

Creating Personal Bandwidth Maps using Opportunistic Throughput Measurements

Ghulam Murtaza, Andreas Reinhardt, Mahbub Hassan, Salil S. Kanhere
School of Computer Science and Engineering
The University of New South Wales, Sydney NSW 2052, Australia
{gmurtaza,andreasr,mahbub,salilk}@cse.unsw.edu.au

Abstract—The ongoing success of smartphones and tablet computers, combined with the widespread deployment of cellular network infrastructure, has paved the way for ubiquitous Internet access. Access to mobile services has become a commodity for many commuters on public transport vehicles. On their daily trips to work and back, however, people often experience varying throughput rates due to the different capacities of network cells and the channel quality to the cell site. Links with reduced or no throughput are clearly unfavorable when users need to download large files or engage in synchronous communication activities. We thus introduce the notion of opportunistic personal bandwidth maps (OPBMs) in this paper. OPBMs allow the user to schedule activities with high throughput demand to parts of their journey where the bandwidth requirements are likely to be met. Users create their own OPBM by means of opportunistically monitoring their throughput during access to the cellular network and consolidating these individual measurements. Due to the opportunistic nature of our approach, no additional data transfers are required. Our measurements for more than 70 commutes show that the achievable throughput for road segments is highly variable across different trips. Still, the availability of OPBMs allows users to make decisions (e.g. to download a large file) when traveling along the segment with highest expected throughput.

I. INTRODUCTION

Although cellular network operators are continuously upgrading the capacity of their networks, it is extremely difficult for them to guarantee a stable and consistent bandwidth to every user at all times and locations. As a result of radio interference, potential scheduling issues, and varying cell utilization, the bandwidth available to each individual user commonly fluctuates over time and space [1]. This uncertainty about the available bandwidth makes it challenging to guarantee quality of service for emerging real-time mobile applications such as video streaming and synchronous voice chats. The spatial dependency of network throughput has thus resulted in the concept of bandwidth maps [2], which provide available bandwidth estimates to users based on previous measurements. A large number of throughput observations are recorded for all locations of a given commute and consolidated into a bandwidth map, which reveals the probability distribution of bandwidth as a function of location. For scenarios with highly mobile users, these bandwidth maps can be used to support applications in many ways, e.g., by pre-fetching data before entering a region with low bandwidth or postponing the download of email attachments to maintain more bandwidth for applications with real-time requirements. In simulations, researchers have demonstrated the huge potential of bandwidth maps for improving application performance [3]–[6].

While the potential applications of bandwidth maps are discussed in numerous publications, little has been reported on the practical methods for constructing such maps. Creating a wide area bandwidth map that stores bandwidth probability distributions for every single possible geographic location is a challenging problem due to its scale. The degree of spatio-temporal bandwidth dependency has been assessed by Deshpande et al., who have shown that the noticeable variation in bandwidth distribution already occurs when moving by only 20 meters [1]. A map that stores bandwidth distribution at this granularity would need to record billions of bandwidth observations to map the entire road network of a modern city. To make the problem worse, the map needs to be dynamically updated and maintained at all times due to possible changes in the bandwidth distribution for some locations arising from network upgrades or the construction of new roads. Finally, the deployment of the resulting map to individual mobile devices would be highly impractical because of its large size. We have deliberately not considered an approach that follows the participatory sensing paradigm [7] due to the low relevance of other people's maps for each individual user.

In this paper, we thus utilize smartphones to develop *personal* bandwidth maps, which contain the routes frequently travelled by the user and during the times the user commonly travels along these routes. In particular, we show how the user's smartphone can opportunistically record the experienced bandwidth in each location of the user's commuting route. By combining all recorded observations, the personalized bandwidth map for the user can be constructed. The repetitive nature of human mobility results in the collection of a large number of bandwidth samples along frequently travelled routes, thus improving the statistical significance of the personal bandwidth map and keeping its storage requirement small. To the best of our knowledge, we are the first to report the construction and evaluation of such personal bandwidth maps for commercial smartphones. We make the following contributions:

- We propose a method to assess the available bandwidth based on opportunistically monitoring TCP flows on a smartphone and detail its implementation on an Android-based handset.
- We use our prototype implementation to construct a bandwidth map for two specific routes in Sydney and demonstrate the practicality of the proposed method.
- We analyze the bandwidth distributions of different segments of the route and propose means to stochastically model the bandwidth distribution.

