Concise Paper: TCP Performance over Mobile Networks in High-speed Mobility Scenarios

Qingyang Xiao*†, Ke Xu*†, Dan Wang‡, Li Li*†, Yifeng Zhong*†
* Tsinghua National Laboratory for Information Science and Technology (TNList)
† Department of Computer Science and Technology, Tsinghua University, Beijing, China
‡ Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China
Email: xqy13@mails.tsinghua.edu.cn, xuke@tsinghua.edu.cn, csdwang@comp.polyu.edu.hk,
ll-12@mails.tsinghua.edu.cn, zhong-yf11@mails.tsinghua.edu.cn

***Abstract*—Recently, the performance of mobile data networks has been evaluated from many aspects, e.g., TCP/IP protocols, comparison with WiFi or even satellite communication, under different movements within a metropolis area. Nevertheless, the result is still unknown in high-speed mobility scenarios and in a scale that crosses different metropolis and geographic areas. To fill in this blank, we carry out a comprehensive measurement study on the performance of mobile data networks under high-speed mobility, i.e., 300 km/h or above. Such speed is the current de facto standard of the China Railway High-speed (CRH) network, the largest commercial high-speed railway network in the world so far. We first present an overview on the TCP performance over LTE networks. We observe that decent throughput may exist under high-speed mobility. However, comparing to the stationary and driving (100 km/h) scenarios, the throughput and RTT not only are worse, but also have a large variance. We then take an in-depth investigation into two key factors affecting the performance, i.e., the wireless channel and handoff. We believe our study on these factors is useful not only for TCP, but also for other upper-layer protocols.**

I. Introduction

Mobile data networks are now expending their domains to every part of our life. In daily life, mobility is a basic requirement, including walking, driving, taking the subway and other public transport. Recent advancements, e.g., 3G/LTE, in mobile data networks try to provide high data rates and seamless mobility support for most of these scenarios.

Along with these development and deployment, researchers have investigated the performance of mobile data networks from many aspects. Tan et al. analyze the performance of UMTS networks in the stationary environment [1]. Tso et al. conduct a large-scale test on HSPA networks to study low-speed mobility impact [2]. Huang et al. carry out a measurement study to quantify the impact of mobile speed for a commercial LTE network [3]. Nevertheless, little is known about the performance of mobile data networks under high-speed mobility scenarios. With the increasing expansion of high-speed trains, especially in China, it is an important task to understand the performance of mobile data networks in such scenarios. High-speed mobility has unique features distinguishing it from other low-speed ones, requiring some existing topics to be reexamined. This motivates us to investigate the network performance and user experience under high-speed mobility.

In this paper, we carry out a measurement study on the performance of mobile data networks in high-speed mobility scenarios, i.e., 300 km/h or above. Such speed is the current de facto standard of the China Railway High-speed (CRH) network, the largest commercial high-speed rail (HSR) network in the world so far. In particular, we focus on the Beijing-Tianjin Intercity Railway (BTR). In our experiments, we study the performance over LTE networks of China Mobile, one of the largest carriers in the world. To obtain specific information of network usage, we develop a measurement tool on the Android platform, *MobiNet*. For comparison, we also conduct experiments for the stationary and driving (100 km/h) scenarios. To have a fair comparison, our driving follows the same path of BTR. We seek to uncover the speed limit for current cellular networks.

Note that, we use TCP as a working example to study the performance of LTE networks under high-speed mobility. Clearly, TCP is the most important protocol, which supports tons of application protocols. We believe understanding certain key parameters of TCP performance is of crucial importance for the design and optimization of upper-layer protocols. However, we strategically try to avoid protocol-specific conclusions. We believe our results are useful for TCP, as well as the design of other upper-layer protocols, for example, the Real-time Transport Protocol (RTP) over UDP, or the protocol used in Skype [4]. We consider that a more TCP-related study, which may include the impact of high-speed mobility on TCP handshake, TCP performance of different flow sizes, and TCP break-down operations, worth a separate study.

Based on our measurements, we first present an overview of the TCP performance over LTE networks under high-speed mobility. We show the throughput and RTT of TCP. Compared to the stationary and driving scenarios, the throughput is worse. Still, we observe that there can be a decent throughput of 1000 kbps. This is a fine throughput for almost all applications to date. However, there is a large variance in both the throughput and RTT. In many occasions, the throughput drops to zero. The RTT can easily become greater than 100 milliseconds, whereas in most driving scenarios, the RTT is only 10 milliseconds.

