. Towards this direction, we develop a service that resembles the offering of PowerBoost by Comcast; delivering files at a much higher speeds for a short amount of time. Implementing PowerBoost is possible in cable networks since they have a shared-medium whose capacity can be allocated and released at short time-scales. ADSL networks, in contrast, have a dedicated twisted pair and need to adhere to long synchronization delays of the modems at the two ends. Therefore, they cannot implement PowerBoost natively. One of the main contributions of our work is the observation that by assigning excess cell network capacity to an individual customer for a limited amount of time we can effectively implement PowerBoost for ADSL with the help of the cellular network. Such temporary assistance could: a) assist the downlink of ADSL to reduce the pre-buffering time for more responsive video playout. and b) assist the uplink of ADSL to source content from home, *e.g.*, uploading photos.

If 4G is available, the concept of 3GOL is even more compelling. With the reduced latency, and the large increase of bandwidth, the period of powerboosting time might be extremely short, reducing the overhead added on the cellular network. In some countries, where the deployment of 4G is faster than fiber (like France, or India), it is largely plausible that 3GOL might be a solution to improve home connectivity.

2.4 Architecture sketch

We want 3GOL to be easily deployable and transparent for the applications. Hence, we propose to implement

it as an Over The Top (OTT) solution, where a typical service like voice or messaging is provided over the network, without using the native functionality of the network. In the OTT architecture (Fig. 2), every device participating in 3GOL is connected to the LAN managed by the residential gateway (GW) using either wireless or wired connections. First, we detail a network-integrated solution, where 3GOL is provided by a network operator that provides both wired and cellular service. Later, we describe how the architecture changes when the two providers are different.

In the network-integrated case, the single operator could provide the 3GOL traffic for free to its mobile subscribers during non-peak hours, *i.e.*, avoid metering and subtracting it from their mobile data plans. This way it would benefit its ADSL business while using available resources in the corresponding 3G network. Network integration refers primarily to the communication between the billing elements inside the wired and wireless network, such that the respective pricing policies can be applied. Each device receives the permission to transmit from the 3GOL backend server, which is revoked by the same when congestion is detected. The backend server interfaces with the 3G network monitoring system and checks whether utilization in the affected area is below an acceptance threshold. If it is, the transmission is authorized and a permit is cached for a certain duration (few minutes). Else, the transmission is denied, and the cellular device does not advertise its availability on the Wi-Fi network.

We introduce two new software components. The first one runs on the cellular device and performs the following tasks: i) implements a proxy that pipes incoming connections through the 3G network, and ii) advertises the device availability through a discovery protocol like Bonjour only if the device has an active permission by the cellular network. The second one runs on the device that needs augmentation; referred to as the “client”, and i) implements a proxy that waits for incoming connections from the applications to be augmented, ii) builds the set of admissible cellular devices (denoted by Φ) by discovering them on the Wi-Fi network, (iii) uses the GW plus the set of admissible cellular devices to transfer required contents using a multipath scheduler. The job of the scheduler is to handle multiple downloads and uploads of generic items over multiple paths. We consider the following setting: we have N available paths, with $N-1$ corresponding to Φ , and M items to download/upload, from/to a given server. We refer to the action of downloading/uploading the set of M items a transaction. The scheduler goal is to transfer the full set of M items as fast as possible, *i.e.* to minimize the total transfer time. In Sec. 6 we describe how we change the software components to handle volume caps, in the case where the network operators providing wired and cellular services are different.

3. 3GOL: WHERE AND WHEN

In the previous section, we described the context for 3GOL and outlined a high-level description of the system architecture, where 3GOL is provided as a network integrated solution. Even if 3GOL is offered by the operator that provides both wired and cellular services, we still need to understand where and when 3GOL can be used. We conduct active measurements to answer these questions.

Where? Location matters: The most common cellular technology deployed in Europe is UMTS/HSPA. Both

Name	Description
3G web traffic	HTTP traffic logs for one large cellular network operated in western European country for 24 hr period, Oct 2011, millions of users
MNO	Per-user monthly data demand generated by nearly one million users in a Mobile Network Operator.
DSLAM	Flow level information for all subscribers connected to one DSLAM in a major European city. The coverage area includes 18,000 DSL connections
Handset experiments	Downlink and Uplink throughput experiments with 10 high end smartphones in six different locations in one European city across five days

Table 1: Data sources used throughout the paper.

the HS-DSCH downlink channel (HSDPA) and the E-DCH uplink channel (HSUPA) are best-effort technologies. Their capacity is shared between users through scheduling policies aimed at optimizing the channel allocation given the available radio conditions. Hence, the per user throughput at the base-station level depends on the highly variable radio conditions as well as on the number of users simultaneously accessing the channels. To glean the available bandwidth under real operating conditions in a metropolitan city, we program 10 high-end terminal devices⁴ to download and upload 2 MB files from/to a server via the cellular interface. The description is in Table 1, where we also describe the other datasets we use in the paper. We perform throughput measurements in a western European city at six different locations with these devices assisting the existing residential broadband connection. One would expect that the degree of augmentation will be higher for residential/rural areas and during the night. We select locations and timing to test this assumption.

Given that cellular networks have middleboxes [26] and dynamic IP address assignment, we use `wget` to estimate downlink bandwidth using a random 2 MB file to avoid biases added by proxies with compression capabilities. Uplink bandwidth is tested with `iperf` [13], although we also use `netperf` [19] for validation, obtaining similar results. We first activate one 3G device and perform uplink/downlink throughput measurements. We repeat such measurements four times in sequence. Every 20 minutes we introduce a new device and run the same measurements for all active devices in parallel. All devices are synchronized using NTP. We carefully select the size of the file and the number of repetitions to collect measurements in all locations without exceeding the data plan cap (10 GB/month) to avoid throttling. For each experiment, we also record the base station serving each smartphone as well as the DSL speed, that we repeatedly measure using <http://www.speedtest.com>.

