

uCap: An Internet Data Management Tool For The Home

Marshini Chetty¹, Hyojoon Kim², Srikanth Sundaresan³, Sam Burnett², Nick Feamster⁴, and W. Keith Edwards²

¹University of Maryland, College Park, marshini@umd.edu

²Georgia Institute of Technology, {joonk; keith}@gatech.edu, sburnett@google.com

³International Computer Science Institute, srikanth@icsi.berkeley.edu

⁴Princeton University, feamster@cs.princeton.edu

ABSTRACT

Internet Service Providers (ISPs) have introduced “data caps”, or quotas on the amount of data that a customer can download during a billing cycle. Under this model, Internet users who reach a data cap can be subject to degraded performance, extra fees, or even temporary interruption of Internet service. For this reason, users need better visibility into and control over their Internet usage to help them understand what uses up data and control how these quotas are reached. In this paper, we present the design and implementation of a tool, called *uCap*, to help home users manage Internet data. We conducted a field trial of *uCap* in 21 home networks in three countries and performed an in-depth qualitative study of ten of these homes. We present the results of the evaluation and implications for the design of future Internet data management tools.

Author Keywords

Data caps; bandwidth caps; home networking tools

ACM Classification Keywords

H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION

Internet service providers (ISPs) are instituting a variety of policies to discourage heavy use to cope with network congestion [10]. One of the more prevalent—and controversial—policies is to impose a data cap on a subscriber; such a cap limits on the amount of data that a customer can consume during a billing cycle. A user who exceeds the cap can face higher usage fees, degraded performance, and even disruption of service [3,12]. Many ISPs around the world already impose data caps on their subscribers [25]. There is even some speculation that the Internet may now be subject to pricing tiers [19], a practice already present in larger “backbone” ISPs in the Internet core [33]. Data caps and other forms of usage-based pricing

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.

CHI 2015, April 18-23, 2015, Seoul, Republic of South Korea.

Copyright 2015 ACM 978-1-4503-3145-6/15/04...\$15.00.

<http://dx.doi.org/10.1145/2702123.2702218>

have angered consumers [27]. For example, Comcast, a major United States (US) ISP, announced in May 2014 that it would institute “data caps” on all of its customers, and later withdrew the announcement after public backlash. More importantly, data caps have highlighted the fact that users have no idea about how much data they are using, let alone how to manage their Internet data usage.

Specifically, users need more visibility into and control over how the applications they use consume their available data budgets. Without better tools, users cannot manage how (and when) data-hungry applications such as video, automated backup, and other cloud-based services use the network [15] and can sometimes make decisions that have more serious ramifications, such as forgoing critical software updates simply because they might consume a large (or unknown) amount of data [20].

In this paper, we report on the user-centered design, implementation, and evaluation of *uCap*, a data cap management system that we deployed in 21 home networks in three countries (South Africa, India, and the United States). In addition to the empirical measurements that we gathered from these homes reported in Grover *et al.* [14], we performed a qualitative study in ten of the homes to evaluate which aspects of the tool users found most effective. *uCap* runs directly on a user’s home router. It provides real-time visibility into and control over network traffic usage in the home. *Visibility* exposes the websites, devices (and by proxy users) that are sending and receiving large amounts of data, and when they are doing so. *Control* allows a user to actually do something about data usage such as by limiting the amount of data a particular device might be able to send. We present three contributions: (1) we introduce a tool for Internet data management in home networks, (2) we present the results of the *uCap* field trial, and (3) we discuss implications for the design of future tools of this nature.

Designing and implementing *uCap* presented many challenges, the most significant of which involved developing a *usable* system on an extremely resource-constrained device. Most home routers have a fraction of the memory and processing power of a “normal” computer and thus cannot support arbitrarily sophisticated monitoring capabilities or elaborate user interfaces. To cope with these constraints, we designed a system that placed only critical

functions on the home router and offloaded more complex functions to a separate software system. This system interfaced with the router and permitted the necessary visibility into data usage and control over Internet traffic, yet offloaded the more expensive operations such as analysis of traffic data and presentation of data in a graphical user interface. This design decision introduced more latency, making it more challenging to offer users real-time feedback and introducing a design tradeoff.

Our field trial of *uCap* suggests users desire more information about how ISPs implement usage-based policies and that the additional visibility that *uCap* provides often gives users a much richer understanding about the applications and devices that consume data on the network. We also found that users generally did not use *uCap* to control the usage of different devices; instead, they relied on social mechanisms to do so. Additionally, in contrast to our expectations about users' desire for privacy of Internet data, we discovered that many wanted more fine-grained information about usage (*e.g.*, the specific application that was consuming traffic, as opposed to just the Internet destination). Overall, our study results imply that users are receptive to Internet data management tools, but the designs can be improved to better fit into the home environment. With Internet pricing mechanisms in flux, our results suggest that developing usable tools that provide consumers visibility and control over their Internet usage is an important area with a rich set of HCI research problems.

BACKGROUND AND RELATED WORK

Various forms of *data caps*, also known as bandwidth or usage caps, exist around the world [25,26]. Many ISPs in US are implementing usage-based pricing [32]; in South Africa, *all* service plans had data caps through the end of 2010 [2]. Caps range from as little as 1–2 GB per billing cycle to several hundred GB. ISPs may enforce caps by degrading (or terminating) a user's connection when they have exceeded the usage cap or by imposing additional fees [18]. Caps that cut users off from the Internet are typically called "hard caps"; other caps are often called "soft caps".

