# Mobile Broadband Performance Measured from High-Speed Regional Trains

Jun Yao, Salil S. Kanhere and Mahbub Hassan
School of Computer Science and Engineering, University of New South Wales, Sydney, Australia 2052
Email: {jyao, salilk, mahbub}@cse.unsw.edu.au

***Abstract*—While mobile broadband performance measured from moving vehicles in metropolitan areas has drawn significant attentions in recent studies, similar investigations have not been conducted for regional areas. Compared to metropolitan cities, regional suburbs are often serviced by wireless technologies with significantly lower data rates and less dense deployments. Conversely, vehicle speeds are usually much higher in the regional areas. In this paper, we seek to provide some insights to user experience of mobile broadband in terms of TCP throughput when travelling in a regional train. We find that (1) using a single broadband provider may lead to a large number of blackouts, which could be reduced drastically by simultaneously subscribing to multiple providers (provider blackouts are not highly correlated), (2) the choice of train route may have a more significant effect on broadband experience than the time-of-day of a particular trip, and (3) the speed of the train itself has no deterministic effect on TCP throughput.**

## I. Introduction

Thanks to the rapid developments and coverage expansions in the wide area wireless networking (WWAN) technology, e.g., GPRS/UMTS/HSDPA, CDMA 1x/EVDO, it is now possible for users to access the Internet directly from their mobile devices even from moving vehicles, i.e., trains, coaches and buses. As a result, passengers can now make use of their long commute time either for work, e.g., by connecting to their corporate Virtual Local Area Networks (VLANs), or for entertainment, e.g., through web-browsing, blogging, social networking, or even watching streaming video.

Due to the popularity of mobile Internet access, there has been a significant interest in recent years to measure and characterize the performance of mobile broadband from moving vehicles [1]–[6]. Because most of the user population of mobile broadband is centered around the cities, the studies are almost exclusively conducted for metropolitan areas. Little has been reported for regional areas. Given that many country residents take long distance trains to commute to the cities, it is important to investigate the regional routes as well. As compared to metropolitan cities, regional areas are often serviced by wireless technologies with significantly lower data rates and less dense deployments. However, trains travel at much higher speeds in regional areas than in cities. The measurement studies from regional trains can help operators design a comprehensive broadband solution that delivers high quality services to both metropolitan and regional areas. Further, they can provide insights for developing robust communication protocols and algorithms between end systems and servers, which can perform reliably in challenging regional environments.

In this paper, we present an empirical study on the mobile broadband performance experiences from regional trains in the state of Victoria, Australia. These measurements, which log TCP throughput over WWAN networks, were conducted from onboard the fastest trains in Australia, travelling up to 160 Kmh. We consider different WWAN variants, e.g., GPRS/UMTS and CDMA1x/EVDO through different provider networks. Our findings can be summarized as:

- Use of a single provider leads to frequent service blackouts along regional corridors. Interestingly, we find that blackouts from different providers are not highly correlated. Our measurements show that use of just one more provider as a backup would significantly reduce the frequency and length of service blackouts.
- A user may experience vastly different broadband experience in different train corridors, but the broadband performance in a given corridor remains relatively more consistent throughout the day.
- When the train crosses city-regional boundaries, the user can expect significant change in TCP throughput due to higher capacity network deployments in the city compared to regional areas.
- Lastly, our results suggest that the train speed has no deterministic effect on TCP throughput. It is feasible to achieve high TCP throughput even when travelling at 130 Kmh.

The rest of this paper is organized as follows. Section II reviews the related works in mobile WWAN bandwidth measurements. We present our measurement methodologies in Section III. The findings from our measurement study are discussed in Section IV. This paper is concluded in Section V.