Fig. 3 shows the aggregate 3G throughput as a function of the number of active devices for the first four locations at a given time of the day (Table 2). This analysis extends previous studies of cellular networks capacity based on individual measurements [12] as it overloads the base stations serving a specific location with multiple devices accessing the channel simultaneously. *Wired capacity can be augmented by up to 14 Mbps in the downlink and 10 Mbps in the uplink for some locations.* In some of them, even when the measurements are taken at peak hour (e.g. Location 2 at 4 pm), the cellular network seems to be well provisioned, leading

⁴Samsung Galaxy S II with MIMO HSDPA Category 20 and HSUPA Category 6.

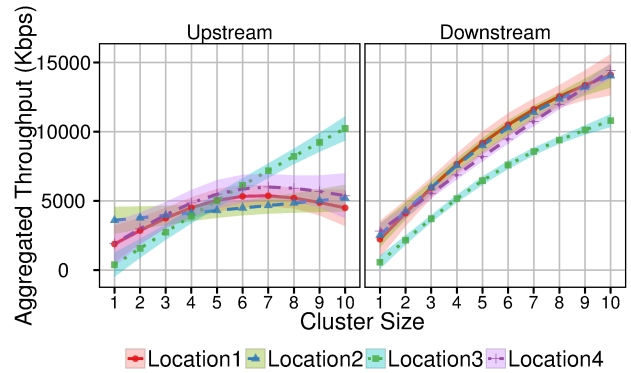


Figure 3: Aggregated throughput (uplink and downlink) for up to ten devices.

to a linear increase in downlink throughput up to ten devices. For two devices we obtain an augmentation around 4.8 Mbps, median. However, the behavior for the uplink is significantly different. For three out of the four locations, we observe a clear plateau in aggregation at five devices, nearly equal to 5 Mbps which is the maximum capacity for HSUPA (5.76 Mbps). However we note that Location 3 exceeds 5 Mbps in the uplink as shown in Fig. 3 also when the devices are connected to the same base station. It is possible that the devices are connected to different sectors of the same base station, allowing a cluster of devices to load balance naturally on a cell basis.⁵ Such a scenario is plausible since Location 3 is a hub for tourists with a large density of cellular infrastructure.

When? For the network: To understand the temporal evolution of the cellular capacity, we programmed five of the devices to download and upload the same file in groups of five, three and one device at every hour over five days. We start downloads and uploads at ten past and twenty past each hour respectively. The results, illustrated in Fig. 4, indicate that *the throughput that can be provided to a single device can be up to 2.5 Mbps both for the downlink and for uplink depending on the hour of the day.* However, as the number of devices sharing the channel increases, we observe a higher per device throughput variability. With five devices, and across all locations, the uplink and downlink per device throughput can vary from 0.65 to 1.12 Mbps and from 0.77 to 1.42 Mbps at 2 pm and 2 am respectively. While the diurnal throughput variations do exist, they are rather small, and indicate low congestion levels in most of the studied locations, even in the peak-hours.

Load balancing: The violin plots⁶ in Fig. 5 show the distribution of the capacity provided to one device by the different base stations covering the areas under study. The solid lines indicate the usual average capacity of dedicated UMTS channels under good radio conditions⁷: 360 kbps and

⁵Unfortunately, our devices can only report the base station they connect to but not the sector which is a known information on the provider side.

⁶The violin plot is similar to box plots, except that they also show the probability density of the data at different values (in the simplest case this could be a histogram).

⁷UMTS capacity can be lower in case the device is having poor coverage.

Location	Time	DSL Mbps (d/u)	3G Mbps (d/u)	3GOL/DSL (d/u)
1. Densely populated residential area (city center)	1 a.m	3.44/0.30	5.73/3.58	2.67/12.93
2. Office area at rush hour	4 p.m	4.51/0.47	2.94/1.52	1.65/4.23
3. Residential area in tourist hotspot	10 p.m	6.72/0.84	2.08/1.29	1.31/2.54
4. Sparsely populated residential area (suburbs)	1 a.m	2.84/0.45	4.55/2.17	2.60/5.82
5. Densely populated residential area (city center)	n/a	8.57/0.63	3.88/2.63	1.45/5.17
6. Densely populated residential area (city center)	n/a	55.48/11.35	2.32/1.52	1.04/1.14

Table 2: Description of the locations we measure with comparison between the DSL and the 3GOL (DSL + 3G) throughput when using three devices.




Figure 4: Aggregated throughput (downlink and uplink) in six different locations.

Cluster Size	Uplink (mean/max/sd) Mbps	Downlink (mean/max/sd) Mbps
1	1.09/2.32/0.72 Mbps	1.61/2.65/0.57 Mbps
3	0.90/2.47/0.60 Mbps	1.33/2.32/0.51 Mbps
5	0.65/2.44/0.50 Mbps	1.16/3.44/0.56 Mbps

Table 3: Average, maximum, and standard deviation of the per device throughput of a HSPA base station for the different groupings.

64kbps in the downlink and the uplink, respectively. Any additional capacity above these thresholds comes from the shared HSDPA or HSUPA channels. From the violin plots we observe that a base station can provide throughput that varies from 0.7 to 2.5 Mbps both in the uplink and downlink. In Table 3 we report the average, maximum and the standard deviation of the per device throughput obtained by a base station, noting that it decreases with the number of grouped devices both in the downlink and in the uplink.

We note that devices are associated with at least two different base stations at all locations. This suggests that when there is a denser 3G deployment, *load can be balanced across different base stations in the network, hence enabling to obtain a throughput beyond the capacity limit of a single base station.*

Main Takeaways: Our results indicate that it is possible to achieve significant benefits both in the uplink and downlink even with a small number of 3G devices, at various geographical locations. The majority of uplink throughput augmentation can be obtained with 2-3 devices during off-peak times, whereas downlink throughput seems to scale up better. We also find that such benefits are pronounced in rural areas, but not negligible in urban areas. In fact, *when we contrast the throughput of 3GOL with that of wired broadband, we find a speedup of $\times 1-3$ downlink and $\times 1-13$ uplink*

even when using just 3 devices (Table 2). Furthermore, these gains are observed not only for rural but also urban areas, and even during peak hours. Even at overloaded locations at peak hour, *it is possible to augment by 30% the downlink capacity and double the uplink one.*

4. IMPLEMENTATION

In the previous sections, we have analyzed the feasibility of 3GOL using active measurements. In this section, we describe the implementation of the architecture for 3GOL (Sec. 4.1.1).