Researchers have demonstrated that users adapt depending on how bandwidth is priced [30,31]. Sen *et al.* performed a study of time-dependent pricing and allowed users to schedule activities depending on the price of bandwidth (*e.g.*, deferring data-intensive activities to off-peak hours). They found that users were often willing to reschedule activities when data prices were cheaper. Instead of experimenting with alternative pricing schemes, *uCap* provides users with better visibility and control over how applications and devices can make efficient use of their *existing* service plans. Rather than requiring users to interact via a single client device, *uCap* runs on the home router and thus can help users monitor and control usage from *all* devices in the home network. Sambasivan *et al.* displayed the cost of accessing a link in search results to mobile Internet users in Ghana, allowing users to choose

whether to click a link based on its cost and their current data balance. This work provided information about pricing for a single mobile device; *uCap* provides information about usage across an entire home network.

uCap applies insights from our previous study of users' experiences with managing data caps in the home [3]. This study highlighted the need for a tool that helps users manage an ISP data cap by tracking usage and identifying the devices, online activities, and users that consume the most data and controlling their usage. Yet, merely demonstrating the need for such a tool does not solve the problem of *creating* a useful one. A usable data cap management tool must grapple with issues of privacy in the home [22]; the flexibility and fine-grained control that the system could provide vs. the risk of distracting users with minutia like bits and bytes [8]; coping with continually changing requirements and usage patterns; and accounting for users of varying technical abilities to both understand the information presented and implement effective usage controls [29]. Such a system must also help users visualize and control "invisible" data consumers, such as automated applications and services that consume data in the background, without the user's knowledge [14].

uCap takes advantage of several trends: (1) the shift towards providing users more control over their home networks, (2) the open-sourcing of software for home routers, and (3) the ability to implement complex network functions in separate network control programs. First, to address various problems users have managing their networks [7,13,15], different systems have addressed problems related to device configuration [34], congestion [4], performance, and access control [5,17,24]. Yet, none of these systems address network management problems related to managing data usage in the face of data caps. Second, the rise of open-source software—specifically OpenWrt [35]—for home routers has made it possible to write software that dramatically improves the monitoring and control capabilities of these (resource constrained) devices. Third, the rise of Software Defined Networking (SDN) [9] and associated software—including OpenFlow-controlled software switches such as Open vSwitch [28] that can run on home routers—makes it possible to implement complex monitoring and control functions and integrate them with elaborate interfaces.

UCAP DESIGN AND IMPLEMENTATION

The conclusions from our prior work [3] provided the initial requirements for a data cap management system. Using those guidelines as a starting point, we performed a user-centered design of *uCap* from June 2011 to June 2012.

System Design

Although home routers provide an ideal location for performing both monitoring and control, they are also underpowered and thus introduce significant constraints. For example, a typical NetGear router might have 16 MB of flash storage, 512 KB of RAM, and a 500 MHz processor,

making it considerably more resource constrained than even a smartphone. There are some trends that suggest movements towards more powerful home routers (e.g., Comcast’s recent acquisition of SkyDog [21]) but these vendors compete on price and have little incentive to create more powerful devices in the short term. Such constraints presented us with some serious choices concerning design tradeoffs: we wanted a responsive, usable, intuitive interface that could operate on a commodity home router.

SDN helped us navigate this tradeoff. Briefly, SDN allows a network device (in this case, the home router) to be controlled by a high-level software program that *runs as a separate control system, typically on a separate device*. Separating the router’s forwarding behavior from its control logic and interfaces allowed us to separate the more complex logic and interface code from the router itself. This design choice did result in more delays between the control systems and the router (and slightly slower responsiveness), but it enabled us to design an interface and control system that would not have been possible on the router itself.

Rather than run a controller in each home, we ran a single controller at a central location that managed the logic and front-end interfaces for all deployed routers. This decision made it easier for us to deploy the system in homes, since deployment only required installing a home router, as opposed to an additional server. We ensured that if a user’s router loses connectivity to the controller, the router would forward traffic normally, to ensure that connectivity problems or controller malfunction do not interfere with the user’s Internet connectivity. More details on the technical implementation are described in Kim *et al.* [16], and all of our code is available on GitHub.

Interface Design

We iterated on the system architecture based on feedback on an initial prototype from network experts at a major conference and an HCI event held at our institution [11,16]. For example, we changed the front-end technologies and back-end communication to improve implementation efficiency. We found that users liked the real-time data usage information and seeing a bar chart of their devices’ data usage against the device’s data cap as well as numeric values of data usage. Potential users also told us that they enjoyed setting and seeing the immediate effects of caps on devices. Once we had a working data tracking and control system, we iterated on interface design ideas from our previous work [3] and this prototype.

We created a series of sketches and higher-fidelity mockups for uCap. Networking and HCI experts performed a heuristic evaluation of these designs in several design-critique sessions. We found that tracking data usage for and setting limits on individual users would require too many changes to the way users currently get online in the home (e.g., requiring all users to use a login mechanism prior to using the Internet). Thus, we shifted our focus to tracking and controlling devices, domain names, and overall

household usage, rather than individual users. We also received feedback on how to improve the overall interaction design, such as placing all the control information in one tab and placing real-time and historical data usage information in separate tabs. We used this feedback to create our final system, *uCap* (a play on “you set the caps”), which allows users to monitor data usage and set mini-caps on home networked devices.




Figure 1. Implementation: Separating control logic and interface from the resource constrained router.

