# Unified Methodology for Broadband Behavior Measurements in the Acreo National Testbed

Christina Lagerstedt, Andreas Aurelius, Hemamali Pathirana, Claus Popp Larsen
Netlab, Acreo AB
Kista, Sweden
christina.lagerstedt@acreo.se
andreas.aurelius@acreo.se
hemamali.pathirana@acreo.se
claus.popp.larsen@acreo.se

Olle Findahl
World Internet Institute
Gävle, Sweden
olle.findahl@wii.se

*Abstract*— **The continuous evolution of Internet and Internet applications increases the demands on access networks. Understanding user behavior and Internet usage patterns is fundamental in developing future access networks that meet technical as well as end user needs, and from a societal point of view it is equally important to correctly recognize and understand user behavior. In this paper, we present measurements from the Acreo National Testbed where we have access to traffic measurements from real end users. We have developed a unified methodology combining traffic measurements with web questionnaires and diaries to compare the results from different methods as well as investigate user behavior. By comparing the different measurement methods we find that the end users have difficulty in estimating the time they spend on different Internet activities although they are fairly well able to estimate the frequency of usage. We also found that though the diaries are quite accurate, the traffic measurements give us a much more detailed picture of the end user activity. The importance of having a testbed with real end users is invaluable to this kind of study and we emphasize the importance of having access to access network traffic.**

***Keywords - Testbed; Traffic measurements; User behavior; FTTH.***

## I. INTRODUCTION

As the Internet continues to evolve and offer new services, it takes up a larger part of our lives. We find new ways to communicate, interact and entertain ourselves. This puts new demands on access networks [1] and requires new insights into the behavior of those who use them. We believe that understanding user behavior and needs is the key to develop future networks and services that are accessible, reliable and that address the needs of real end users. There are several ways to study user behavior. From the technical side, the data traffic can be measured and analyzed. Other common ways are to use surveys or diaries. Traffic measurements are routinely performed by all larger operators, but results are rarely published because the operators don't want to share this information with competitors. On the other hand, published behavior studies are almost always based on surveys with individuals (e.g., telephone interviews with a large population). Such surveys often attract considerable interest among in the public debate, and far-reaching conclusions may be drawn.

The question is now whether such surveys are reliable. People can forget about their Internet behavior, they may not know what they did, they may not know what the children in the household have done, they may lie about sensible subjects, etc.

The purpose of this study is to develop a unified methodology where different kinds of surveys are combined with actual traffic measurements. We compare three different methods of looking at Internet user behavior; 1) Internet protocol (IP) traffic measurements, 2) web questionnaire and 3) diary. This will lead to a more detailed knowledge of the behavior, and, by evaluating results from different methods, we will obtain a better knowledge of their respective limitations. Furthermore, this way it can be verified whether surveys are correct or whether there perhaps is a systematic bias in survey answers that thus leads to misleading results.

There are few comprehensive traffic measurement studies in the literature and most of these are based on traffic from a campus areas [2], [3] or measure aggregated traffic [4], [5]. Questionnaires and diaries are well known and often used methods [6], [7], [8], but, as far as we know, have not been used together with traffic measurements. In the following sections, we will first describe the measurement setup and population. We will then present and discuss our results and finally presents our conclusions.

Combining technical measurements with surveys require test subjects in order to get statistical data. At Acreo we put a lot of effort into developing the Acreo National Testbed, ANT, which enables us to perform in-depth measurements and test new technology and equipment as well as to interact with end users. Here we can study user behavior in a unique way as we have access to real user traffic measured on a household level.

We have signed agreements with the end users where they agree to give feedback and participate in surveys and investigations. This means that the network conditions such as network topology, link speeds, service setup, etc as well as user metadata such as the number of people in a

 ISBN: 978-1-61208-193-9 

household, age, etc are known. This gives us a unique opportunity to perform measurements of user behavior, and to compare the results from different measurement techniques in order to evaluate the validity of the results.