4.1 System Implementation

We implement the architecture shown in Fig. 2 under the worst case scenario: we consider all devices connected to the gateway (GW) through the local Wi-Fi. In this way we expose our system to packet losses due to the wireless channel, which is also affected by interference generated by co-located overlapping Wi-Fi networks. In the best case, clients could be wired and cellular devices could be docked to the GW. Hence we upper bound the backhaul aggregation to the local Wi-Fi maximum goodput, which for TCP flows is around 24 Mbps and 110 Mbps for 802.11g and 802.11n respectively⁸. In this paper we do not examine the energy related issues of 3GOL. 3GOL is not likely to be continuously used over long periods of time, and also 3GOL devices are often connected for recharging while at home, hence energy consumption is not a primary concern in the 3GOL design space.

Given that most residential traffic today is HTTP [9], we select two HTTP applications to augment via 3GOL: Video-

⁸The maximum goodput may decrease according to hardware quality, channel conditions, number of transmitting stations, average bitrate, average payload size used in the BSS

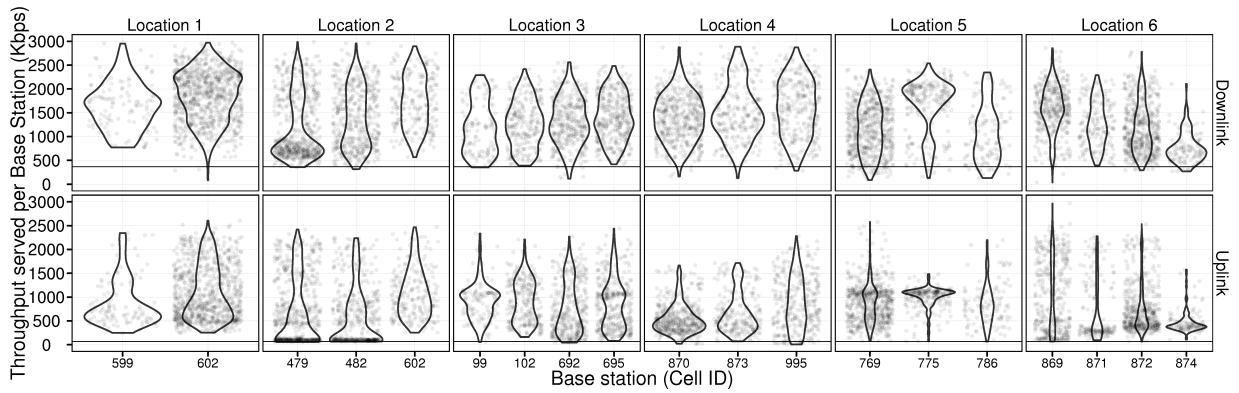


Figure 5: Throughput served per base station (downlink and uplink) in six different locations over five days.

on-Demand (VoD) for the downlink and picture upload for the uplink case. We choose a VoD application that uses Apple’s HTTP Live Streaming (HLS) protocol [3] that is supported by all Apple devices, including the Apple TV, and is also implemented in other video players (Android ≥ 4.0 , VLC ≥ 2.0).

HLS is similar to Dynamic Adaptive Streaming over HTTP (DASH) [22] and divides the video in segments of short size that are separately requested by the player with one HTTP GET request for each segment. The list of segments is retrieved through an extended M3U (m3u8) playlist and this is the first downloaded element. Next, the player sequentially requests the segments, one at a time, in the same order in which they will be required by the decoder. The video starts after a pre-buffering phase which is application dependent. We study the amount of time that 3GOL takes to fill the player pre-buffer as this gives the startup waiting time for the user.

For the uplink case we select to upload a set of pictures because it is a common use-case in social networks like Facebook, as well as in specific photo sharing applications such as Flickr and Picasa. All native clients of the aforementioned applications use multipart HTTP POST request to upload the pictures and they perform sequential uploads, one file at a time.

We implement the mobile component as an Android application that includes a basic HTTP proxy to serve the requests coming from the Wi-Fi using the 3G interface. Hence, the device must be connected both to the Wi-Fi network and to the 3G network as a client. This operation mode, that is different from the commonly implemented tethering, is not natively supported by Android and we had to root the phones to enable it. Instead, the VPN API introduced from Android version 4, could be used to perform route manipulation on the flight. We implement an HLS aware HTTP Proxy, an HTTP uploader, and a multipath-scheduler in the client component that we detail next. The client component intercepts the extended M3U (m3u8) playlist, and using the scheduler it pre-fetches the segments by performing parallel downloads. Also the HTTP uploader uses the scheduler to perform parallel multi-part POST requests to upload a set of selected picture on a web server. Both client components can easily be installed by the user as they do not require the server to change neither specific drivers or protocols. Fur-

thermore, this implementation is completely transparent to the residential gateway.

4.1.1 Scheduler

Since we are dealing with devices having different capacities, bandwidth variability over time, TCP flows over 802.11 and HSPA channel allocation schemes, the application level goodput of each path could vary in time, especially for paths over the cellular network. A first option is to use a round-robin scheduling policy. However, the peak capacity of the ADSL link is generally very different from the peak capacity of HSPA and hence round-robin cannot be expected to maximize gains. One can use history to predict future conditions and to schedule. However, estimating available capacity under rapidly changing network conditions can result in inaccurate estimates [18].