Implementation

We implemented the uCap interface in HTML5 and JavaScript; the router runs custom firmware based on OpenWrt [35]; this firmware runs an OpenFlow-enabled software switch that exchanges both traffic statistics and control messages with an SDN controller that is based on Pyretic [23]. As shown in Figure 1, the router collects the data, aggregates traffic data at a one-minute granularity, and sends reports to a controller that maintains usage statistics and implements controls according to policies that the user sets. A user interacts with the front-end webpage to monitor usage patterns and set usage policies through a webpage that communicates the resulting policies with the controller. If a user sets a usage cap for a specific home device, uCap tracks the cap and prevents a device from accessing the network if it exceeds its cap.

Although the statistics we collect anonymize most destinations for user traffic, uCap helps users track their usage to popular sites by “whitelisting” destinations that are in the Alexa top 100 sites in the United States (excluding pornographic and spam sites) [1]. Users can add or remove sites from this list. uCap only tracks the second-level domain for each site. For example, uCap monitors usage to *.google.com, but not finer-grained information about traffic to “subdomains” (e.g., mail.google.com). We logged all user-initiated actions in the uCap system to help us study how users interact with the system.

To protect user privacy, uCap hashes the second half of each device’s Media Access Control (MAC) address with a cryptographic hash that is based on a unique key on the router. The router maintains a mapping between the original and hashed MAC addresses, but it only communicates with

the controller using the hashed (*i.e.*, anonymized) MAC addresses of the devices; the router never stores the actual device MAC addresses.




Figure 2. Household and device usage for current month.




Figure 3. Monthly usage history timeline per-device.

MONTHLY HOUSEHOLD CAP			
Total Monthly Bandwidth Cap: 250 GB			
The Next Billing Cycle Ends On: 22 Feb 2014			
MONTHLY DEVICE CAPS			
Cap Enabled	Device	Device Usage	Cap Amount
<input type="checkbox"/>	001f16e2366a (N/A)	0 MB	MB
<input type="checkbox"/>	001f5c9a8b6 (N/A)	51.5 MB	MB
<input type="checkbox"/>	0026bb08487e (N/A)	61.5 MB	MB
<input type="checkbox"/>	00405d701857 (N/A)	67 MB	MB
<input type="checkbox"/>	008865031b76 (N/A)	120 MB	MB
<input type="checkbox"/>	00886560d3bd (N/A)	0 MB	MB
<input type="checkbox"/>	20c9d0778a26 (N/A)	0 MB	MB
<input type="checkbox"/>	20c9d0b33667 (N/A)	60 MB	MB
<input type="checkbox"/>	246152140307 (N/A)	1046 MB	MB

Figure 4. Per-device usage and cap settings.

uCap Features

Overall Household Usage: Users first need to understand how the household is consuming data, in near real time. Once users log into *uCap*, they land on a *Network Tab* by default. As seen in Figure 2, in the *Overview* part of this tab, we provide users with a bar graph of the percentage of the total household data cap used to date. The bar reflects cumulative usage from the start of the billing cycle (or the start of the month, if the billing cycle is unspecified) and updated in real time. The data usage resets at the beginning of the billing cycle. This tab also shows a pie chart illustrating each device’s contribution to overall data usage.

A user can set billing cycle information, household cap, and time zone in the *Network Manager* tab. When the household usage reaches the cap, the bar is colored red. We do not enforce the household cap because *uCap* may not be exactly aligned with the ISP’s tracking system, and we did not want to prevent users from using the Internet.

Data Usage of Devices and Online Activities: Once users know the overall trends of data usage, they need to see specific information about data usage of devices and online activities in real time and over time so that they can adjust their usage patterns. Thus, in *uCap*, in the *Network Tab*, a list of devices currently online on the home network is always visible. When a user clicks on a device, he or she is taken to a page for that particular device. This screen shows the amount of data used by the device in real time in megabytes (MB). If a user sets a usage cap for a specific device, the percentage of usage towards the cap is shown as a bar graph. When the device reaches or exceeds the specified cap, this bar is colored red. Users can also choose to set a notification threshold (*e.g.*, 50%, 80%). If device usage reaches this percentage of the cap, *uCap* sends the user a warning. Finally, users can rename their devices (*e.g.*, “Jack’s laptop”) and personalize each device’s avatar.

uCap also provides a historical view of usage in a *History & Status Tab*, as shown in Figure 3, which allows a user to view data usage by devices on the home network over time. The usage is shown as a line graph over the default period of the last week. Users can also select any date range over which to view their data and choose which devices to visualize. The page also shows top ten domains responsible for the most data usage in the selected date range.

Controls to Set Device Specific Data Caps: To effect change over data usage, *uCap* allows users to set mini-caps on devices on their home network and our system enforces these caps. In the *Network Tab->Manager*, shown in Figure 4, the user sees a list of devices on the home network and each device’s current data usage in real time. Users can set caps on the devices and remove the caps for one or multiple devices at a time. Existing caps can also be enabled or disabled. If a device reaches a cap set by the user, it is blocked from accessing the Internet. The users can still access the *uCap* website even if they are capped to allocate more data to devices or disable the cap. On the *History & Status Tab*, users can see a log of all device caps that were set or removed over the selected date range.

Users can personalize their profiles in a *Settings Tab*. In a *Support* tab, *uCap* provides a help manual and allows users to both adjust system settings and send bug reports.