## II. Related Works

There are many existing measurement studies in the literature that investigated WWAN performance under mobile conditions in metropolitan areas. In [6], the authors found that the mobile channel data rate of 3G HSDPA is significantly more variable during 1-hour of driving along a route than that observed from stationary locations. They revealed that the rate fluctuation can be attributed to the variety of radio conditions that the vehicle encountered along the route. Justin et al. [2] present a measurement study over CDMA 1xEVDO from two buses in the Madison city. They reported that the TCP throughput seems to be correlated with the location due to the location-specific signal strength characteristics. The works in [1], [3], [4] reported the mobile measurements results on 3G WWAN networks (both CDMA 1xEVDO and HSPA) from different transport means, i.e. the mix of cars, trains, ferries and subways. These studies evaluated the performance of the 3G networks using various user applications, e.g, FTP and VoIP. The results concluded that link conditions are highly unstable in a mobile environment. Interestingly, [3] reported that even though the vehicular mobility generally degrades WWAN service quality, it surprisingly improves the fairness in bandwidth allocation among users and traffic flows. Rodriguez et al. [7] conducted mobile measurement studies on different WWAN networks using their MAR system. Their interesting results highlighted the performance diversity between different operators and technologies. They showed that, by exploiting the diversity, data download performance can be significantly improved through multihoming techniques, i.e., using multiple links simultaneously to access the Internet. The above studies have primarily concentrated on the WWAN performance within the populous metropolitan area, which is extensively covered by dense 3G networks. This paper, on the other hand, investigates the WWAN bandwidth conditions experienced from high-speed trains that travel into the regional areas.

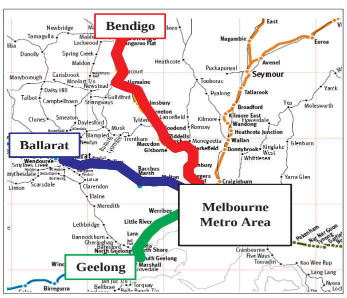


Fig. 1: Train network map of Melbourne and surroundings

## III. Measurement Methodologies

In this section, we briefly discuss the setup of our WWAN measurement studies from trains travelling between Melbourne and various regional cities in the state of Victoria. For the measurements, we have cooperated with a major train operator in Victoria, who provided us security clearances, free travel passes and location log files for each trip we made.

We focused on measuring the downlink TCP throughput from moving trains, since the downlink direction makes up the bulk of the traffic of a user. To study provider/network diversities, we conducted measurements on two WWAN networks that operate on different WWAN technologies. We label the providers as A and B, which operate on CDMA 1x/EVDO and GPRS/UMTS, respectively. For each measurement, we boarded a train with two laptops. Each computer was fitted with one PCMCIA WWAN modem (corresponding to one provider). Our goal is to study the WWAN reception experienced by users within the cabin, thus the modems were not connected to any external antennas.

During measurements, the computers continuously download the same large file (90405 KB) from a FTP server at the University of New South Wales (UNSW) in Sydney. The chosen file was large enough for the download to last for an entire trip. We commenced the file download on boarding the train and used the downloading manager Flashget [8]. Flashget automatically maintained one FTP session to download the file whenever possible. We used the popular packet analyzer tool, Wireshark [9], to record the traces of packets transmitted in each trip. The traces were analyzed offline at a later time. From the trace files, we calculate the instantaneous throughput of the downloading session for every second.

TABLE I: Observed average bandwidth in each trip

| Corridors | Trips | Avg. Tput. (Kbps) | |
|---|---|---|---|
| | | Prov. A | Prov. B |
| Mel-Gee | 1 (12:00-12:53) | 110.70 | 75.73 |
| | 2 (14:31-15:33) | 127.19 | 79.08 |
| | 3 (18:56-19:49) | 99.77 | 86.72 |
| | 4 (20:23-21:18) | 110.39 | 92.79 |
| Mel-Bal | 1 (08:05-09:31) | 88.66 | 47.42 |
| | 2 (10:10-11:26) | 60.84 | 33.78 |
| | 3 (13:08-14:27) | 91.41 | 38.26 |
| | 4 (15:00-16:27) | 66.21 | 32.48 |
| Mel-Ben | 1 (07:10-09:17) | 72.23 | 29.84 |
| | 2 (10:30-12:23) | 69.04 | 34.09 |
| | 3 (14:15-15:58) | 72.19 | 39.13 |
| | 4 (16:30-18:18) | 71.53 | 34.82 |