Hence we design a simple greedy scheduling algorithm that can be described as follows. First, an item is assigned to each path. Then, if there are any remaining items ($M \geq N$), they are scheduled by order, on the first available path. It is easy to see that this simple greedy scheduling policy leads to all paths being busy, achieving maximum throughput and minimum download/upload time. Optimal performance will be achieved if all paths end at the same time, which is very unlikely to happen in our case. Hence, when all items have been already scheduled and a path becomes idle before the transaction is completed, we reassign the oldest scheduled item among the ones being transferred by the other $N - 1$ paths. We keep doing this until the transaction ends, *i.e.*, until all the items have been transferred. The amount of bytes wasted is upper bounded by $(N - 1)S_m$, where S_m is the maximum item size. However, this waste is generally much smaller than this upper bound as when a rescheduled item completes, all other ongoing transfers of that item are aborted.

We could modify the scheduler to cover also the play-out phase, but given the wide amount of proposals in this area [15], we leave this extension as future work.

5. EVALUATION OF 3GOL

We present results of evaluating a prototype implementation of 3GOL as described in Sec. 4.1. We first look at the general performance of the scheduling algorithm, and then look more specifically on the applications described in Sec. 4.1 by evaluating them “in the wild” at different res-

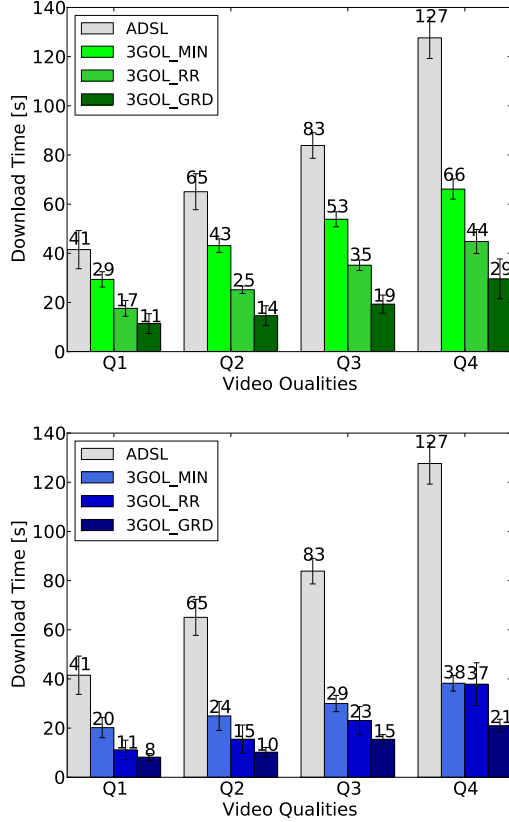


Figure 6: Scheduler comparison in downloading an HLS video lasting 200s with a 2Mbps ADSL using one device (top) and two devices (bottom).

identical locations. The 3GOL setup we use in all evaluations consists of a notebook client running Ubuntu 10.04 LTS, and up to two Android 2.3 smartphones⁹ and we do not introduce any change in the residential gateway. We use 802.11n compliant devices. Furthermore, we use a dedicated well provisioned web server, featuring a stable bandwidth of 100 Mbps in download and 40 Mbps in upload, and we disable caching on the content we transfer.

5.1 Scheduler Performance

We compare our greedy scheduler (GRD) with two common schedulers: round robin (RR) and minimum time scheduler (MIN). The round robin scheduler cyclically assigns one item to each path, while the minimum time scheduler assigns the items to the path that minimizes the estimated transfer time, computed by using the estimated available bandwidth of each path. For the MIN scheduler we assign the first N items in a round robin fashion to initialize and then estimate the bandwidth using exponential smoothing filtering. We set the filter parameter to 0.75 to maintain a high level of agility.

First we setup a 3GOL test environment using an ADSL line at 2 Mbps downlink and 0.512 Mbps uplink. We measure the average download time of an HLS video with different resolutions. The video sequence is the sample HLS video

⁹Same devices, Galaxy S2, used in Sec. 3.

Location	DSL (d/u)	3G Signal (dBm/ASU)
loc1	6.48/0.83 Mbps	-81/16
loc2	21.64/2.77 Mbps	-95/9
loc3	8.67/0.62 Mbps	-97/8
loc4	6.20/0.65 Mbps	-89/12
loc5	6.82/0.58 Mbps	-89/12

Table 4: Measured ADSL speed (downlink/uplink) and 3G signal strength of all locations evaluated with 3GOL.

provided by Apple (bipbop). We do not change the segmentation that was set to 10 seconds of video for each segment but we set the video duration to 200s, which is the median video length of a YouTube video [2]. We also keep the original qualities of the video (Q1=200 kbps, Q2=311, kbps, Q3=484, kbps, Q4=738, kbps) as they reflect commonly used bitrates according to the same study. Since the aforementioned qualities are lower than the ADSL speed, we are able to play the video without pauses. For each video quality we run 30 repetitions using first the ADSL alone, and then enabling 3GOL with each scheduler, one at a time. We repeat this cycle using both one and two phones. In order to minimize bandwidth fluctuations both in the Wi-Fi and in the 3G network, we start the scheduler comparison at night (1:00 am) In Fig. 6 we plot the results obtained, reporting both the average and the standard deviation for each case. For all video qualities we observe the same trend: the greedy scheduler is the best one, followed by the RR and the MIN scheduler. The MIN scheduler performs the worst due to the high variability of the available bandwidth of phones. The high variability of channel conditions results in poor estimates, leading to suboptimal decisions. Changing filter and/or sampling criteria was not helpful in improving the performance of the MIN scheduler. As expected, the benefit does not linearly scale with the number of phones used for the same reasons mentioned in Sec. 3

5.2 3GOL evaluation in the wild

We perform extensive evaluations of 3GOL at five different residential locations, both for video download and for photo uploads. Since we want to extend our sample points, we consider a different set of locations from the ones reported in Sec. 3 in the same city. We report in Table 4 the ADSL speeds, repeatedly measured as in Sec. 3 before starting each evaluation, as well as the signal strength measured for each location. We measure total upload/download time achieved with and without 3GOL, and for VoD we also study the pre-buffering time, *i.e.*, the measured delay from the initial request of the video to the first frame displayed by the player. As the amount of pre-buffering is application dependent, we vary it throughout the experiments.