UCAP STUDY

Recruiting Process: We created a promotional website and Facebook page for *uCap* and a Google Forms signup page. We then publicized these sites and actively recruited participants via Facebook, Twitter, mailing lists, broadband forums, and our contact networks. We attempted to recruit

#	Household Composition	Annual Household Income (USD)	# Internet Enabled Devices	Cost	Cap (GB)	Location	Approx. Length of Participation in Study
H1	Male [UX research] (25-34)	\$35K-\$50K	>10	~\$62	250	US	15 months
H2	Male [CTO of ISP] (35-44) & partner	Undisclosed	>10	~\$48	300	US	19 months
H3	Female [Research Scientist] (35-44) & two children	\$50K-\$75K	9	~\$40	40	US	16 months
H4	Male [Grad student] (25-34) & partner	\$50K-\$75K	>10	~\$47	U	US	19 months
H5	Male [Grad student] (25-34) & 3 roommates	\$100K-125K	9	~\$90	300	US	19 months
H6	Male [Research manager] (35-44) & wife	\$100K-125K	>10	~\$32	75	IN	3 months
H7	Male [IT engineer] (25-34) & wife	Undisclosed	7	~\$110	U	ZA	11 months
H8	Male [Technician] (18-24), mother & brother	<\$8K	7	~\$49	U	ZA	4 months
H9	Male [IT support] (18-24), mother, aunt, & sister	<\$8K	6	~\$66	U	ZA	11 months
H10	Male [IT manager] (25-34) and wife	>\$25K	7	~\$50	130	ZA	4 months

Table 1. Participant Demographics (occupation, age range and household members, Cost is self-reported and may include Internet and phone costs, line rental, and bundles per month, U is uncapped, US is USA, IN is India, ZA is South Africa).

more technical users who could easily install the *uCap*-enabled router and who were already interested in their home networks, regardless of whether they had a data cap. By focusing first on this demographic, we could work on the design of the system itself, with our ultimate goal being to extend this work to the general population. We also looked for participants in multiple geographic locations to determine whether different Internet environments affect *uCap* use. Although this approach yielded fewer users in any single country, it provided more diversity in the types of environments we could study.

We recruited 21 households: 13 in United States, 6 in South Africa, and one in India. Ten of these households also participated in an in-depth qualitative study. These participants used *uCap* between June 2012 and February 2014, some over the entire period and some for parts of the period. After we gained informed consent from participants' entire households, we mailed them a router with instructions on how to swap it for their existing router (or, in South Africa, to plug it in behind their existing ADSL router).¹ When participants received and installed their routers, we created *uCap* accounts for them and showed them how to use the system. All participants were compensated with the router.

Privacy Challenges: *uCap*'s collection of personal information introduces challenges related to privacy. To mitigate risk, all of our data collection and consent processes were reviewed by our institution's review board. The lengthy consent process and user anxiety about data collection made it difficult to recruit a large sample of

users. We expect future studies involving collection of user data in home networks to face similar challenges.

Method: All ten households that participated in the qualitative study completed a demographic survey and a pre-study survey about their Internet usage habits, their service plan, and their experience with data caps. We interviewed participants both before and after they used *uCap* and had them complete a post-study survey about their experiences. The final interviews focused on how participants used *uCap* and their experiences with monitoring data usage. All interviews were conducted over Skype and audio-recorded. We also logged all user-initiated actions on the *uCap* website and collected per-device usage information, as well as data usage of whitelisted websites.

We transcribed all pre- and post- study interviews and used a thematic inductive analysis to find the themes across the data [6]. The primary author performed the initial analysis, and the themes were discussed with the research team to arrive at a consensus. We also performed a triangulated analysis, comparing measured network usage and interaction with *uCap* to the users' reports of their behavior.

Participants: Table 1 summarizes the participant demographics. We had a mix of household types (couples, one single person, and families), but the primary person we interacted with and interviewed was typically a (male) head of the household who had a technical background. Four households had an unlimited service plan, and six households had a bandwidth cap.

FINDINGS

We discuss users' desires for a data cap management tool prior to using *uCap*, the results of the field trial, and users' privacy concerns about Internet data management tools. We did not observe any significant differences between countries; capped and uncapped users had different concerns, as we explain below.

¹ Because the South African users did not necessarily use the router as their main router, not all of their home devices may have connected to the network via this router. As such, *uCap*'s measurements of their data usage may not have been as accurate as for the US users.

Before uCap: Desires for a Data Cap Management Tool

Our participants all reported using the Internet regularly in the interviews and surveys for email, work activities, and entertainment including either gaming, file sharing, and downloading or streaming movies and music. Echoing some findings from our earlier study [6], they all struggled to understand their data usage over the course of the month, what activities and which devices were using up data on their networks and desired more *control* over data usage.

Generally, participants signed up to use *uCap* to gain more insight into their data usage, beyond what was available to them already via their ISP or existing tools. Some used *uCap* to learn more about data usage of various devices and applications. Users sometimes wanted to ensure that devices or websites were not consuming excessive data. In other cases, users wanted to know more about the relative costs of comparable services. For instance, one household (H4) mentioned they wished to compare usage of two streaming services, HBO versus Netflix, to determine which was worth the subscription fee.

Users with “hard caps” wanted to use *uCap* to see why they were exceeding their quota. Households that had used previous tracking tools (7 of the 10 in the study) wanted to see how *uCap* compared to these tools, particularly since setting up tracking often involved significant effort or time investment. For instance, H8 developed his own monitoring tool to show him graphs of data usage in his home. Others experimented with existing tools such as NetWorx (H9) but found that these tools were only able to track one PC not other devices such as an Xbox [36] or experienced issues installing the tools on machines used for work (H6). Only one household (H5) had previously used a router-based tool to monitor data usage. Users noted that existing tools were difficult to set up, did not provide adequate visualization, or did not present per-device usage.

To sum up, all participants expressed a desire to track their data usage and most had tried to use existing tools but found these tools were lacking in terms of either setup or the information and control that the tool offered.