To correlate throughput data to geographic locations along each corridor, we need to also collect location information during each measurement. For this, we initially planned to carry a hand-held GPS device on-board to record location coordinates during each trip. However, we found that that the GPS signal is not available inside the cabin. To resolve this issue, we obtained location logs for the trips from the train operator. Each train is equipped with a positioning system, which automatically records the GPS time, relative distance from the Melbourne terminal and instantaneous train velocity pulse into a log file for each trip. We used these log files in conjunction with the processed TCP throughput traces in our analysis in the next section. Note that, this requires the system clocks on the measurement laptops to be synchronized with the GPS time available in the trip log files. To achieve this, before we boarded onto the train for each measurement, we connected the laptops to a hand-held GPS device to calibrate their system clocks. Thus, for each trip, we are able to calculate the instantaneous TCP throughput value for each 500m location segment travelled, by averaging all raw throughput samples observed within that segment.

We conducted measurements on the fastest trains in Australia, V/Locity trains [10], with a maximum speed of 160 Kmh. We conducted measurements on three popular corridors originated from Melbourne. The first corridor runs to Geelong (the largest regional city in Victoria), which is about 80 Km south-west to Melbourne. The second corridor runs to Ballarat (the third largest city in Victoria), which is about 120 Km west of Melbourne. The last corridor runs to Bendigo in central Victoria, about 180 Km north of Melbourne. We conducted four measurement trips on each of the corridors. The first and second columns in Table I summarize the details of the trips.

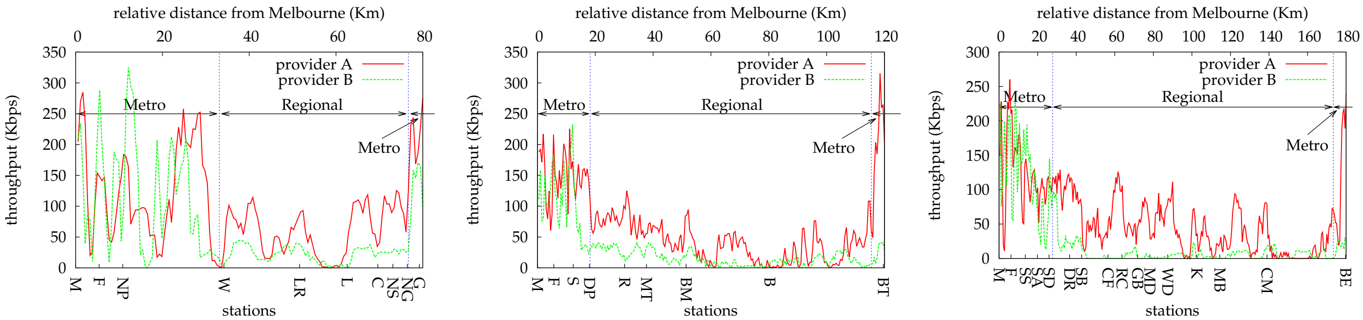


(a) geelong (b) ballarat (c) bendigo

Fig. 2: Instantaneous bandwidth along different corridors.
(Acronym for stations: B-Ballan, BE-Bendigo, BM-Bacchus Marsh, BT-Ballarat, C-Corio, CF-Clarkefield, CM-Castlemaine, DP-Deer Park, DR-Diggers Rest, F-Footscray, G-Geelong, GB-Gisborne, K-Kyneton, L-Lara, LR-Little River, M-Melbourne, MB-Malmsbury, MD-Macedon, NG-North Geelong, MT-Melton, NP-Newport, NS-North Shore, R-Rockbank, RC-Riddells Creek, S-Sunshine, SA-St Albans, SB-Sunbury, SD-Sydenham, SS-Sunshine, W-Werribee, WD-Woodend).