The rest of the paper is organized as follows. Related work is analyzed in Section II. We present our proposed opportunistic location-based bandwidth measurement methodology and implementation details in Section III. Field trials and data collection are reported in Section IV, and we conclude our paper in Section V.

II. RELATED WORK

From completely independent measurements, researchers from different continents have discovered that distributions of observed bandwidth from 3G networks are strongly correlated to locations of a road network. Yao et al. [8] measured bandwidth of three different mobile operators from a moving vehicle repeatedly traveling through the same route in parts of Sydney. When bandwidth observations were analyzed for every 500 meter road segments, they found that the distributions of bandwidth observations were different for different road segments for all three major cellular network providers. Later, Deshpande et al. made the same observation in New York [1], but this time they demonstrated that the phenomenon holds even for 20-meter road segments. The implication of this independently verified observation is that the uncertainty in mobile bandwidth can be reduced using *location-specific* bandwidth distributions for a given operator, instead of using a general distribution. This has motivated researchers to study the application of using bandwidth maps, which store bandwidth distribution information for every single road segment of a given route or road network.

Several new applications of bandwidth maps have been considered by different research groups. Using simulations, Yao et al. [8] demonstrated that access to a bandwidth map could significantly improve the scheduling performance of a multi-homed mobile router, which connects an onboard vehicular WiFi network to the Internet via multiple mobile networks. Sing et al. [3] and Curcio et al. [4] have independently demonstrated that use of bandwidth maps leads to smoother video streaming in mobile environments. Recently, Siris and Kalyvas [6] proposed a new method to improve mobile data offload performance using bandwidth maps. Finally, Hojgaard-Hansen et al. [5] have shown that mobile network performance maps that store information beyond bandwidth, such as round-trip-time, can be used to reduce TCP overhead in mobile communications. Bandwidth estimation is an active area of research. Researchers have used active probing based techniques, which calculate the bandwidth by explicitly transferring some data and measure the amount of time taken to transfer that data. Popular techniques for estimation include Packet Pairing [9, 10], Packet Trains [11] and Self Loading [12]. However, we consider it unlikely that users would be interested in applications that intentionally use up data volume, because most mobile phone plans are limited in their included traffic. The aforementioned active probing methods are thus not suitable for the creation of personal bandwidth maps. As an alternative, in this paper, we propose a passive measurement approach.

While new applications of bandwidth map continue to appear, to the best of our knowledge, practical methods for constructing a bandwidth map are yet to be reported in open literature. There are some commercial products, such as routemetrics [13], which allow users to download a network

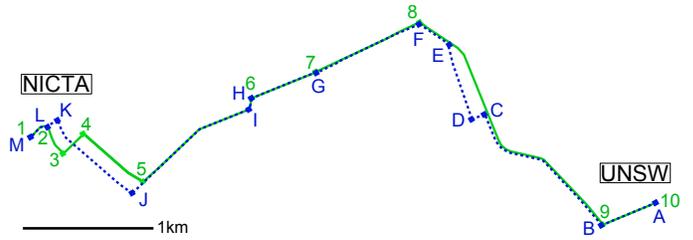


Fig. 1. Map for UNSW to NICTA with waypoint annotations.

testing application and report back network performance to a central database if the users use the application to test the network. Such data is then used to create a *network coverage map* of a city for a given operator. However, these coverage maps only report the average performance of the operators to help users compare different operators for locations where they use mobile network most frequently. Our proposed OPBMs are different from these centralized measurements in nature, as our method measures bandwidth passively and transparently each time the user browses the Internet. The bandwidth map we propose can be created from such passive user data and it can provide information useful for both humans and machines. Additionally, in contrast to centrally collected network coverage maps, OPBMs provide information at a much finer granularity.

III. OPBM CONSTRUCTION

We propose a scheme for bandwidth estimation, which is based on opportunistic throughput measurements. The methodology has specifically been chosen to allow for its implementation on a mobile phone, where it executes as a background process and collects data from users traveling on any means of individual or public transportation. The construction of a user's OPBM is a two-step process. First, using this opportunistic bandwidth measurement scheme, we collect many bandwidth measurements during the user's travel on individual or public transport. In a second step, we use these bandwidth samples to construct the OPBMs for personal use. When traveling in a city, there is high possibility that different routes share a number of sub-routes, e.g., individual road segments. We thus perform our analysis on the granularity of these segments of a route. This decision makes it possible to use measurements made for same segment across different routes, and thus to aggregate readings collected at different times.