We first run the experiment with one cellular device then with two devices, each time starting with devices from idle (3G) first and then from a connected mode (H)¹⁰. For VoD, we vary the video quality from Q1 to Q4 in each case and for each quality we vary the pre-buffer amount from 20% up to 100% of the video length, which is the inner most

¹⁰To bring the phones into a connected mode (FACH or DCH) we issued a train of ICMP packets spaced at 0.1 s before starting the 3GOL transaction, interrupting them right afterwards.

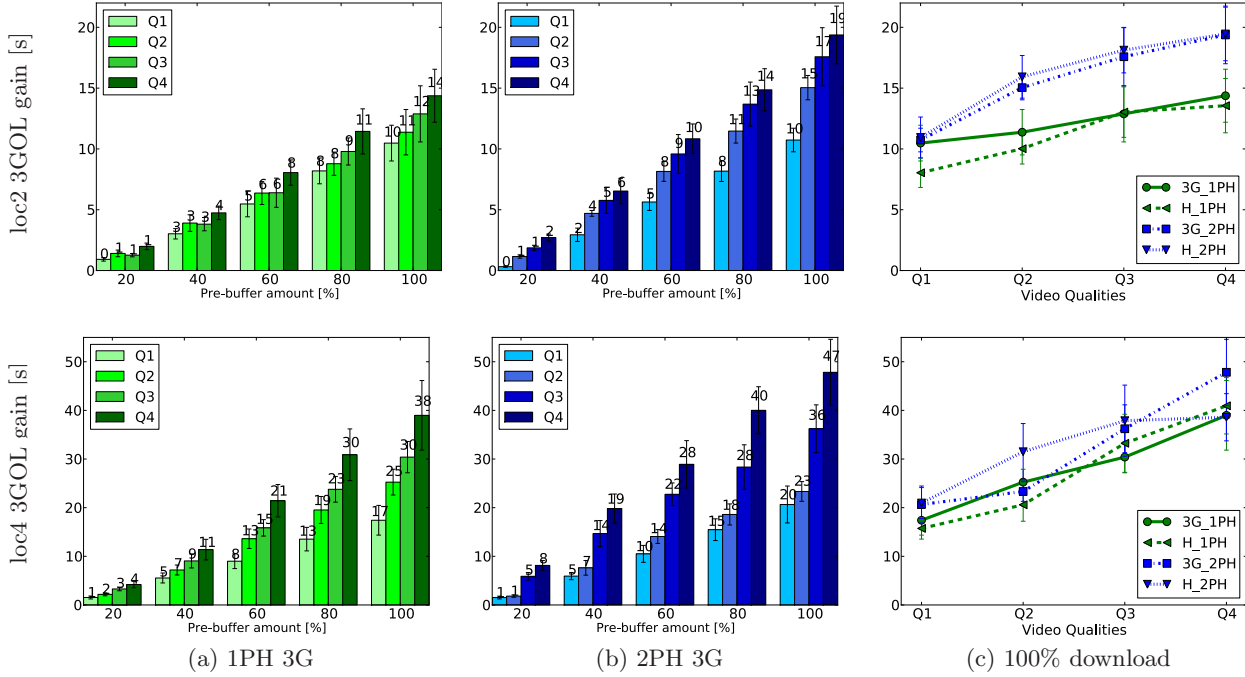


Figure 7: Video 3GOL performance comparison between location with fastest (loc2) and slowest (loc4) ADSL, using one (1PH) or two (2PH) phones, and starting from IDLE (3G) or from a connected mode (H).

loop. For each setting we repeat the experiment 30 times, first running it without 3GOL and then enabling it using the greedy scheduler. On average, the duration of both the download experiment and of upload one was approximately 12 h, depending on both ADSL and device speeds. In all locations, we run both download and upload during weekdays at daytimes, starting each one around 9.00 am. Note that we evaluate 3GOL under worst case conditions according to the findings reported in Sec. 3. Fig. 7 shows 3GOL gains, computed as the reduction in seconds with respect to the ADSL alone, at the fastest and the slowest location. The former corresponds to the location with the fastest ADSL (loc 2 with 21.64 Mbps), the latter to the location having the slowest one (loc4 with 6.2 Mbps downlink).

We observe that 3GOL gain increases with video quality as well as with pre-buffer amount. Indeed, both the number of segments and the size of each segment amplify the benefits brought by 3GOL, when increased. As we test 3GOL starting from very few pre-buffering segments (min. 4, max. 20) as well as from small segment size (min. 0.2 MB, max 0.95 MB), the observed trend suggests that 3GOL could always be beneficial for HLS and similar applications. From Fig. 7 it is evident that the addition of a second device does improve the overall performances, up to +35% of the best gain with 1 phone for location 4 and up to +26% for location 2. However, the obtained gains are high with even one device, leading to a reduction of 14s for location 4 and 38s for location 2.

In Fig. 7 we show the performance improvement when devices start from a connected state. The expected performance boost, due to the absence of the channel acquisition delay, does not clearly appear. This is partly due to the higher variability of environmental conditions. However, it

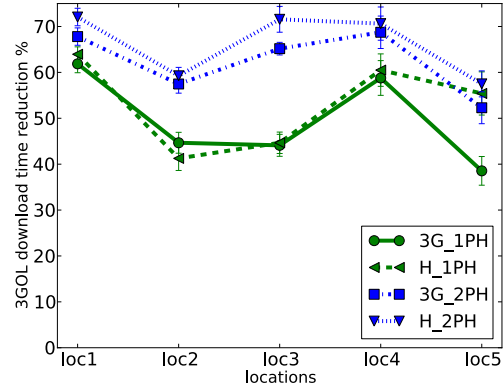


Figure 8: 3GOL % total video download time reduction in using one (1PH) and two phones (2PH) starting from idle (3G) and from a connected mode (H).

can be noted that when present this gain is limited and it shrinks as the transaction duration increases. We can infer then, that the channel acquisition delay has little or no bearing on the performance of 3GOL for this application. Considering all pre-buffer setting from 20% to 80% across all locations, corresponding to an average and maximum speedup of $\times 2.1$ and $\times 3.8$ respectively, we obtain an average transaction time reduction of 47%.