## IV. MEASUREMENT FINDINGS

In this section, we present the major findings from our measurements. We first investigate the effect of location and time on TCP throughput. Next we study the performance diversity between different WWAN providers and technologies. Lastly, the effect of speed on TCP throughput is discussed.

### A. Spatial and Temporal Variations

In Fig. 2, we plot the average TCP throughput (over all four trips) along each corridor for both providers. These plots show that the TCP throughput fluctuates frequently and significantly for both providers at different locations. We suggest the fluctuations are primarily caused by the constant changes in the radio signal quality, which are dictated by the varying distance between the train and the base station, multi-path fading and co-channel interference conditions.

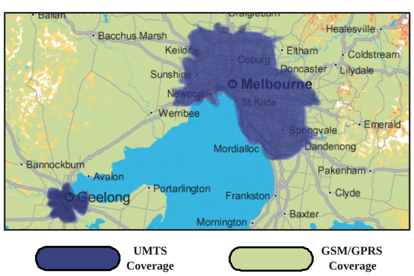


Fig. 3: Coverage map of provider B along the Melbourne-Geelong corridor [11]

Further, the graphs suggest that the TCP throughput is higher (i.e. 200-300 Kbps) at some portions of the corridors. For example, Fig. 2(a) shows that the TCP throughput of provider B drops noticeably from 200 Kbps to 50 Kbps when the train approaches the Werribee station. Also, when the train enters the metropolitan area of Geelong, the TCP throughput increases sharply. This observation is consistent with the geographical deployment of the network providers. For example, country areas are primarily serviced by 2G networks with lower data rates, due to the sparse population density and lower bandwidth demands in those areas. To confirm this, we show provider B's coverage map along the Melbourne-Geelong corridor in Fig. 3. Observe that only the metropolitan areas, i.e., Melbourne-Werribee and Geelong, are covered by UMTS. A large portion of this corridor is served by the GPRS. We are able to confirm similar observations in other corridors and for provider A.

In addition, we notice from Fig. 2 that there are a few service blackouts, where TCP throughput is zero due to coverage holes, along each corridor. In particular, we observe most of the blackouts occur in regional areas. This shows the reliability of data services is better in the metropolitan areas as compared to regional pockets. Interestingly, the service blackouts cannot be observed from Fig. 3. This suggests that coverage maps provided by service providers may not be always a reliable indicator of the actual user experienced network performance at a particular location. This highlights the need for the community to profile the actual WWAN performance at each location through participatory measurements from voluntary users [12]. For example, independent organizations, e.g., [13], now provide street-level coverage maps of WWAN and Wi-Fi networks. The maps are freely available online, which provide an unbiased comparison of different networks and aid consumers in selecting providers.

Recall that we made four trips on each corridor. Table I summarizes the average TCP throughput observed for each corridor, trip, and provider. It is evident that the downloading experience while travelling along different corridors is vastly different. Clearly, the Geelong corridor provides the highest TCP throughput, for both providers. This is consistent with the WWAN providers' network coverage maps. Since Melbourne and Geelong are the largest cities in Victoria, the 80 Km Melbourne-Geelong corridor is more populous as compared to the other two corridors. As a result, more area within this corridor has higher capacity 3G coverage. In fact, Fig. 2(a) shows that a user can experience higher 3G data rate at nearly 40% of this corridor for both providers. In comparison, due to the sparse population along the Ballarat and Bendigo corridors, 2G technologies are mostly deployed in these areas.