A. Road Segments

When searching for a route on Google Maps¹, the resulting route is typically divided into segments, generally spanning from 10 to 500 meters within city (larger on big highways). We propose to use this waypoint-based segmentation of the road segments in different routes of a city. For our experiments, we have used two routes in Sydney, one from The University of New South Wales (UNSW) to National ICT Australia (NICTA) and another from NICTA to UNSW, which takes a slightly different route. Figure 1 shows these two routes returned by Google Maps, super-imposed on each other. The route in green line is from NICTA to UNSW and dotted blue line shows the route from UNSW to NICTA. We label the waypoints

¹<http://maps.google.com.au>

from UNSW to NICTA as A through M and NICTA-UNSW waypoints are labeled as 1-10. As a result we get 12 segments and 9 road segments for both routes respectively. The exact coordinates of the waypoints are also given in Appendix A. This waypoint-based division is possible even if the user is traveling on other means of transport, e.g., by bus.

B. TCPDUMP-based Passive Bandwidth Measurement

In this work, we focus on TCP based traffic due to overwhelming use of TCP in today's Internet-based applications. The proposed passive bandwidth measurement is based on capturing all TCP packets downloaded during a trip. After computing the bandwidth for all road segments in the trip, the captured TCP packets are discarded. We have found that most smartphone browsers launch multiple TCP connections in parallel to speed up the download of a page, which typically contains multiple objects. Monitoring TCP throughput of individual TCP connections would therefore not provide accurate estimates of the available bandwidth on any given road segment, because multiple TCP connections could overlap at the same time. Our methodology thus addresses such TCP connection overlaps. Another observation we made is that users tend to spend some time consuming the contents of a page once downloaded before starting new downloads. Therefore, we can expect that there would be periods when there is no active traffic. Our methodology takes these *idle periods* into consideration when calculating the bandwidth for a given road segment. To explain the proposed bandwidth computation algorithm, we use the following notations:

- m_i : total number of TCP packets downloaded in i^{th} road segment
- d_{ij} : number of bytes in j^{th} packet received in i^{th} road segment
- n_i : number of idle periods in i^{th} road segment
- t_{ij} : duration of j^{th} idle period in i^{th} road segment
- T_i : trip duration of i^{th} road segment

The estimated bandwidth for i^{th} road segment is then derived as shown in Equation (1).

$$B_i = \frac{8 \times \sum_{j=1}^{m_i} d_{ij}}{T_i - \sum_{j=1}^{n_i} t_{ij}} \text{ bps} \quad (1)$$

Once a new sample is available for a given road segment in the bandwidth map, the distribution parameters, such as mean and standard deviation of bandwidth, are updated. This way, the bandwidth map can be kept up-to-date as the user travels through the same road segments over and over again.

Let us make an example. Consider a set of TCP flows (some of which can be parallel) 1, 2, ..., n . Flow i starts at time s_i and finishes at time $f_i > s_i$. Two flows i and j overlap if $[s_i, f_i] \cap [s_j, f_j] \neq \emptyset$. Each TCP flow i has m total packets where m varies for each TCP flow. The objective is to find the combined throughput for all TCP flows. To this end, we sort the TCP flows on their start times such that $s_i < s_j$ as shown in Figure 2. As all the packets are TCP packets, flows can be differentiated by unique socket addresses even between same machines. We scan all the data transferred in these flows

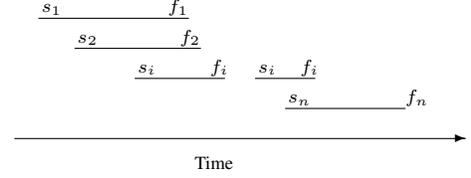


Fig. 2. Example for the sorting of TCP flows.