## II. TEST ENVIRONMENT AND POPULATION

As mentioned in the introduction, the measurements in this study were performed in the Acreo National Testbed (ANT), which has previously been described in [9] and [10]. Contrary to lab based testbeds, this is a live network with real end users or *test pilots*. In return to being test pilots the end user households are given free access to services like Internet and IPTV. Fiber to the home (FTTH) is the main access technology in the testbed, and a schematic picture of the network is shown in Fig. 1.

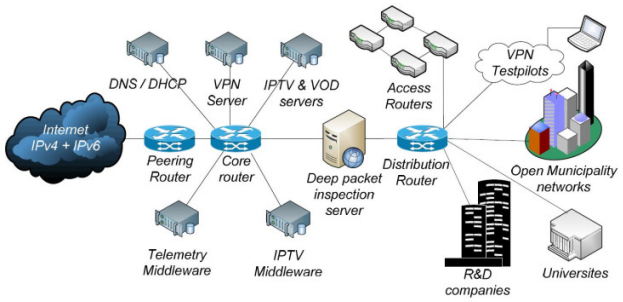


Figure 1. Schematic picture of the Acreo National Testbed, ANT.

The FTTH installation at the test households is active Ethernet providing 100 Mbit/s symmetrical connections to each household.

The number of households in the testbed changes over time according to the current tests that are being performed. At the time of this study, there were approximately 40 active households in the testbed. Approximately 20 were apartments in a building centrally located in the Swedish town Hudiksvall. The rest of the households, approximately 20 single dwelling units, were connected to the testbed via a fixed wireless access network, depicted as HSPA (High Speed Packet Access) in Figure 1.

The measurements in this paper are based on data from 5 testbed households. The reason for using a small population in the study is that we wanted to perform a *qualitative* analysis of their Internet usage, i.e., the main point is not to gather data for statistical analysis, but rather to study the user behavior *in depth*, as we have not found any similar studies in the literature. As this is a novel comparative methodology it was also important to develop methods and analysis tools that can be scaled up to considerably larger populations [13] where the amount of data to be handled will be much more extensive. That is, we do not claim to be able to make statistical conclusions in this study, but interesting trends for further study will be pointed out. Nevertheless, we will develop a unified methodology that, at a later stage, can be used for more substantial conclusions.

A letter of invitation was sent out to all households in the testbed, asking for participation in the study. The participants were selected from those who responded. In the selection process, household details such as number of people in the household, number of computers, ages, etc were taken into account, in order to get a varied test population with different household constellations. The test pilots were assured that their answers would be treated anonymously. This resulted in a test group of 5 households of which 2 were single households and 3 family households. Of the family households, one had small children (<7 years old), one had teenage children and one both teenagers and small children.

## III. MEASUREMENTS

The combined measurement methodology consists of three parts: traffic measurements, web questionnaires and diaries. In this section we will describe these methods.

### A. *Traffic Measurements*

The traffic measurements have been performed using PacketLogic (PL) [12], a commercial traffic management device used in many commercial broadband access networks all over the world. Traffic is identified based on packet content (deep packet inspection and deep flow inspection) instead of port definitions. The device can identify more than 1000 Internet application protocols, and the signature database is continuously updated.

Since the PL is a commercial product, the details of its functions are proprietary. However, the identification process is connection-oriented, which means that each established connection between two hosts is matched to a certain application protocol. When a new connection is established, the identification of this connection begins. The identification algorithm searches for specific patterns, signatures, in the connection. The patterns are found in the IP header and application payload. The PL uses the traffic in both directions in the identification process. The measurement point is depicted as “deep packet inspection” in Fig. 1.

The PL can track and identify several hundred thousand simultaneous connections, storing statistics in large databases. The statistics database records the short-time average amount of traffic in inbound and outbound directions as well as the total traffic for all nodes in the network. The data is averaged over 5 minute periods. Data concerning which web sites have been visited is stored in the connection log.

The measurement setup, although giving detailed measurements, has certain constraints. First, the traffic is measured per household and not per person and the analysis in this paper is therefore done on a per household basis.