978-1-4799-6204-4/14 $31.00 © 2014 IEEE
DOI 10.1109/ICNP.2014.49

We then look into two key factors underlying the performance, i.e., the wireless channel and handoff. We find that for the wireless channel, the signal strength becomes worse. We conjecture that this is due to the Doppler effect. There are even coverage holes in 2G EDGE networks only. We further investigate the handoff, and find that there are both inter-system handoff and intra-system handoff. In addition, the handoff frequency is much greater. In particular, 45.1% of handoffs can be less than 10 seconds. This may seriously affect TCP retransmission, etc. As a matter of fact, we observe that there are periods when TCP transmission suspends, i.e., no packet sent or received.

To the best of our knowledge, this is the first public large-scale empirical study on network performance in high-speed mobility scenarios. Our empirical study results are useful for network operators to provide extensive support for high-speed mobility. It is also meaningful for network protocol and application designers to take high-speed mobility condition into consideration. Our work makes a first step to explore these issues focusing on high-speed mobility. We also add driving forces to this study from two aspects. One is the popularity of HSR, and the is the popularity of mobile broadband services (i.e., data access).

II. Related Work

There is great interest in the performance of mobile data networks. We categorize the related work into network performance, mobility and handoff.

Network performance: Tan et al. investigate the performance of UMTS networks in terms of throughput, latency, video, and their ability to provide service for different traffic classes under various loading conditions [1]. Huang et al. analyze general performance with various impacts, including signal strength, time of day and carriers [5], [6]. In addition, to measure network usage and performance, *MobiPerf* [7] tests network throughput and latency.

Mobility: Tso et al. present an empirical study on the performance of HSPA networks in Hong Kong via extensive field tests, including many possible mobile scenarios in urban areas, i.e., subways, trains, off-shore ferries and city buses [2]. A. Schulman et al. show that signal strength varies across locations and has a non-negligible impact on the performance [8]. P. Deshpande et al. show that over a wide geographic region and under vehicular mobility, networks exhibit very different throughput and coverage characteristics [9]. Measurements have been carried out at different mobile speed [3]: stationary, 35 mph and 70 mph. It shows that there is no major performance downgrade at different speed for LTE networks, and RTT remains stable at different speed with small variation. This greatly differs from our observations.

Handoff: Tu et al. analyze how mobility affects mobile data accounting and discovers the root cause of handoff [10]. Studies [11] and [12] analyze the handoff influence in detail based on the ns-2 simulator and prove handoff has a bad impact on network performance. K. Jang et al. examine the impact of high-speed mobility and show that mobile nodes have far worse performance than stationary nodes in the same condition [13]. Our work shows that variance (instability from the upper-layer point of view), rather than the sheer throughput, is the key problem. We investigate the root causes on wireless channel and handoff.

III. Experiment Methodology

In this section, we present our methodology for measuring mobile network performance over 4G LTE networks. Our dataset is based on a series of network tests at different mobile speed.

A. Measurement Tool

To conduct an in-depth study and collect packet traces, we develop an Android application named *MobiNet* for measuring cellular network performance. *MobiNet* achieves a series of network tests on throughput and latency, as well as other useful network metrics. *MobiNet* also records the variation of signal strength, the information of base stations, fine locations based on GPS, mobile speed, data connection state, and handoff occurrence. We trace communication route from clients to local servers and record the ping latency of each hop.

Our test algorithm is based on *iperf* [14], a famous tool to measure the maximum TCP bandwidth. Our work improves the packet-sending method in order to control the TCP flows by limiting the transmission rate and TCP window size. In our experiments, we generate four types of traffic: TCP long-lived traffic, TCP intermittent traffic with different flow sizes, concurrent TCP traffic, and UDP traffic with constant bit-rate.

Further, we have shared *MobiNet* to the public[1].

B. Measurement Setup

To explore the impact of high-speed mobility, we carry out a comprehensive measurement study on TCP performance over LTE networks using *MobiNet*. We conduct three types of evaluations: downlink only, uplink only, as well as simultaneous downlink and uplink. At the same time, packet traces are collected from both sides of the client and server. We take throughput and latency as our main metrics. Through a fine-grained packet analysis, we calculate the *throughput* and *RTT*. In addition to packet traces, we collect real-time data connection state, network type, and the variation of signal strength. We also record the mobile speed, the information of base stations, cell identifier (CID), and location coordinate based on GPS (i.e., latitude and longitude).