and time taken to transfer this data and we exclude the idle time between these flows. Our algorithm maintains an active set A as well counters for the total transferred data D and the total time T . A contains the identifiers for active TCP connections. D is total data downloaded by the device using TCP connections in A , and T is the time for which these TCP connections stay active. Our algorithm examines each packet and looks for the start of TCP flows. Whenever it encounters the first three-way handshake it record the timestamp and puts the socket address into A . It keeps on capturing data packets which belong to any of the TCP flows in A . We record the timestamp of the first packet of the first TCP flow in A . Then we keep examining the remaining data packets and keep record of number of bytes downloaded in D . Whenever we encounter a FIN packet, we remove the corresponding socket address from A . Once the last active flow in A has ended, we use its timestamp to calculate the total session time and add it to T . Finally, we calculate the estimated bandwidth by dividing the total data transferred by the session duration.

C. Implementation Details

To implement the algorithm of Equation (1), we need to derive all the parameters of the algorithm first. This is achieved using the well-known TCPDUMP tool. We use it to capture each TCP packet downloaded during the trip, for which it records the size, the timestamp, and all TCP flags. Android does not have a built-in TCPDUMP utility, so we have used the freely available version from [14]. Our application records location information every 500 ms. From this collected location information, we match (using the closest match) the user's current GPS position to known road segment separators (i.e., pre-loaded waypoints) and derive *time boundaries* for each road segment. Once the time interval for a given road segment is determined, we identify all packets (using their timestamps) that were received within this road segment, which allows us to compute the numerator of Equation (1). The calculation of the denominator in Equation (1), however, involves the identification of idle periods. Since no TCP connection is active during idle periods, we make use of TCP flags to this end. The start of a TCP connection is identified by detecting the SYN flag in a packet. Similarly, the end of a TCP connection is identified by a FIN packet. Since all TCP packets for a given TCP connection carry a unique socket number, we can easily match a FIN to its corresponding SYN for same TCP connection, which allows us to work out how many and which TCP connections are active at any given time. An idle period starts when the number of active TCP connections in the system becomes zero and ends as soon as a new TCP connection starts. We trace the time interval of all idle periods (t_{ij}) for all road segments, which allows us to compute the denominator of Equation (1).

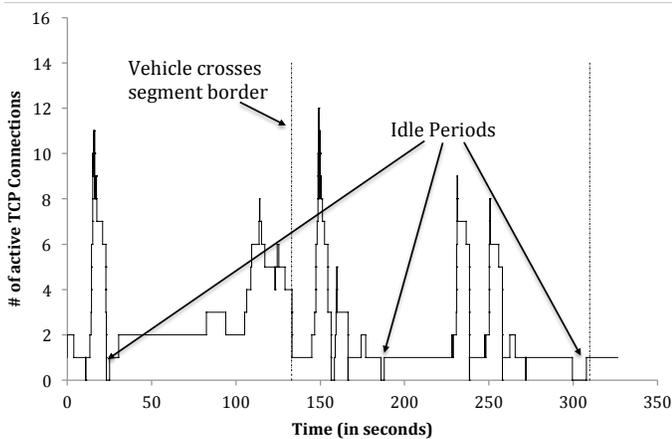


Fig. 3. Number of parallel active TCP connections for the first two segments.

IV. FIELD TRIALS AND ANALYSIS

We have implemented our previously described algorithm for an Android-based mobile phone and used it for collecting bandwidth data for the two routes introduced in Section III-A. This section provides the details of these field trials and a subsequent analysis of the collected data.

A. Trials and Data Collection

We have 12 and 9 segments for the UNSW-NICTA and NICTA-UNSW routes (cf. Figure 1) respectively. Two volunteers were recruited to conduct the tests, one of whom was driving a car, while the other was using the Android phone in the passenger seat. The user had to start the application at the start of the trip and they were asked to use a number of Internet applications. At the end they stop the application when the vehicle had reached the destination. After each trip, the output of these two files along with the preloaded route information files were analyzed to compute the bandwidth for each of the road segment. The data connection used was provided by the Australian operator OPTUS, and the 3G service used was a mix of HSDPA and UMTS during these trips. We have made a total of 74 trips for each routes, and as a result collected 74 data points for each road segment of both routes. These trips were made on 13 consecutive days, on both weekdays and weekends. On any of these days, our volunteers chose time of their convenience and drove from UNSW to NICTA and the back. The average time of each round trip was 30 minutes. This way, they made around 6 trips everyday. During these trips, the passenger was asked to use YouTube, Facebook, and Web browsing, all of which use TCP for their data transfer.