 ISBN: 978-1-61208-193-9 

There is also a 5 minute resolution in the measurements, which may have an impact on measurements of applications that are used in short time periods such as instant messaging. The data cut-off is 1 kbps, which may influence the measurements of certain applications such as gaming where the amount of data is generally very low. The signature database in this study was not up-to-date due to old hardware, which may result in a larger amount of unknown traffic. An upgrade will be performed before follow-up studies are performed.

The Internet traffic of each household was measured both during the days when the household recorded their diaries and for a complete month (May 2009) to get enough statistics to compare with the web questionnaire. Statistics on what web sites were visited by the different households were monitored for two weeks.

### B. *Web questionnaires*

Each household member was asked to answer a web questionnaire regarding their Internet activity and behavior. In the case of small children (<7 years), the parents were asked to answer for them. The web questionnaire contained basic questions concerning family situation, education and occupation as well as questions pertaining to computer knowledge and Internet activity. The Internet activity questions include questions on what kind of medium is used to access Internet applications and frequency of use of different applications.

The questionnaire also surveys how often different types of web sites are visited such as banks or newspapers, etc. Finally the respondents are asked to estimate the amount of time they spend on different Internet activities. The questions are multiple choice with an additional field for comments.

### C. *Diaries*

The members of each household were asked to fill in a diary logging their Internet activity during two consecutive days, 17-18 May 2009. Each day was divided into 15 minute intervals. The diary had four columns that the test pilots were asked to fill in:

1. Daily activity (sleep, work/school, leisure time activities, meal times etc.)
2. Media usage (TV, newspaper , radio, book, etc)
3. Internet activity when at home (web browsing, playing games online, community, downloading material from the Internet, etc)
4. Web address or service used

## IV. RESULTS

From the traffic measurements, we find that the Internet activity of the households in the study occurs mainly during afternoons and evenings with shorter bursts of traffic during the morning and lunch hours. This is consistent with the traffic patterns established both in ANT and in municipal networks of similar characteristics, but with much higher populations [10]. The daily traffic pattern in the testbed during May 2009 is shown in Figure 2.

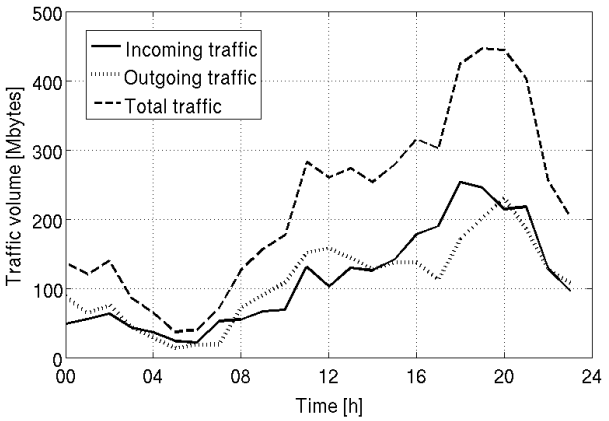


Figure 2. Daily traffic pattern of all active households in ANT during May 2009.

We also note that the average time spent online calculated from the traffic measurements is greater during the weekends than the weekdays for the family households while the opposite is true for the households without children, see Table I.TABLE I. The assumption here is that for the family households, the time of day when household members use Internet applications will be more spread out during the weekends and of course there are more people at home with leisure time. However, this assumption should be confirmed for a larger population.

TABLE I. AVERAGE TIME PER DAY SPENT ON INTERNET APPLICATIONS PER HOUSEHOLD.

| Household | All days [min/day] | Weekends [min/day] |
|---|---|---|
| *1* | 253 | 111 |
| *2* | 432(IP1) / 75(IP2) | 614 (IP1) / 66(IP2) |
| *3* | 588 | 496 |
| *4* | 196 | 154 |
| *5* | 1047 | 1162 |

Concerning application usage, the frequencies of use reported by the household members in the web questionnaires corresponds well with the measured data for the most part. The test pilots were able to give estimations of how often they used specific applications such as Spotify or file sharing applications and how often they visited certain types of web sites such as newspapers or banks. Deviations in the estimation of frequencies were however seen for more general questions about usage. For example, the users were able to estimate how often they use Spotify but had a harder time answering questions about how often they listened to streaming music or watched streaming media. Here it is important how the questions are posed and if the user understands which medium he/she is using.