Fig. 8 shows the reduction achieved by 3GOL in downloading the entire video for all locations, where each point is the average of the percentage reduction across all video qualities. Reductions span from a minimum of 38% up to a maximum of 72%, corresponding to a speedup ranging

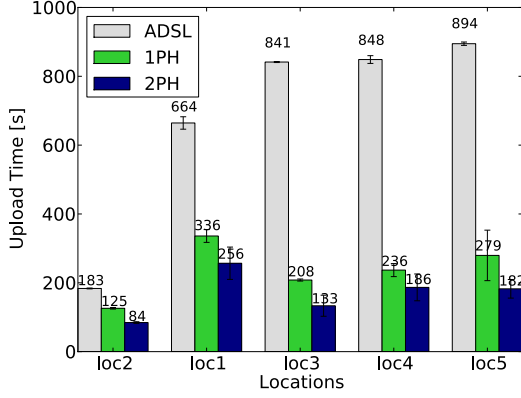


Figure 9: ADSL vs 3GOL upload time using 1 and 2 device starting from idle.

from a minimum of $\times 1.5$ up to a maximum of $\times 4.1$. We can see the benefit of adding the second device and of starting with devices on a connected mode. The second device always aids in increasing download time reductions, ranging from $+5.9\%$ up to $+26\%$ with respect to the case where only one device was used, while starting from a connected mode mostly brings marginal gains.

Upload:

We evaluate upload at the same five locations and using the same methodology. We repeatedly upload a set of 30 pictures with average size of 2.5 MB and standard deviation of 0.74 MB. We obtain these values from a set of 200 pictures taken with iPhone 5 and iPhone 4S, that are the most used devices for uploading photos to Flickr. Compared to the download case, we increase the number of items (from 20 to 30) we transfer, as well as their size (from a maximum of 0.9 MB up to a maximum of 3.9 MB).

The limited ADSL uplink bandwidth in our locations range between 0.58 Mbps and 2.77 Mbps resulting in a huge 3GOL gain. Using one device the total upload time is reduced from 31% up to 75%, corresponding to a speedup between $\times 1.5$ and $\times 4.0$, while using two devices the upload times decrease from 54% up to 84%, corresponding to a speedup between $\times 2.2$ and $\times 6.2$.

Also in this case, gains are not directly proportional to the number of devices used, and one device is enough to achieve high gains. Fig. 9 shows the average total upload time obtained with ADSL and with 3GOL, both using one and two device starting from idle.

To exploit the bandwidth on multiple paths one can use MP-TCP [11] as well, although then both client and server need to support it. This requires the presence of a third party server that handles MP-TCP flow for legacy TCP stack. We experimented with MP-TCP and it provided no benefit due to the issues probably related to the Coupled Congestion Control (CCC) algorithm of MP-TCP that is not optimized for wireless use yet. We omit these results for brevity. However, with the arrival of Android version of MP-TCP stack¹¹ as well as on Apple iOS7¹², we hope to revisit and hopefully obtain better performance.

¹¹<https://github.com/mptcp-galaxy2>

¹²<http://perso.uclouvain.be/olivier.bonaventure/blog/html/2013/09/18/mptcp.html>




Figure 10: The distribution of fraction of used cap. 40% of customers use less than 10% of their cap. 75% of customers use less than 50% of the cap.

6. DEALING WITH MULTIPLE PROVIDERS

Up to this section, we have been looking into 3GOL as a network integrated service, offered by a provider with both cellular and wired footprint. The fundamental technology behind 3GOL is, however, not limited to such a scenario since it does not require explicit cooperation between the two networks beyond billing for cellular data discounting. In this section, we explore the scenario with different cellular and wired operators. In such a case, one needs to be careful using 3G data in order to avoid penalties associated with exceeding the enforced cellular data plans [23]. Second, due to caps, users may find it difficult to allocate volume to 3GOL and may find it burdensome, leading to ‘price of uncertainty’ [4]. Third, even if users have volume they can use towards 3GOL, we need to make sure the load 3GOL puts on the cellular network does not hurt the primary 3G service. We deal with these issues next.

Do users have spare volume under the cap?

We analyzed data collected from a medium sized Mobile Network Operator (MNO) (Table 1). The data set contains mobile data sessions by all mobile broadband customers. In Fig. 10 we show the distribution of data usage as a fraction of the contracted cap. We find that 40% of the customers use less than 10% of their cap, and 75% of the customers use less than 50%. On average, it would appear that 20 MB of volume can be used per device per day by 3GOL. Similar observations about users having disposable volume have been reported elsewhere [24].

How to allocate volume towards 3GOL?

In order to alleviate the burden of the user choosing the allowance to be used for 3GOL, we propose a simple estimator with a ‘guard’ parameter that can minimize volume cap overruns. To compute the suggested 3GOL allowances $3GOL_a(t)$ we use a simple technique that estimates the free capacity at month t of user u . We compute the average free capacity over τ months prior to t :

$$\bar{F}_u(t) = \frac{\sum_{s=1}^{\tau} F_u(t-s)}{\tau}.$$

To minimize cap overruns, we set $3GOL_a(t)$ to the estimated free capacity $\bar{F}_u(t)$ discounted by a guard/margin:

$$3GOL_a(t) = \bar{F}_u(t) - \alpha \bar{\sigma}_u(t),$$

where $\bar{\sigma}_u(t)$ is the sample standard variation of the free capacity during τ months prior to t and α is a tunable guard parameter. By running this estimator on the MNO dataset,

we find that using $\tau = 5$ and choosing $\alpha = 4$ allows around 65% of the available free capacity to be used by 3GOL with expected overrun time of under 1 day per month overall. Hence, there is room for exploiting the free/unused 3G capacity with no or very little disruption of the primary 3G service, for the user. We consider the load on the network next.

We can easily modify the architecture presented in Fig. 2 to handle caps. The component running on the cellular device can track 3GOL data usage $U(t)$ and estimate the 3GOL allowance $3GOL_a(t)$. If the available quota $A(t) = 3GOL_a(t) - U(t)$ is greater than zero, the device advertises itself. All devices that advertise themselves, become part of the admissible set Φ . Thus, we need no input from the network.