WWAN bandwidth not only varies across locations, but also may vary with time of day, due to possible varying network loads, i.e., number of active users, in the cellular towers. However, from Table I, we do not observe significant throughput disparity across different times of day in the same corridor. This may be because of the low number of active

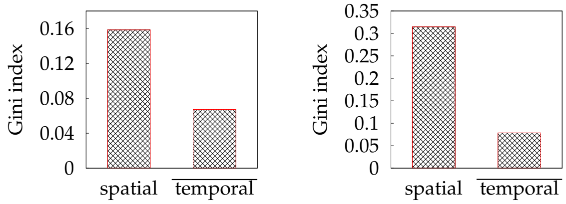

(a) Provider A (b) Provider B

Fig. 4: Comparisons of spatial and temporal disparities of TCP throughput.

users in the regional areas when we conducted measurements. Hence we do not observe the significant effect of network load. In comparison, recall that we have observed noticeable disparity in the TCP throughput across different corridors. This suggests that the effect of location on TCP throughput is more prominent, as compared to time-of-day. To further demonstrate this, we evaluate the disparity of TCP throughput conditions between different corridors and between different times of day (for the same corridor) using the popular inequality analysis metric Gini index [14]. Recall that the Gini index scales between 0 and 1, wherein 0 indicates that all measurements are equal and 1 means maximal inequality. We calculate the *spatial* Gini index using the mean throughput values measured from different corridors. For each corridor, we also calculate the *temporal* Gini index over the mean throughput data measured at different times of day. Fig. 4 compares the spatial Gini index against the mean temporal Gini index, which is averaged over the temporal results from all three corridors. The graph clearly confirms that the disparity of TCP throughput between different times of day is significantly lower as compared to that between corridors.

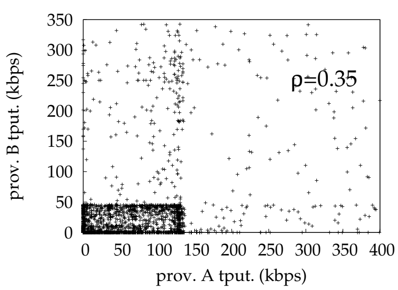

Fig. 5: Scatter plot of TCP throughput between providers.

## B. Provider and Technology Diversity

Our results also indicate the technology and provider diversities in TCP throughput. First, Table I shows that TCP throughput observed from different providers are different. Provider A, which runs on CDMA 1x/EVDO results in higher TCP throughput as compared to provider B running on UMTS/GPRS. Further, Fig. 2 suggests that the TCP throughput experienced from two different networks at the same location is not correlated. That is, if the throughput from provider A is very high at location X, then it does not necessarily imply that the throughput from provider B at this location will also be high. To better demonstrate this, Fig. 5 shows the scatter plot of TCP throughput between provider

TABLE II: The comparison of blackout statistics.

| | # of blackouts | | | avg. blackout length (m) | | |
|---|---|---|---|---|---|---|
| | A | B | ***Comb.*** | A | B | ***Comb.*** |
| Mel-Gee | 1 | 1 | ***0*** | 1000 | 1000 | ***0*** |
| Mel-Bal | 2 | 7 | ***1*** | 500 | 710 | ***500*** |
| Mel-Ben | 10 | 17 | ***4*** | 1350 | 2550 | ***1000*** |

A and B, using data at each 500m location from all three corridors. The corresponding Pearson's correlation coefficient $\rho$ is found only 0.35, which shows minimal correlations. This finding can be attributed to the difference in base stations layouts, configurations and service load conditions between providers. Thirdly, Fig. 2 suggests the service from provider A clearly suffers less blackouts along a route, as compared to that from provider B. We summarize the average statistics of blackouts of both provider A and B in Table II. In particular, we show the average number of blackouts and the average length of a blackout for each corridor. We see that provider B clearly suffers significantly more blackouts along all corridors. However, we also observe that the blackouts from provider A and B are not correlated with each other. It is evident in Fig. 2 that both networks rarely experience blackout simultaneously at the same location. Table II shows that by considering both providers together (shown in "combined") the effect of blackouts can be significantly alleviated. The above observations are consistent with the findings reported in [7]. This highlights that a mobile device or even an on-board communication network may achieve better reliability and throughput if it can support subscribing to multiple WWAN links concurrently. To address the inter-operator handover issue when shifting traffic among WWAN networks, it may be possible to use the existing techniques proposed for end-host multi-homing [7], [15].