 ISBN: 978-1-61208-193-9 

TABLE II. USER PENETRATION OF SPECIFIC APPLICATIONS SEEN IN THE TRAFFIC MEASUREMENTS.

| Application Household | HTTP | HTTP media stream | BitTorrent | Spotify | SSL | Flash video | MSN messenger | Skype |
|---|---|---|---|---|---|---|---|---|
| *1* | x | x | x | | x | x | x | |
| *2* | x | x | x | x | x | x | x | x |
| *3* | x | x | | x | x | x | | |
| *4* | x | x | | | x | x | | |
| *5* | x | x | x | | x | x | x | |

The user penetration of a number of applications seen in the traffic measurements is found in Table II. This is in good agreement with the answers from the web questionnaires. As is expected, HTTP is used by all households as well as the SSL protocol, which is used by for example Internet shops and banks. All of the households also use HTTP media stream as well as flash video meaning that they look at streaming material on the Internet.

Although the estimation of frequencies agreed fairly well with the traffic measurements, there were discrepancies between the approximations of the time spent on different applications and the measured time. For example, three of the households were able to give a reasonable figure as to how much time they spent using Internet applications though tending to slightly underestimate the actual figure, see Table III. The deviation between measurements and questionnaire was especially large for household number 2. We note that this is a family household and the total time is more difficult to estimate.

TABLE III. TIME SPENT ON INTERNET APPLICATIONS PER HOUSEHOLD AND PER WEEK.

| Household | Questionnaire [h/w] | Measured HTTP [h/w] | Measured ALL [h/w] |
|---|---|---|---|
| *1* | 14 | 16 | 30 |
| *2* | 97 | 28 (IP#1) | 50 |
| *3* | 8 | 10 | 69 |
| *4* | 7 | 14 | 23 |
| *5* | 49 | 48 | 122 |

Comparing the diaries with the traffic measurements we find that they are mainly in accordance with each other. However, the traffic measurements add details to the picture given by the diaries, showing the potential that the method has in analyzing user behavior. A limitation of a diary is that users may not record all activities, either because they are done very frequently or because they are a natural part of your daily life that you do without reflecting that you are actually using the Internet. For example, one of the respondents had noted down “watching TV” in the diary. From the traffic measurements, we find that she was watching SVT play (streaming video from the Swedish state television) while at the same time being active on several community sites. This raises, among other things, an interesting question: What is actually meant by watching TV in the future when even more TV material will be available on-demand over the Internet?

The questionnaire contains questions concerning file sharing. The answers correspond well with what is measured and the frequencies of use recorded in the questionnaires agree well with the measurement. Although the use of file sharing applications is a sensitive question in public debate, the test pilots seem comfortable answering questions about this. Three of the households have used file sharing applications during the measurement period but none of the households are heavy users. It should be noted here that the test pilots are used to answer questionnaires, so they may not be representative in the sense that they may be less shy than other users when sharing sensible information.

## V. CONCLUSION AND FUTURE WORK

In this paper, we have proposed and applied a unified methodology using three different methods to study Internet user behavior: traffic measurements, web questionnaires and diaries with the purpose of verifying and comparing the different methods as well as gaining more insight into user behavior.

From the measurements, we conclude that the test pilots are well able to describe some of their short term behavior seen in the diaries, although some activities were not noted in the diaries. The long term behavior seen in the web questionnaires are fairly accurate in describing frequencies of use specific applications and visits to specific web sites. The estimation of the amount of time spent on different activities was seen to differ from that of the traffic measurements, with a slight tendency to under-estimate the time spent. An even more powerful conclusion is the complex and rich picture of user behavior, which is obtained via traffic measurements. Here, details and behaviors that are not exposed in diaries or questionnaires are visible. This gives new insights into user behavior as well as valuable feedback for better construction of question based investigations in the future.