Lack of Transparency into ISP Policies and Accounting

Our participants also complained about the lack of transparency that ISPs provide into data usage. For instance, H8 told us that his ISP profiles “heavy” users to determine how to shape users connections including those on unlimited plans. However, this profiling was not transparent, making it difficult for the user to trust the ISP: *“They also implemented a threshold for a week period: They used to have a threshold, every seven days which was for my account 20 Gigs in seven days. If you pass that you get warning. Now they work on a 5 star system. Depending on network load and your usage, it figures out your star [rating] as well as what shaping moves to apply to your account.”*

Five households (H2, H5, H6, H8, and H9) discussed using their ISPs’ websites to check aggregate use but also commented on how these tools did not satisfy their needs.

For example, H2 lamented: *“I don’t feel like I have a lot of control in that situation. Because all they just say is a net number and I don’t know where it’s coming from and if I go over the number and I am charged, I don’t really have a lot of control on that.”* Participants often did not trust the accuracy of ISP tools because they did not give details at a fine level of granularity; for example, they lacked a daily usage breakdown. In one case, H5 reached the usage quota and checked the router’s logs to better understand the usage patterns: *“There was one day a few months ago where we ended up blowing through the bandwidth cap and when we checked on the router, the router said that we had consumed something like 80 gigs in a part 24 hour period, which, given our speeds, I don’t think was actually possible.”*

Others questioned whether ISPs have an incentive to notify them about whether they are going over the cap because many ISPs collect “overage charges” when the usage cap is exceeded. Some simply did not understand how or when ISPs would enforce the caps. An example of this concern is expressed by H1 who talked specifically about running way over the cap and not feeling any consequences: *“They haven’t actually been enforcing it. So I have been definitely passing over a terabyte a month sometimes and yet I’ve not received a few notifications.”*

Some participants worried about whether they had subscribed to the right data plan and whether a data cap would affect them. For instance, H4 wanted to make sure that the household would be within the cap: *“I really had no way of knowing how much I should be signing up for and what plan I should be signing up.”* In some cases, participants wanted to know if they could switch to a cheaper plan or decrease costs by sharing the connection with others: *“I live in a fairly small apartment building and I have a neighbor upstairs and one of the things I’ve been wondering was whether or not we could share Internet.”* (H3)

Thus, before using *uCap*, participants expressed a strong desire for transparency into how ISPs track and enforce the cap, a distinct distrust in ISPs in general, and doubt about how to pick the service plan that best suits their needs.

uCap Usage in the Field Trial

Table 2 shows the number of times each participant accessed *uCap* during the study, according to our logs. This usage is slightly lower than the usage reported on the surveys but from both data sources, it appears that users accessed the web interface infrequently. Yet, all but one of the households (H3) told us that they found the information on *uCap* to be helpful and that they actively used the system when they wanted more information on data usage (much like one checks a bank statement).

In particular, self-identified “heavy” users who faced the possibility of exceeding their usage caps used *uCap* regularly. For instance, when their ISP notified them that they would enforce the cap in December 2013, H5 turned to *uCap*: *“Since then we’ve been monitoring our usage on uCap pretty closely and obviously uCap had months of data prior to December of last year and so we were sort of comparing the*

numbers over the summer 2013 of Comcast versus uCap”. The exception, H3, felt that the system was not helpful because her household was not hitting its cap. In another case, an OpenWrt bug causing instability in Wi-Fi connections led to one participant (H6) to stop using uCap altogether, even though this household found uCap helpful.

	# Logs	Estimated Use	Actual Use (uCap)
H1	112	1 TB	916 GB
H2	21	40-50 GB	72 GB
H3	9	-	-
H4	39	80-100 GB	105 GB
H5	19	400-450 GB	529 GB
H6	7	50-60 GB	66 GB
H7	4	-	-
H8	9	-	-
H9	11	-	-
H10	5	-	-

Table 2. uCap Usage Logs and Estimated vs. Actual Monthly Data Usage Per Household.

All participants said that they found the system was easy to install, since enabling uCap merely entailed replacing the home router and did not require other configuration changes to the network or devices. H1’s quote summarizes the collective positive feedback on why our participants found it helpful: “I don’t think that anything else has provided that level of transparency to home networks.”

Another quote suggests the type of use envisioned for uCap beyond the study period by H4: “I anticipate that I’m going to keep using it. Kind of more infrequently than I used to. Just checking in maybe once every few weeks.” He emphasized: “So it’s the kind of tool I really liked having. I can see it becoming an invaluable resource for me.”

Many participants asked for a version of uCap for mobile devices that provided more options for push notifications about data usage. Overall, uCap was well received.

Tracking Overall Household Usage and Learning Patterns

Most participants told us that they found tracking the overall household data usage useful for learning patterns of use. For example, in a quote typical of those we heard, H1 explains: “So I was actually kind of surprised with how much data I was using. I posted it on some of my social media accounts. I was like check it out guys I’m using over a terabyte a month.”

For the half of our households that used uCap as their primary router—for which we also had a self-reported estimate of monthly data usage—we calculated the actual monthly data usage (Table 2). Interestingly, all of these five households gave fairly accurate estimates of monthly data usage. For households H2, H5, and H6, we suspect that this is because they had used previous tools to track their usage.

Others may simply be well-informed, or it may be that uCap gave them more awareness about their usage.

Participants were sometimes confused because the real-time data and the historical data were not always consistent. Two participants (H2 and H5) were worried that the usage reflected did not match up with the reported usage from their ISP. This inconsistency resulted because our real-time and historical data are tracked by two separate components, creating slightly different data counts. Furthermore, uCap itself consumes data, but we opted not to show this data usage because it represented a minimal 0.4% of total usage on average. For these reasons, the data usage reported by uCap differed slightly from the user’s actual usage.