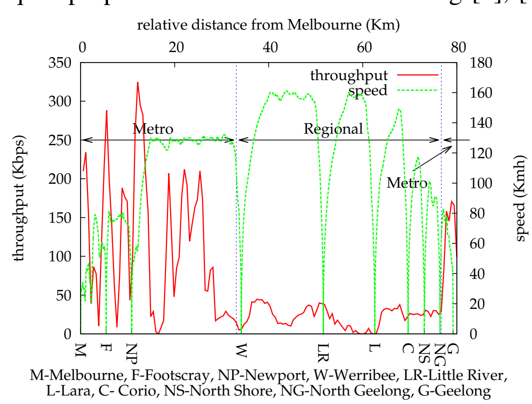

Fig. 6: Average TCP throughput and speed along the Melbourne-Geelong corridor.

## C. Effect of Train Speed on WWAN Bandwidth

Now, we explore the effect of the train speed on the TCP throughput. Recall that, we acquired the train velocity data from the trip logs provided by the train operator. Fig. 6 plots the average TCP throughput for providers B and the corresponding train speed at different locations along the Melbourne-Geelong corridor. Note that we have similar observations for provider A and other corridors. We can see that the train speed remains relatively low in the metropolitan areas, i.e. the section between the Melbourne and Werribee, and the section between North Geelong and Geelong. This is due to the reason that the train speed is restricted to comply with the safety and noise regulations applied in those areas. As a train travels out of the city areas, it can increase its speed until reaching the maximum 160 Kmh.

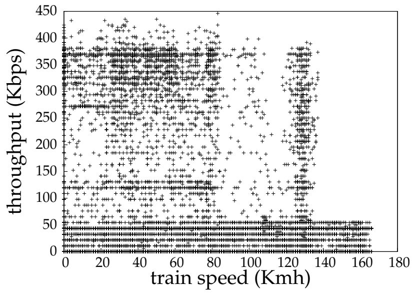


Fig. 7: Scatter plot between TCP throughput and speed over the Melbourne-Geelong corridor.

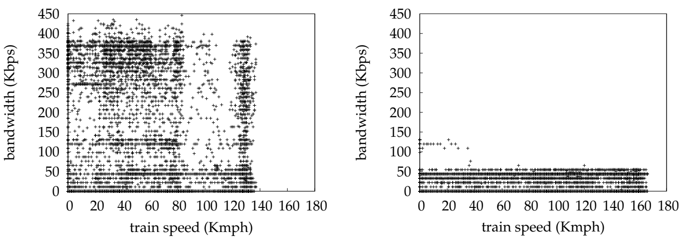


(a) Metropolitan area (Melbourne-Werribee & North Geelong-Geelong)

(b) Regional area (Werribee-North Geelong)

Fig. 8: Scatter plots of TCP throughput and Speed for metro and regional areas for the Melbourne-Geelong corridor.

To understand the correlation between the train speed and experienced TCP throughput, we plot the throughput-speed scatter plot for provider B for the Melbourne-Geelong corridor in Fig. 7. At first glance, the figure suggests that WWAN cannot support high TCP throughput under high train speed (over 130 Kmh). However, our further analysis revealed that this is merely the effect of the geographic network layouts of WWAN networks. Recall that the train speed in metropolitan areas is restricted. However, we have shown in Section IV-A that only the metropolitan areas are covered by the 3G networks with high bandwidth. Conversely, when a train is able to achieve top speed in the regional areas, a user can only experience 2G services. Hence the TCP throughput is consistently low. Fig. 8(a) and 8(b) plot the scatter plots using only the portion of data points observed in the metropolitan and regional areas, respectively. We can see that within the metropolitan area (Fig. 8(a)), the TCP throughput can still achieve over 300 Kbps under maximum allowed train speed (130 Kmh). In the regional areas, Fig. 8(b) shows the TCP throughput primarily fluctuates between 0-50 Kbps regardless the train speed. These observations confirm our above reasoning. Finally, another observation from Fig. 8 is that even under the same speed, the throughput can vary. This effect can be better shown in Fig. 6. For example, between the Newport and Werribee stations, observe that the TCP throughput is fluctuating significantly, while the train speed remains consistent at 130 Kmh. The above observations suggest that the train speed does not have a deterministic effect on the WWAN performance.