As mentioned in section III-C, Android optimizes TCP throughput by starting multiple parallel TCP connections. Figure 3 shows the number of parallel TCP flows for the first two segments on the first trip of the UNSW-NICTA route. This graph shows the data from one trip only and makes clear that multiple concurrent TCP flows as well as some idle periods exist. The two vertical lines at 133 and 309.8 seconds indicate the times when the vehicle crossed a waypoint. The figure confirms that periods exist during which multiple parallel TCP connections are used for Internet activity, as well as times when there are no active TCP connections at all. When calculating

TABLE I. BANDWIDTH (KBPS) STATISTICS OF UNSW-NICTA ROUTE FROM EACH SEGMENT OVER ALL 74 TRIPS

Seg. No.	Mean	ST.Dev	W	p-value
A-B	212.01	262.08	0.733	0.000
B-C	315.9	174.48	0.982	0.423
C-D	198.82	135.6	0.93	0.0002
D-E	372.36	322.47	0.740	0.000
E-F	242.77	222.28	0.861	0.000
F-G	253.31	295.65	0.630	0.000
G-H	240.21	218.34	0.856	0.000
H-I	226.39	239.05	0.813	0.000
I-J	312.79	339.48	0.756	0.000
J-K	229.32	152.42	0.918	0.0002
K-L	191.5	159.22	0.905	0.000
L-M	419.19	280.16	0.739	0.000

the amount of time, we exclude these idle periods. For some trips, some road segments recorded zero bandwidth due to not receiving any TCP packet at that time as can also be seen for trip 43 in Figure 5b. We exclude these zero bandwidth measurements from our subsequent analysis.

B. Analysis

In order to build the personal bandwidth maps, the bandwidth readings for all the traveled road segments need to be stored in an efficient way. Although memory restrictions rarely exist on current smartphones, the analysis of all stored TCP connection information can be time-consuming when a large number of samples have been collected. As a result, we store both the raw data (in order to calculate statistics after further data have been added) as well as their processed representations, as outlined below.

1) *Mean and Standard Deviation:* We have tabulated the mean and standard deviation of all individual segments for both routes in Tables I and II. From the tables, it can be seen that the standard deviation is relatively high, which points out the limited use of the mean value for making bandwidth predictions. Further analysis whether the data can be approximated using a normal distribution (by analyzing their W and p-value according to [15]), has shown that the data does not follow a Gaussian distribution and that additional data points are required to model it properly. The results of the tests for both routes can also be seen in the tables. By looking at Figures 5a and 5b, one can furthermore deduce that the mean value is not a good representative of expected bandwidth for particular segment because of the high bandwidth variability for the same segment at different times. Therefore we need a better representation of this bandwidth data to estimate the probable future bandwidth expectation for a segment.

TABLE II. BANDWIDTH (KBPS) STATISTICS OF NICTA-UNSW ROUTE FROM EACH SEGMENT OVER ALL 74 TRIPS

Seg. No.	Mean	ST.Dev	W	p-value
1-2	330.48	369.6	0.806	0.000
2-3	288.04	202.88	0.943	0.002
3-4	157.51	297.84	0.702	0.002
4-5	212.82	176.06	0.910	0.000
5-6	219.08	159.81	0.940	0.001
6-7	213.61	193.31	0.764	0.000
7-8	197.71	174.84	0.802	0.000
8-9	175.08	163.48	0.865	0.000
9-10	203.73	198.89	0.589	0.000