During the test, we let smartphones communicate with the benchmark sever directly via TCP/UDP connections. Our deployed server can provide a sufficient bandwidth of 100 Mbps. Thus we ensure the server bandwidth is not the bottleneck in our tests. In the experiments, we use four popular Android phone models as main devices, including Samsung Galaxy Note 3, Samsung Galaxy S4, Google Nexus 4, and HTC Desire, with the same runtime environment. During the experiments on trains, all the devices are set on the table close

[1]Our measurement tool is available at: http://www.thucsnet.org/projects.html

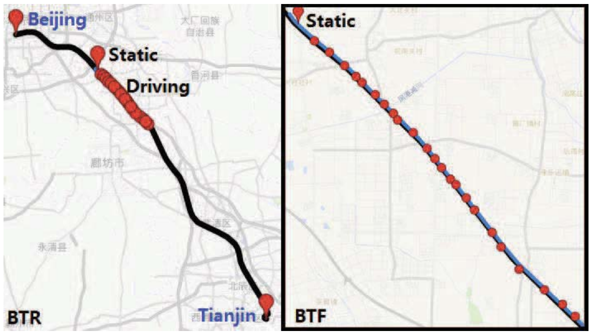

Fig. 1. Measurement traces. Left: BTR; Right: the overlapped area of BTR and BTF.

TABLE I
EXPERIMENT SCENARIOS

	Stationary	Driving	HSR
Test location	Toll station	**BTF**	**BTR**
Mileage (km)	-	20	115
Duration (min)	5	12	33
Speed (km/h)	0	100	300
Test times	6	6	42

to the window in order to ensure an identical test environment. We also take tests with 3G/LTE USB modem on laptops to verify the same experiment results can be obtained whether using smartphones or laptops.

C. Dataset

We perform experiments in three scenarios: stationary, driving, and high-speed mobility. We use Beijing-Tianjin Intercity Railway (**BTR**) as our primary testbed, which connects two largest metropolitan areas in China, Beijing and Tianjin with a distance of 115 km, as shown in Fig. 1. The total time for the high-speed train to complete this journey is 33 minutes. Besides the starting and ending phase, the speed of BTR reaches 300 km/h.

We also conduct measurement by driving a car. To have a fair comparison, we choose Beijing-Tianjin Freeway (**BTF**). This freeway is in parallel with the BTR line for a subset area, as shown in Fig. 1. In this period, we observe that the same base stations are used, which is called an overlapped area of BTR and BTF. We label the base stations along BTR and BTF, and show the location where we conduct the measurements in the stationary scenario. Thus, we believe the network environment is the same. A more complete profile of our experiments is shown in Table I. We conduct the measurements for seven months, from December 2013 to June 2014 and collect 49.25 GB packet traces.

IV. AN OVERVIEW ON TCP PERFORMANCE

We first give an overview of the TCP performance under high-speed mobility. For comparison, we also plot the stationary and driving (100 km/h) scenarios. We study two important metrics, the throughput and RTT. Table I shows the locations and ranges where we extract the data of the stationary, driving and high-speed train.

We plot the results of TCP throughput in Fig. 2, where Fig. 2 (a) (b) (c) show the stationary, driving and high-speed mobility cases respectively. The vertical dotted lines represent where a handoff occurs. Clearly, we can see that the throughput of the stationary case is the best, whereas under high-speed mobility, the throughput is much worse. TCP downlink throughput on HSR is only 575.8 kbps on average. For comparison, we observe a throughput of 1663.4 kbps in the stationary scenario and 1509.5 kbps in the driving scenario. We can also see that under high-speed mobility, the throughput can reach 1000-1500 kbps from time to time. As a matter of fact, such throughput, if stable, can actually satisfy most applications to date. Yet, there is a large variation in the throughput. Most notably, we see that there are many cases that the throughput drops to nearly zero. Such variation, if not understood or handled appropriately, may significantly affect the performance of upper-layer protocols.

We next show the RTT under high-speed mobility. In Fig. 3, we show the CDF of RTT. We see that 11.4% of RTT can be as long as 0.4 seconds. In such scale, the RTT can be beyond the tolerance of TCP. For comparison, in Fig. 3(b), 99.6% of the RTT in the stationary scenario is less than 30 milliseconds and 97.3% of the RTT in the driving scenario is less than 30 milliseconds.

We next show the RTT jitter in Fig. 4. We see that for stationary and driving at the speed of 100 km/h, for 90% or above, the RTT is only 7.6 milliseconds, yet for high-speed train, 56.6% of the RTT is greater than 100 milliseconds. In addition, the RTT variance is 1636 times greater than the other two scenarios. Note that an average RTT of 200 milliseconds may still be tolerable (we measure a connection between Beijing and New York can be 214.9 milliseconds), yet the large RTT jitter may seriously affect the performance.

Our conclusion is that the variance, rather than the sheer throughput or RTT is damaging for the wireless data network performance under high-speed mobility. In this paper, we are not trying to provide solutions, rather, we are looking into the root causes of such performance. We discuss wireless channel and handoff respectively.