## V. Conclusion

We have presented measurement results for mobile broadband access from high-speed regional trains, which travel through both city and regional areas. We find that compared to city commuters, regional commuters are expected to experience significantly lower data rates and more frequent service disruptions. While operators are unlikely to spend a fortune for upgrading their networks for regional areas with lower population and commuters, it may be worth exploring policies and technologies that can enhance data rates and reliability in regional areas by seamlessly combining services from multiple providers. Our data shows that indeed the blackouts from different providers are not highly correlated and that use of just one more provider as a backup would significantly reduce the frequency and length of service blackouts in regional areas. While these initial results are encouraging, more measurements in different parts of the world should be carried out before making a general conclusion.


## Acknowledgements

We thank V/Line Passenger Pty. Ltd. for their assistance with our measurements and providing trip log data for analysis.


## References


[1] K. Jang, M. Han, S. Cho, H.-K. Ryu, J. Lee, Y. Lee, and S. B. Moon, "3G and 3.5G Wireless Network Performance Measured from Moving Cars and High-speed Trains," in *Proc. of MICNET*, Beijing, China, 2009.

[2] J. Ormont, J. Walker, S. Banerjee, A. Sridharan, M. Seshadri, and S. Machiraju, "A City-Wide Vehicular Infrastructure for Wide-area Wireless Experimentation," in *Proc. of WinTech*, San Francisco, 08.

[3] F. P. Tso, J. Teng, W. Jia, and D. Xuan, "Mobility: a double-edged sword for hspa networks: a large-scale test on hong kong mobile hspa networks," in *MobiHoc 2010*, 2010.

[4] F. P. Tso, L. Zhang, J. Tengy, W. Jia, D. Xuany, and F. Zhang, "An Empirical Evaluation on the Performance of Mobile HSPA Networks," *Technical Report, Dept of CS, City University of Hong Kong*, Jul. 2009.

[5] J. Yao, S. S. Kanhere, and M. Hassan, "An Empirical Study of Bandwidth Predictability in Mobile Computing," in *Proc. of ACM WinTech 2008*, San Francisco, CA, USA, Sep. 2008.

[6] J. Derksen, R. Jansen, M. Maijala, and E. Westerberg, "HSDPA Performance and Evolution," *Ericsson Review*, no. 03, pp. 117–120, 2006.

[7] P. Rodriguez, R. Chakravorty, I. Pratt, and S. Banerjee, "MARS: A Commuter Router Infrastructure for the Mobile Internet," in *Proc. of ACM MobiSys 2004*, Boston, Massachusetts, USA, Jun. 2004.

[8] Flashget. [Online]. Available: http://www.flashget.com

[9] Wireshark. [Online]. Available: http://www.wireshark.org

[10] V/Locity Trains. [Online]. Available: http://www.vline.com.au/about/ourfleet/vlocity.html

[11] Vodafone Australia. [Online]. Available: http://www.vodafone.com.au

[12] J. Yao, S. S. Kanhere, and M. Hassan, "Geo-intelligent Traffic Scheduling for Multi-Homed On-Board Networks," in *Proc. of ACM MobiArch09*, Krakow, Polland, Jun. 2009.

[13] Root Wireless. [Online]. Available: http://rootwireless.com/

[14] C. Gini, "Measurement of Inequality of Incomes," *The Economic Journal (Blackwell Publishing)*, 1921.

[15] N. Thompson, G. He, and H. Luo, "Flow Scheduling for End-host Multihoming," in *Proc. of Infocom 06*, Barcelona, Spain, 2006.