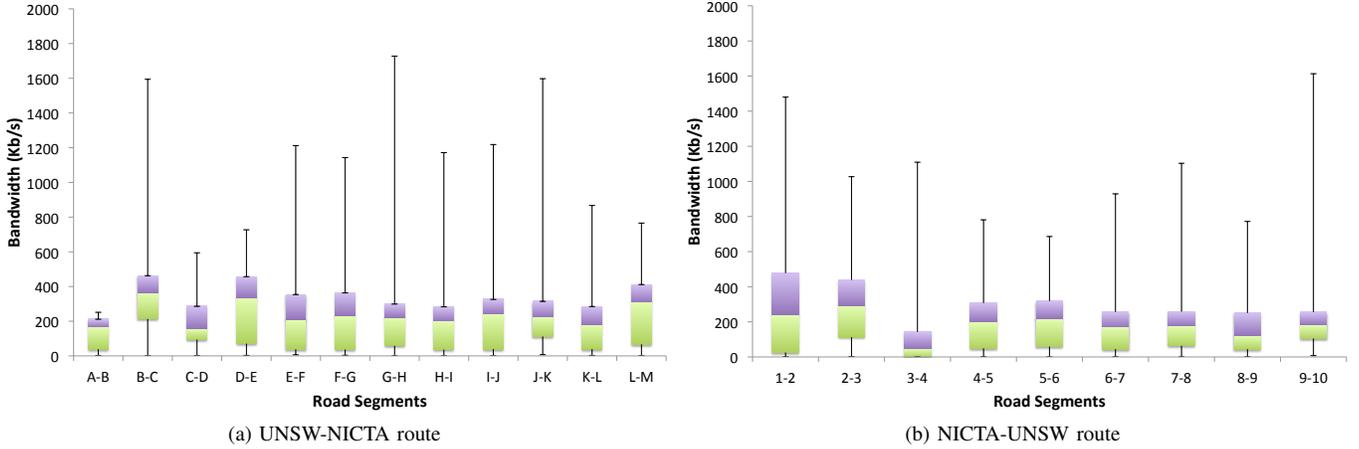


Fig. 4. Bandwidth statistics obtained in form of box-plots over all 74 trips

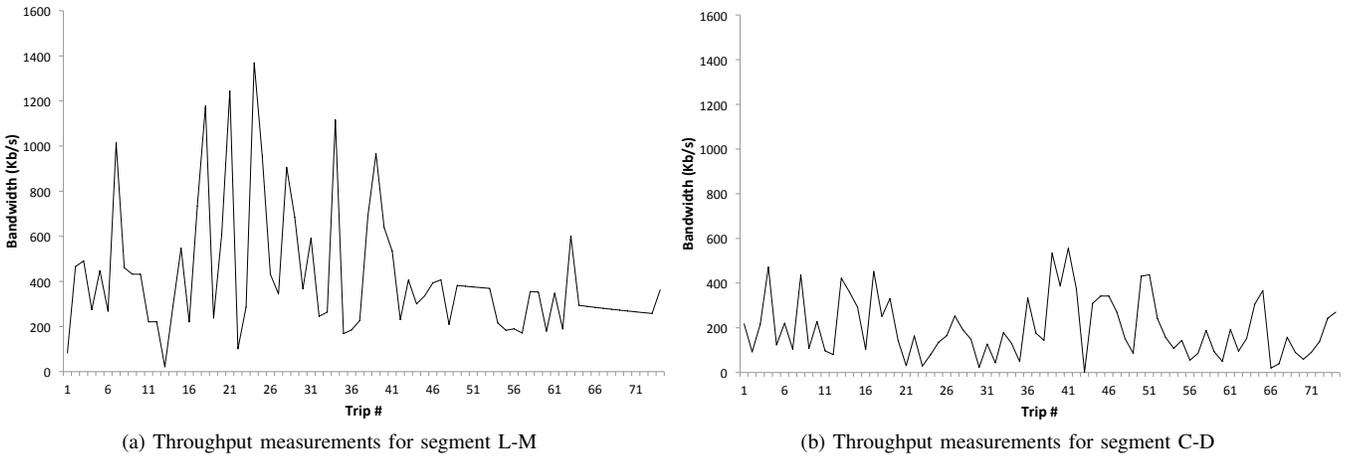


Fig. 5. Visualization of raw bandwidth measurements for a high-throughput segment (L-M) and a low-throughput segment (C-D) across all 74 conducted trips.