Participants also felt that the real-time usage units (GB for the household and MB for individual devices) were too coarse-grained. For example, H1 talked about how incremental changes were hard to notice and correlate with the real-time usage, since many actions such as checking email do not consume much data and were thus difficult to see on the progress bar. “Most applications don’t use the multiple gigabytes or megabytes when you’re watching it. So it’s hard to notice the incremental changes that smaller actions on the network do, checking emails and things like that.”

Participants liked the pie chart showing the device breakdown in the Network Overview Tab but found that it became cluttered after awhile, as it continued to accumulate many unidentified devices. We did not provide a way to delete or remove transient devices (e.g., those belonging to guests), but doing so would improve the user experience. For example, Grover et al. [14] revealed that one dominant device typically consumes the bulk of traffic in a home (60% on average), thus showing the top 5 devices may provide sufficient information to identify data “hogs”.

Most participants did not mention frequently using the Network Tab->Devices page for individual devices. A few had tried to set notification emails but had trouble with receiving emails. Others did not realize that it was possible to set notifications. In summary, participants found household-wide tracking helpful but noticed discrepancies between the real-time and historical information. They also suggested ways to de-clutter the interface.

Tracking Device Data Usage and Starting Conversations

Participants consistently told us that they found the history view most useful for tracking individual devices’ data usage. For example, H6 talked about comparing different devices’ usage over time. H5 talked about how he liked uCap’s breakdown by MAC address and how he would use this information as a conversation starter with other household members: “Because uCap breaks down the usage by MAC address if there was a huge spike in bandwidth, I would ask the roommate associated with that MAC address if they had done anything. And the reason for that was not to restrict them from using the Internet right. Because it becomes pointless if basically you prohibit yourself from using the Internet in the normal way.

But the question was more to make sure that they were the ones responsible for it as opposed [to processes they are unaware of]”.

Similarly, H2 described how he discovered that their overall data usage was not very high and that the data “hogs” he suspected were not using as much as they thought: *“I actually discovered my Apple TV used less data than it actually used. I was expecting it to be more chatty than it actually was. And watching movies didn’t kill me as bad as I thought.”* In fact, after using *uCap*, this participant said his usage pattern changed because he was *less* worried about exceeding his usage cap. Now, he said, *“When a provider tells me that you know [you have] 300 gigs or whatever I guess. And then I do a simple calculation, then I get worried about using it. Having information saying that it really wasn’t as impactful as I thought it was, it actually made me use the service more.”*

In many cases, *uCap* confirmed what users perceived about usage of specific devices before they used the tool. Many of our participants also mentioned that they discovered new and interesting things about their home network data usage patterns. A participant in H4 explains: *“Every once in a while there would be weird spikes where a particular device would use twice as much data as it had ever used before. And so there was always a little mystery with that. Usually it was just things like the device downloading updates, things like that.”*

Others said that *uCap*’s list of online devices and their usage history gave them a sense of security about whether the data cap was being abused by unauthorized users. In H7, *uCap* alerted the household to a virus: *“I was actually able to find out that my wife’s laptop was infected. Actually, yes because the amount of data it was using was far more than her looking at YouTube”*.

This same participant and others expressed surprise at how much data an iPad and an Android phone use. *“When you actually see how much data that device uses just when you’re at home sleeping and the phone is charging. It’s quite scary”*. Participants also commented on how they enjoyed seeing what devices such as the Wii or Xbox used and how easy *uCap* made it to see the usage: *“It makes it a bit easier for normal home users to actually monitor all their various devices that are actually using bandwidth. So you know things like smart phones and stuff like that where you can’t specifically monitor that kind of traffic without having some kind of a proxy”*. (H9)

In summary, participants found *uCap* useful for illuminating per-device usage (particularly background traffic) and the tool reassured them that trusted devices were on the network using data. Participants also suggested that showing usage by time-of-day would be helpful.

Tracking Data Usage for Online Activities

Participants had mixed feelings about the panel showing the top-ten domains using data in a month. Some users liked the information that it provided but also noticed that the “other” category reflected a lot of usage. (*uCap* categorized traffic to any destination not explicitly whitelisted as “other”, since any destinations outside of the Alexa top 100 were anonymized and, hence, could not be categorized.)

Others felt the top-ten domains did not tell the users much about their data usage. For example, H4 noticed NetFlix did not show up as the biggest data hog: *“It didn’t really give me as much information as I was hoping for. Like whenever I would use NetFlix, it didn’t show up as coming from NetFlix, cause the data is sort of by Amazon usually, but so was a lot of other stuff on the Internet. So basically we would just end up with Amazon, 2 Gigabytes”*. In fact, since NetFlix is actually hosted on a content distribution network (CDN), such as Amazon Web Services, the domain reflecting NetFlix usage would show up as *amazonaws.com* or *akamai.com*, not as *netflix.com*. (In fact, *amazonaws.com* consistently shows up in the top five most popular domains according to Grover *et al.* [14].)

Most of all, participants wanted more detail on top website and application usage, contrary to what we would have expected given privacy concerns. For example, H2 wanted a per-protocol breakdown of traffic. Others wanted usage in terms of activities such as “watching a YouTube video”.

Participants liked having the ability to modify the whitelist, but some were suspicious that what was displayed was not the same as what was tracked. To sum up, participants wanted more detailed information on data usage, even at the expense of privacy. Additionally, they wanted data usage to reflect the everyday online activities they engage in to make the information easier to interpret.