Another major result of the study is the importance of the testbed to the study. Here, we have the possibility of making measurements in a controlled environment with real end users. We gain much more insight into the behavior of the end users than can be obtained from only questionnaires or diaries. We also gain understanding of what the end user experiences that complements the traffic measurements. From the network side, this can be used to improve the quality of service both from the technical and the end user perspective.

Our future work will continue with a wider study to follow up on the results presented here as well as a continued development of our testbed, which makes these types of measurements possible.

 ISBN: 978-1-61208-193-9 

## ACKNOWLEDGMENTS

The work in this project has been partly financed by the VINNOVA project Broadband Behavior and the European Regional Development Fund project Acreo National Testbed. Also, it has been partly funded by the Celtic IPNQSIS project and Swedish National Strategic Research Area (SRA) within the program TNG (The Next Generation) and the project eWIN.

## REFERENCES

[1] M. Chesire, A. Wolman, G. M. Voelker, and H. M. Levy, "Measurement and Analysis of a Streaming-Media Workload", Proceedings of the 3rd conference on USENIX Symposium on Internet Technologies and Systems - Volume 3, 2001.

[2] M. Alvarez-Campana, A. Azcorra, J. Berrocal, D. Larrabeiti, J. I. Moreno, and J. R. Pérez, "CASTBA: Internet Traffic Measurements over the Spanish R&D ATM Network", 5th HP Openview University Association Workshop, 1998.

[3] Y. Bhole and A. Popescu, "Measurement and Analysis of HTTP Traffic", Journal of Network and Systems Management, pp. 357-371, Vol. 13, No. 4, December 2005.

[4] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot, "Packet-level traffic measurements from the sprint ip backbone", IEEE Network, pp. 6-16, vol. 17, no. 6, 2003.

[5] J. Zhang, J. Yang, C. An, and J. Wang, "Traffic Measurement and Analysis of TUNET," 2005 International Conference on Cyberworlds (CW'05), Nov. 2005.

[6] J. George, "Researching Life in E-Society with Diary Studies", Paper 67, Proceedings of the 2006 European Conference on Information Systems, Gothenburg, Sweden, June 2006.

[7] K. Vermaas and L. van de Wijngaert, "Measuring Internet Behavior: Total Time Diary and Activity Diary as Research Methods", *Journal of Information Technology Theory and Application (JITTA)*: Vol. 7: Iss. 1, Article 11, 2005.

[8] O. Findahl, "Vad säger Internetstatistiken? (What do the Internet Statistics Tell Us?)", Stiftelsen för internetinfrastruktur .SE, 2008.

[9] C. P. Larsen, C. Lindqvist, H. Pathirana, R. Lindström, E. Modin, and A. Aurelius, "ANT: The Acreo National Testbed – Status and Plans", in proc. *NOC 2007*, Stockholm, Sweden, June 18-21, 2007.

[10] http://www.acreo.se/en/Technology-Areas/Broadband-Technology/Projects/Current-Projects/Acreo-National-Testbed [Accessed 10th April 2012].

[11] M. Kihl, C. Lagerstedt, A. Aurelius, and P. Ödling, "Traffic analysis and characterization of Internet user behavior", In Proc. ICUMT 2010, pp. 224-231, Moscow, Oct. 2010.

[12] Procera networks, http://proceranetworks.com [Accessed 10th April 2012].

[13] W. S. Cleveland and D. X. Sun, "Internet Traffic Data", Journal of the American Statistical Association, 1995, pp. 979-985. Reprinted in *Statistics in the 21st Century* , edited by A. E. Raftery, M. A. Tanner, and M. T. Wells, Chapman & Hall/CRC, New York, 2002.

 ISBN: 978-1-61208-193-9 