2) *Box and Whisker Representation*: Given the observed bandwidth variations, it is clear and consistent with previous findings [1] that mobile bandwidth is extremely difficult to predict. Solely relying on two data points, namely mean and standard deviation, is clearly insufficient. Instead, we have plotted the bandwidth data collected for two routes over all 74 trips in the form of box-plots in Figure 4. They present a very interesting view of overall bandwidth data and at the same time give us an insight towards storing them in a compact manner. In the box-plots shown in the figure, the x-axis shows the segment and the y-axis represents the bandwidth in Kb/s for each of the 74 trips conducted during the field trials. It is interesting to compare the input data for segments C-D and L-M, as visualized in Figure 5. One can see that box and whisker representation for C-D shows that the 75th percentile is 250 Kb/s, its median lies at 150 Kb/s, and the 25th percentile value is 95 Kb/s, despite its comparably high observed maximum value of 550 Kb/s. In comparison to segment L-M, where the median values is 310 Kb/s at the 75th percentile ranges at 410 Kb/s, segment C-D is thus of less favorable throughput. As a general observation, the distance between first quartile and median is often highly different from the distance between third quartile and median, confirming the insufficiency of

modelling the distributions using mean and standard deviation.

Using the box-plot representation, we can make the decision to schedule the tasks with high bandwidth requirements, such as synchronous video chats or downloads of large files can be scheduled in segments with a high expected throughput. Like the segment between waypoints L and M, one can see that segments B-C and D-E have also very high median and 75th percentile values with a small difference between them. It shows that these segments are also more likely to have higher bandwidths in subsequent trips. In contrast segments A-B and J-K show more resemblance to C-D in terms of bandwidth statistics and one should not prefer these segments for bandwidth-intensive tasks. Similar deductions can be drawn about the segments of NICTA to UNSW route by looking at Figure 4b. Segments 1-2 and 2-3 are segments with better bandwidth possibility and segments 3-4 and 9-10 should not be preferred for bandwidth intensive tasks.

We can differentiate between a positive and a negative variation of the readings by considering the 25th and 75th percentile marks. While the standard deviation combines both higher and lower bandwidth readings into a single value, these additional data points allow for modelling the distributions at

a finer granularity. Let us assume a scenario where bandwidth for a segment is mostly low and only very high for few trips. As a result, the mean will rise measurably, but so will the standard deviation. By looking at the size of the interquartile range, more accurate predictions can be made.

Therefore, we propose to use box-plots as the representation for our OPBMs, which store five values for each segment. These are minimum, 25th percentile, median, 75th percentile and maximum. We argue that the box-plot representation is better suited for the purpose of modelling the distribution of the encountered bandwidth measurements and gives us a better insight to expected bandwidth for a segment as compared to when we store simple mean and standard deviation.

V. CONCLUSION AND FUTURE WORK

In this paper, we have presented a practical implementation of personalized bandwidth maps. Instead of utilizing throughput probing sequences, we have opportunistically collected statistics about the TCP traffic generated by the users passively whenever they surf the Internet. This allows us to create personalised maps without incurring additional bandwidth overheads. We have conducted a field study in the city of Sydney, Australia, where a team of volunteers were driving a car along two pre-defined routes for more than 70 times in order to collect the bandwidth data that was used in our evaluation.

Our findings show that available bandwidth varies significantly, yet trends can be determined for individual route segments. By analyzing the collected throughput measurements by means of box-and-whisker plots, patterns for some segments became clear. Considering the requirement of only five parameters, the memory footprint of our approach is small.

We argue that OPBMs are better at capturing the expected network performance than maps that characterise radio characteristics. While there certainly exists some correlation between signal quality and bandwidth, there are several other factors such as network load and operator scheduling policies which may influence the available bandwidth. OPBMs allow us to capture the cumulative effect of all these factors on network performance.

In future we plan to build a complete system which does the task scheduling along with building bandwidth maps.

VI. ACKNOWLEDGMENTS

We are very thankful to Bandar Alqahatani and Ayub Bokani for helping us in field trials. Without field trials, it would not have been possible to complete this work.

REFERENCES

- [1] P. Deshpande, X. Hou, and S. R. Das, "Performance comparison of 3g and metro-scale wifi for vehicular network access," in *Proceedings of Internet measurement conference (IMC)*, Melbourne, Australia, 1-3 November 2010.
- [2] J. Yao, S. Kanhere, and M. Hassan, "Improving qos in high-speed mobility using bandwidth maps," *IEEE Transactions on Mobile Computing*, vol. 11, no. 4, pp. 603–617, April 2012.
- [3] V. Singh, J. Ott, and I. D. D. Curcio, "Predictive buffering for streaming video in 3g networks," in *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, San Francisco, USA, 25-28 June 2012.