Load on the cellular network

When users operate under caps, we first need to check if users can still see a tangible benefit. Further we need to ensure that the cellular network does not get overloaded, even if caps are being respected.

We use trace-driven analysis to quantify the benefits even under caps. The trace we study is collected from a DSLAM in a major metropolitan city over 24 hours in April 2011. The ADSL speeds were 3 Mbps. The data contains HTTP transactions, (userid, time of request (UTC), URL) from which we separate out all sessions pertaining to video. We find that 68% of all users in the dataset have seen at least one video. The data contains the HTTP response headers, from where we can obtain the size of the video files requested. We note that users may not watch the entire video that is requested, however our choice of using the size of the entire file is *conservative* as we shall demonstrate.

In our dataset, we find that an end user¹³ views 14.12 videos (mean, 6 median, 30.13 std) per day. Therefore, if we aim to accelerate every video by even 20 MB (as obtained above) would lead to a daily consumption of 280 MB for the 3G data plan, thus being clearly unacceptable. Thus, we focus on using 20 MB per device per day that corresponds to 600 MB per device per month, the average amount of free/unused capacity in the MNO dataset (Table 1). With two devices, this gives an available daily 3GOL budget of 40 MB, that we use to boost any/all videos the end-user views. We plot the savings *per-user* over using just DSL ($\frac{DSL \text{ latency}}{3GOL \text{ latency under budgets}}$) in Fig. 11 (a). We find that 50% of the users can see at least 20% speedup, while 5% of the users can see a speedup of 2. We note here that our results are conservative. Indeed, we can accelerate via 3GOL only the bytes required to fill the playout buffer, achieving an earlier playout for the user.

There is no fixed size of the playout buffer but the general consensus is that it is an order of magnitude smaller than the file size. If the average size of a video is 50 MB [2], and the playout buffer is 5 MB, then 3GOL can speed up 8 videos per day per user—a tangible benefit. From Sec. 3, we saw an increase in capacity with more devices, leading to even higher gains. Likewise, with LTE deployed, we can expect much higher benefit to the end user, *even* with budgets in place.

Now we turn to quantifying the resulting impact to the 3G network in terms of load. Using the same DSLAM dataset, we compute the amount of traffic that would be “onloaded”

to the 3G network if we were to accelerate the *first* video that could benefit from 3GOL (with a size greater than 750 KB, that would require more than 2 seconds on DSL). We further assume we have two 3G devices (with HSPA+) that can be used to accelerate up to 20 MB each per day. The represented geographical area would typically be covered with 2 towers (Sec. 2.1).

In Fig. 11 (b) we show the traffic onloaded to the cellular network and compare it with the backhaul capacity of the cellular deployment that could cover a similar geographical area (2 base stations with 40 Mbps backhaul). We note that if 3GOL would operate without caps, then the 3G network will be guaranteed to be overloaded. However, when 3GOL operates within caps then the additional load introduced on the 3G network could be reasonable. We find that on average, a user would onload 29.78 MB per day, if operating with caps and using two devices. Note again that our results are conservative since we accelerate the entire video transfer, whereas prior work has demonstrated that only 10% of the whole file is downloaded [2].

In the previous analysis, we did not consider the existing load on the mobile network. We also focused on one application – video-on-demand. Taking this further, we evaluate how much traffic will be added to the *existing* load of the cellular network in function of the fraction of users adopting 3GOL. As before, we consider a 20 MB daily allowance per user. Using the MNO dataset (Table 1), we plot (Fig. 11) the relative increase of traffic due to 3GOL, assuming a uniform distribution of 3GOL users the network customer base. We observe that the relative increase in aggregate traffic is relatively modest, when 3GOL adoption is low. In the case of 100% adoption, the increase in traffic is around 100%. Thus if all users adopt 3GOL consuming 20 MB per day via 3GOL, the 3G traffic is expected to double. This increase can be undesirable and hence in this case 3GOL should be provided as a network integrated service (we described such a service in Sec. 4.1.1).

If we assume that 3GOL demands are generated based on the DSLAM flow-level trace, as described above, then if we take into account the non-alignment of peaks from Fig. 1, the superposition of existing 3G traffic and 3GOL traffic onloaded will benefit from this non-alignment. The relative traffic increase due to 3GOL in peak-hour is smaller than the relative increase of the aggregate traffic, albeit the difference between the two is rather small; see Fig. 11.

We remark that by having a network integrated solution, described in Sec. 4.1.1, 3GOL traffic will not harm existing 3G traffic. The previous analysis shows that as long as the 3GOL traffic remains below the threshold set by existing leftover capacity, there will be a modest traffic increase, as we are using already-paid-for capacity. However, we argue that, as with latest versions of peer-to-peer [10], OTT services have the incentive to be friendly with the network and avoid tension that could lead to rate-limiting or blocking. We therefore expect that successful OTT realizations of 3GOL will either self-moderate their consumption or reach out to 3G providers for integration.

7. RELATED WORK

The key idea underpinning our system is to share existing cellular connection(s). Sharing unused cellular bandwidth between mobile connections was proposed in Shair [17], where the authors propose a DTN-based approach built on blue-

¹³a DSL “subscriber”




Figure 11: CDFs of improvement of 3GOL in terms of latency per video: (a) budgeted with 40 MB daily allowance, (b) Load on cellular network by using 3GOL (budgeted) and 3GOL (unbudgeted). Y-Axis is logscale. Solid horizontal line is total backhaul capacity (2×40 Mbps), (c) Relative 3G traffic (total and peak-hour) increase due to 3GOL as a function of the fraction of users adopting 3GOL.

tooth to share bandwidth. Our goal is different, we augment existing wired connections with one or more cellular connections, while being agnostic of the underlying technology.