Controlling Data Usage By Mediating Behavior

Interestingly, none of our participants regularly used the features to set caps on their devices. Only four of the ten (H1, H6, H7, H8) tried the capping functions at all; those four felt that imposing a “hard cap” on a device was not useful or necessary for controlling use within the home. Some suggested that capped devices be throttled instead or redirected to a page explaining what happened. Participants also told us they often used the per-device data to decide whether to start a dialogue about data usage with other household members.

In a typical example given by participants, when Apple TV was discovered not to be a data hog as suspected in H2, the participant consciously did not mention this to his partner, the primary Apple TV user: *“He watched a ton of movies on Apple TV and I thought that that would put us in jeopardy and when the number showed that it didn’t, I never said anything to him about it.”* In summary, rather than relying on the system for control, users relied on conversations to manage data usage, with *uCap* providing a starting point to chat.

Privacy Concerns With Tracking Home Data Usage

In designing *uCap* to protect privacy, we inadvertently compromised usability. For example, obscuring device MAC addresses made it more difficult to identify and rename devices, and withholding certain data about per-application usage also obscured useful information.

Most participants said they did not want ISPs tracking the same data *uCap* collects; a few stated they only trusted the research team with the data. Some felt that the ISP should

not have access to the type and number of devices a user has. Others felt the ISPs could monetize the information or charge them for using certain types of traffic. A few worried that ISPs would track illegal downloads or sell the data. Most mentioned that cloud storage presented a security threat from outsiders. For example, thieves might target households with many devices.

In particular, storing uCap data outside the home at a central location created anxiety: “*Having all your data logged on a sever, your whole usage patterns, and everything it’s a bit nerve-wracking.*” (H8). H6 reflects these concerns: “*Who has access to it? And what is the control, my control over that data? And what are they using the data for?*” Users also worried about how an ISP could use the data to control usage, such as blocking traffic. Our participants had conflicting ideas about what they wanted uCap to show, notions of what data should be private, and who should have access to this data. These attitudes are striking because ISPs can *already* gather much of the data that was cause for concern. This naïveté highlights a gap between users’ perceptions and discomfort about what data might be collected about them and ISPs’ *current* capabilities for gathering this very data.

DISCUSSION

Our field trial showed that uCap helps users understand more about their households’ data use, as well as what devices (and often by proxy, people) and activities consume network data. This visibility helped participants react to in-the-moment situations (e.g., “Do I need to curb usage now?”) and see trends over time. Our participants also suggested minor improvements, such as allowing users to remove devices no longer present on the network or providing more visibility into how ISPs charge for data to enable better auditing. To our surprise, participants also wanted more information about the data usage of online activities that was in direct conflict with their anxieties about keeping this data private from others and ISPs.

As the Internet of Things and a greater reliance on the cloud introduces more background traffic, and as ISPs continue to implement data caps and other pricing mechanisms [19], developing tools that help users manage their Internet data will only become more important. Given that the home network is a complex socio-technical system, these solutions must not just provide technical capabilities; they must also mesh with how people operate in home environments. For instance, our users did use *uCap* for controlling data usage, but not with the mini-caps as we expected; instead, users either self-mediated (*i.e.*, reduced usage) or mediated others’ behavior with social means.

Implications for Design of Data Cap Management Tools

Placing the *uCap* controller in a central location made *uCap* easier to deploy, but users remained concerned about privacy and the granularity of the data that the tool collects. Future Internet data cap management tools should store the tracked data in the home to minimize privacy and security risks. Future tools can improve tracking online activities by

providing more detail on time-of-day usage and intuitive summaries about usage of specific applications (e.g., Netflix). Doing so requires better activity characterization of web traffic in real time, which should be easier if the controller is situated in the home. It also requires better application identification mechanisms, to allow flows to be associated with “watching YouTube”.

Finally, since our study focused on technically savvy users, the interface and visualizations may need to be adjusted for a more general audience in future studies. For instance, *uCap* could include cost information, as in SmartBrowse [26], or compare usage to activities that use the same amount of data (e.g., “you used 1 GB of data, which is roughly one episode of a TV show”). The tool could also depict usage averages and reference points such as past consumption patterns. These mechanisms can help make the tool more accessible to a more general audience. Ultimately, the lessons from the *uCap* study may apply to other settings beyond the home where there are also no dedicated system administrators such as libraries and small businesses.

CONCLUSION

We presented the creation and evaluation of *uCap*, a home Internet data cap management tool. In our field trial, we found users want better ways to manage data and our participants were positive about using *uCap* to help them do so. Our findings suggest that creating tools that fit into users’ everyday Internet usage patterns is difficult without extensive field trials of this nature. Future research could improve the existing interfaces we designed, address the challenges of providing detailed data usage information without compromising user privacy, and extend our work to a more general user population. Our study also emphasizes the increasing value of creating tools to provide consumers with visibility and control over Internet data usage.

ACKNOWLEDGMENTS

This work was supported by NSF Awards CNS-1059350, CNS-1346723, CNS-1405781, CNS-1422680, and a Google Focused Research Award. We thank our participants and other contributors: Victoria Ayo, Michael Dandy, Zuiena Kabir, Rebecca Rouse, Boris de Sousa, Hyewon Suh, Bethany Sumner, and Steve Woodrow.

REFERENCES

1. Alexa. Top websites in USA. <http://www.alexametric.com/topsites/countries/US>.
2. Chetty, M., Banks, R., Bernheim Brush, A.J., Donner, J., and Grinter, R.E. Under Development: While the meter is running: computing in a capped world. *Interactions* 18, 2 (2011), 72–75.
3. Chetty, M., Banks, R., Brush, A.J., Donner, J., and Grinter, R. You’re capped: understanding the effects of bandwidth caps on broadband use in the home. CHI 2012, ACM (2012), 3021–3030.