- [4] I. D. Curcio, V. K. M. Vadakital, and M. M. Hannuksela, "Geo-predictive real-time media delivery in mobile environment," in *Proceedings of the 3rd workshop on Mobile video delivery (MOVID)*, Firenze, Italy, 25 October 2010.
- [5] K. Hojgaard-Hansen, T. K. Madsen, and H.-P. Schwefel, "Reducing communication overhead by scheduling tcp transfers on mobile devices using wireless network performance maps," in *Proceedings of the 18th European Wireless Conference*, Poznan, Poland, 18-20 April 2012.
- [6] V. A. Siris and D. Kalyvas, "Enhancing mobile data offloading with mobility prediction and prefetching," in *Proceedings of the seventh ACM international workshop on Mobility in the evolving internet architecture (MobiArch)*, Istanbul, Turkey, 22-26 August 2012.
- [7] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory Sensing," in *Proceedings of the International Workshop on World-Sensor-Web (WSW)*, 2006, pp. 117–134.
- [8] J. Yao, S. S. Kanhere, and M. Hassan, "Geo-intelligent Traffic Scheduling For Multi-Homed On-Board Networks," in *4th International Workshop on Mobility in the Evolving Internet Architecture (MobiArch)*, Kraków, Poland, 22 June 2009.
- [9] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, ser. IMC '03. New York, NY, USA: ACM, 2003, pp. 39–44. [Online]. Available: <http://doi.acm.org/10.1145/948205.948211>
- [10] K. Lai and M. Baker, "Nettimer: a tool for measuring bottleneck link, bandwidth," in *Proceedings of the 3rd conference on USENIX Symposium on Internet Technologies and Systems - Volume 3*, ser. USITS'01. Berkeley, CA, USA: USENIX Association, 2001, pp. 11–11. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1251440.1251451>
- [11] N. Hu, P. Steenkiste, O. Spatscheck, and J. Wang, "Optimizing network performance in replicated hosting," in *Proceedings of the 10th International Workshop on Web Content Caching and Distribution*, ser. WCW '05. Washington, DC, USA: IEEE Computer Society, 2005, pp. 3–14. [Online]. Available: <http://dx.doi.org/10.1109/WCW.2005.15>
- [12] M. Jain and C. Dovrolis, "End-to-end available bandwidth: measurement methodology, dynamics, and relation with tcp throughput," *Networking, IEEE/ACM Transactions on*, vol. 11, no. 4, pp. 537–549, 2003.
- [13] ROOTMETRICS, accessed on 07 September, 2012. [Online]. Available: www.rootmetrics.com
- [14] Accessed on 28 October, 2012. [Online]. Available: <http://www.mediafire.com/file/njyqyxu3nmx/tcpdump.zip>
- [15] S. Shapiro, *An Analysis of Variance Test for Normality (complete Samples)*, ser. General Electric Company. Advanced Technology Laboratories. Technical information series. General Electric, 1964. [Online]. Available: <http://books.google.com.au/books?id=MkiZtgAACAAJ>
- [16] A. C. Berry, "The accuracy of the gaussian approximation to the sum of independent variates," *Transactions of the American Mathematical Society*, no. 49, 1941.

APPENDIX A WAYPOINT COORDINATES

TABLE III. WAYPOINT COORDINATES USED IN THE FIELD TRIALS.

UNSW-NICTA			NICTA-UNSW		
Waypoint	Longitude	Latitude	Waypoint	Longitude	Latitude
A	151.23033	-33.91993	1	151.19611	-33.89507
B	151.22639	-33.91942	2	151.19746	-33.89713
C	151.22332	-33.90901	3	151.19928	-33.89666
D	151.22239	-33.90887	4	151.20116	-33.90145
E	151.22362	-33.90378	5	151.20976	-33.90008
F	151.22256	-33.90156	6	151.21029	-33.90023
G	151.21492	-33.90094	7	151.21494	-33.90085
H	151.21029	-33.90023	8	151.22266	-33.90146
I	151.20976	-33.90008	9	151.2265	-33.91943
J	151.20014	-33.90175	10	151.23033	-33.91993
K	151.19823	-33.89499			
L	151.19743	-33.89505			
M	151.19611	-33.89507			