Recent work on exploiting different interfaces to increase overall performance includes [6, 1]. Sivakumar et al. [6] describe a method to utilize all available interfaces on the same mobile device. This idea of channel bonding is not new and has been proposed for Ethernet, DSL and Wi-Fi as a means to increase throughput by combining available network interfaces. Badam et al. [1] use a reliable but slow wired connection to exchange control information about a lossy but fast wireless connection in order to obtain a reliable and high speed network. Our solution is different, as we focus on augmenting the bandwidth of wired connections with that of cellular connections.

Our work shares specific commonalities in exploiting off-peak bandwidth [16], specifically on how to transfer data while keeping bandwidth costs low. 3GOL utilizes cellular bandwidth whenever available and, by performing active measurements as well as large datasets analysis, we show that it is feasible to offload data onto the cellular network without overloading it. The closest work to ours, is the work by Rodriguez et al. [25]. We extend this work in significant ways: we explicitly deal with economic constraints like volume caps, we perform extensive measurements, we implement a fully working prototype that we test in the wild. Furthermore, we also present arguments about the future viability of 3GOL, that can be seen as an appealing long-term technology.

8. CONCLUSIONS

In this paper, we have proposed 3G Onloading, a strategy that uses the cellular network to improve application performance in the home, when constrained by the wired network. Our contribution is in studying the feasibility of 3G Onloading, in terms of technical and economic factors. Using multiple data sets, and active experiments in urban and rural locations, we have demonstrated that the benefits to application performance could be significant with reasonable overhead to the cellular infrastructure. Such overhead could further be minimal if 3GOL was rolled out as a network integrated service, offered only when the cellular infrastructure is lightly utilized. We develop a 3Gol prototype

and test the performance in the wild. We find x4 and x6 speedups to applications needing downlink and uplink.

Acknowledgements

The authors would like to thank the anonymous reviewers and our shepherd Ruben Merz for constructive feedback.

9. REFERENCES

- [1] D. Han K. Papagiannaki D. G. Andersen S. Seshan A. Badam, M. Kaminsky. The hare and the tortoise: taming wireless losses by exploiting wired reliability. In *ACM MobiHoc 2011*.
- [2] M. Munafò T. Ruben S. G. Rao A. Finamore, M. Mellia. Youtube everywhere: impact of device and infrastructure synergies on user experience. In *ACM IMC 2011*.
- [3] Apple. HTTP Live Streaming - draft-ietf. <http://tools.ietf.org/html/draft-pantos-http-live-streaming-07>.
- [4] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. The price of uncertainty. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 285–294. ACM, 2009.
- [5] B. Nechaev V. Paxson C. Kreibich, N. Weaver. Netalyzr: illuminating the edge network. In *ACM IMC 2010*.
- [6] R. Sivakumar C-L. Tsao. On effectively exploiting multiple wireless interfaces in mobile hosts. In *ACM CoNEXT 2009*.
- [7] S. Lu H. Luo H. Li C. Peng, S-B Lee. Traffic-driven power saving in operational 3g cellular networks. In *ACM MobiCom 2011*.
- [8] Cisco. Cisco Mobile Data Traffic Forecast. http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-520862.html.
- [9] V. Paxson M. Allman G. Maier, A. Feldmann. On dominant characteristics of residential broadband internet traffic. In *ACM IMC 2009*.
- [10] A. Krishnamurthy Y. Liu A. Silberschatz H. Xie, Y. R. Yang. P4P: Provider portal for applications. In *ACM SIGCOMM 2008*.

- [11] Huaizhong Han, Srinivas Shakkottai, C. V. Hollot, R. Srikant, and Don Towsley. Multi-path tcp: a joint congestion control and routing scheme to exploit path diversity in the internet. *IEEE/ACM Trans. Netw.*, 14(6):1260–1271.
- [12] P. Barford J. Sommers. Cell vs. wifi: on the performance of metro area mobile connections. In *ACM IMC 2012*.
- [13] Jon Dugan , et al. iperf project. <http://sourceforge.net/projects/iperf/>.
- [14] Y. Yi I. Rhee S. Chong K. Lee, J. Lee. Mobile data offloading: how much can wifi deliver? In *ACM CoNEXT 2010*.
- [15] I. Stavrakakis N. Laoutaris. Intrastream synchronization for continuous media streams: a survey of playout schedulers. IEEE.
- [16] X. Yang P. Rodriguez N. Laoutaris, M. Sirivianos. Inter-datacenter bulk transfers with NetStitcher. In *ACM SIGCOMM 2011*.
- [17] K. Xu J. Crowcroft V. O. K. Li P. Hui, R. Mortier. Sharing airtime with shair avoids wasting time and money. In *ACM HotMobile 2009*.
- [18] E.W. Biersack P. Rodriguez, A. Kirpal. Parallel-access for mirror sites in the internet. In *INFOCOM 2000*.
- [19] Rick Jones. netperf project. <http://www.netperf.org/netperf/>.
- [20] C. Joe-Won Y. Im M. Chiang S. Ha, S. Sen. Tube: Time-dependent pricing for mobile data,. In *ACM SIGCOMM 2012*.
- [21] N. Feamster R. Teixeira S. Crawford A. Pescapè Antonio S. Sundaresan, W. de Donato. Broadband internet performance: a view from the gateway. *ACM SIGCOMM 2011*.
- [22] I. Sodagar. The mpeg-dash standard for multimedia streaming over the internet. *MultiMedia*, 2011.
- [23] S Soumya et al. Pricing data: A look at past proposals, current plans, and future trends. *CoRR*, abs/1201.4197, 2012.
- [24] Rade Stanojevic, Vijay Erramilli, and Konstantina Papagiannaki. Cognitive bias in network services. In *HotNets*, 2012.
- [25] Narseo Vallina-Rodriguez, Vijay Erramilli, Yan Grunenberger, László Gyarmati, Nikolaos Laoutaris, Rade Stanojevic, and Konstantina Papagiannaki. When david helps goliath: the case for 3g onloading. In *ACM HotNets*, 2012.
- [26] Q. Xu Z. Mao M. Zhang Z. Wang, Z. Qian. An untold story of middleboxes in cellular networks. *ACM SIGCOMM 2011*.