4. Chetty, M., Banks, R., Harper, R., Regan, T., Sellen, A., Gkantsidis, C., Karagiannis, T., and Key, P. Who's hogging the bandwidth: the consequences of revealing the invisible in the home. *CHI 2010*, ACM (2010), 659-668.
5. Chetty, M., Haslem, D., Baird, A., Ofoha, U., Sumner, B., and Grinter, R. Why is my internet slow?: making network speeds visible. *CHI 2011*, ACM (2011), 1889-1898.
6. Corbin, J. and Strauss, A. *Basics of Qualitative Research*. SAGE, 2008.
7. Crabtree, A., Mortier, R., Rodden, T., and Tolmie, P. Unremarkable networking: the home network as a part of everyday life. *DIS 2012*, ACM (2012), 554-563 .
8. Davidoff, S., Lee, M.K., Yiu, C., and Zimmerman, J. Principles of smart home control. *UbiComp 2006*, Springer-Verlag (2006), 19-34 .
9. Feamster, N., Rexford, J., and Zegura, E. The Road to SDN. *Queue 11*, 12 (2013), 20-40.
10. Federal Communications Commission. Policy Issues in Data Caps and Usage-Based Pricing. (2013), 1-26.
11. Georgia Institute of Technology. GVU Fall Research Showcase. 2011.
12. Government Accountability Office. Eshoo Shares Preliminary Findings of Study on Broadband Data Caps with FCC 2014. <http://eshoo.house.gov/press-releases/eshoo-shares-preliminary-findings-of-study-on-broadband-data-caps-with-fcc/>.
13. Grinter, R.E., Edwards, W.K., Chetty, M., et al. The ins and outs of home networking: The case for useful and usable domestic networking. *ACM TOCHI* 16, 2 (2009), 8-28.
14. Grover, S., Park, M.S., and Sundaresan, S. Peeking behind the NAT: an empirical study of home networks. *IMC 2013*, ACM, 377-390.
15. Kawsar, F. and Brush, A.J.B. Home computing unplugged: why, where and when people use different connected devices at home. *UbiComp 2013*, ACM (2013), 627-636.
16. Kim, H., Sundaresan, S., Chetty, M., Feamster, N., and Edwards, W.K. Communicating with caps: managing usage caps in home networks. *SIGCOMM 2011 Demo*, ACM (2011), 470-471.
17. Yiakoumis, Y., Katti, S., Huang, T., McKeown, N., Yap, K., and Johari, R. Putting Home Users in Charge of their Network. *UbiComp 2012*, ACM (2012), 1114-1119.
18. Lasar, M. *It could be worse: data caps around the world*. Ars Technica, 2011.
19. Lohr, S. *F.C.C. Is Deluged With Comments on Net Neutrality Rules*. *New York Times*, 2014. <http://www.nytimes.com/2014/07/16/technology/a-deluge-of-comment-on-net-rules.html>.
20. Lucey, P. How Data Caps Are Bad for Cybersecurity. http://oti.newamerica.net/blogposts/2014/how_data_caps_are_bad_for_cybersecurity-117644.
21. Lunden, I. Comcast Acquired PowerCloud Systems, Gears Up For Smart Home Play. *TechCrunch*, 2014. <http://techcrunch.com/2014/07/14/comcast-skydog/>.
22. Mazurek, M.L., Arsenault, J.P., Bresee, J., et al. Access Control for Home Data Sharing: Attitudes, Needs and Practices. *CHI 2010*, ACM (2010), 645-654.
23. Monsanto, C., Foster, N., Rexford, J., and Walker, D. Composing Software Defined Networks. *NDSI 2013*, ACM (2013), 1-13.
24. Mortier, R., Rodden, T., Tolmie, P., and Lodge, T. Homework: Putting interaction into the infrastructure. *UIST 2012*, ACM (2012), 197-206.
25. OECD. *Average data caps by country (GB) logarithmic scale, September 2012*. 2012.
26. OECD. *Prevalence of explicit bit/data caps amongst surveyed offers by country, September 2012*. 2012.
27. OToole, J. Comcast plans data caps for all customers. *CNN*, 2014. <http://money.cnn.com/2014/05/15/technology/comcast-data-limits/>.
28. Pfaff, B., Pettit, J., Amidon, K., Casado, M., Koponen, T., and Shenker, S. Extending Networking into the Virtualization Layer. *Hotnets 2009*, (2009).
29. Poole, E.S., Chetty, M., Grinter, R.E., and Edwards, W.K. More than meets the eye: transforming the user experience of home network management. *DIS 2008*, ACM (2008), 455-464.
30. Sambasivan, N., Lee, P., Hecht, G., and Aoki, P.M. Chale, how much it cost to browse?: results from a mobile data price transparency trial in Ghana. *ICTD 2013*, ACM (2013), 13-23.
31. Sen, S., Joe-Wong, C., Ha, S., and Bawa, J. When the price is right: enabling time-dependent pricing of broadband data. *CHI 2013*, ACM (2013), 2477-2486 .
32. Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. A survey of smart data pricing: Past proposals, current plans, and future trends. *ACM Computing Surveys (CSUR)*, (2013).
33. Valancius, V., Lumezanu, C., Feamster, N., Johari, R., and Vazirani, V.V. How many tiers?: pricing in the internet transit market. *SIGCOMM 2011*, 194-205.
34. Yang, J., Edwards, W.K., and Haslem, D. Eden: supporting home network management through interactive visual tools. *UIST 2010*, ACM, 109-118.
35. OpenWRT. <https://openwrt.org/>.
36. NetWorx. <http://www.softperfect.com/products/>