V. ROOT CAUSE I: WIRELESS CHANNEL

We first present the wireless signal strength, which is the underlying support for wireless communication quality. There are many metrics for signal strength, such as RSSI, Ec/Io, SNR, RSRP (Reference Signal Receiving Power), and RSRQ (Reference Signal Receiving Quality). Each of these metrics is used for different scenarios, in particular, RSRP is used primarily for LTE networks [15]. In our work, we focus on LTE services; as such, we use RSRP as our major signal strength measurement metric. There are three levels in RSRP, Good (-85 dBm), Average (-95 dBm), and Weak (-105 dBm).

Fig. 5 shows the signal condition of the BTR line. Fig. 5(a) shows an experiment from 6:50 am to 7:23 am and Fig. 5(b) shows an experiment from 8:31 am to 9:04 am, both from Beijing to Tianjin. The three red lines in the figures show the Good, Average, Weak indices of RSRP.

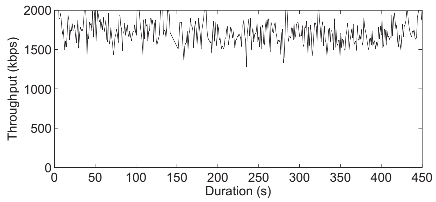

(a) Stationary

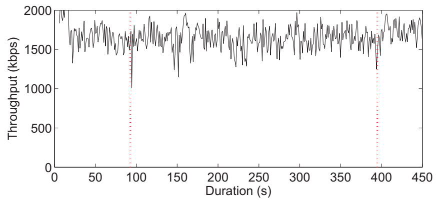

(b) Driving at the speed of 100 km/h

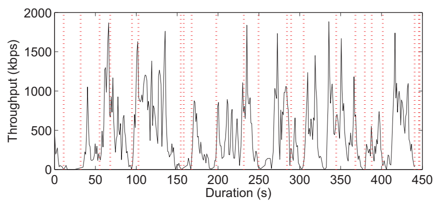

(c) HSR at the speed of 300 km/h

Fig. 2. TCP downlink throughput comparison with different speed.

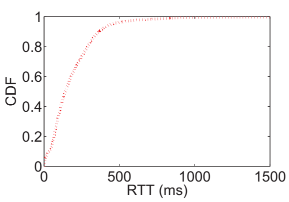

(a) RTT of HSR

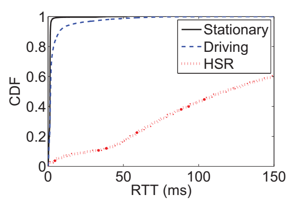

(b) RTT comparison

Fig. 3. RTT comparison with different speed.

First, we can see that the results of the two experiments are consistent. As a matter of fact, we have seen similar effects in the other experiments as well. Second, we observe coverage holes. Most notably, we observe that there is no 4G LTE coverage between the 462th second and the 591th second in the BTR line. This translates to a coverage hole of approximately 2.15 minutes. Third, we can see that the signal strength is primarily between Good and Average. This assures that data transmission is possible. And only 3 of 17 minutes are within the Good level. Although the general signal quality is acceptable, signal strength is dynamically changing. Since the

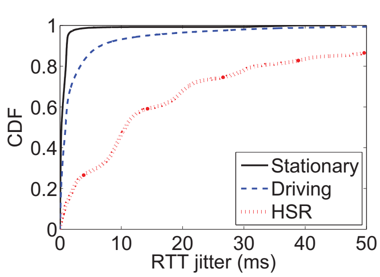

Fig. 4. RTT jitter comparison with different speed.

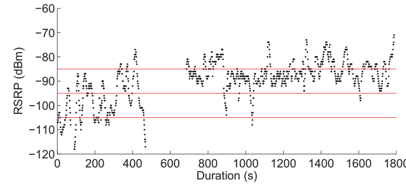

(a) LTE case 1: 6:50 am-7:23 am

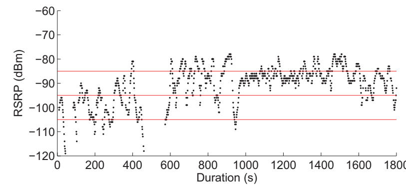

(b) LTE case 2: 8:31 am-9:04 am

Fig. 5. An example of signal condition under high-speed mobility.

experiments at different time are consistent, in the remaining part of the paper, we use the data of case 1, 6:50 am - 7:23 am.

We next compare the signal strength between the high-speed train and driving scenarios. In Fig. 6, we plot the signal strength between the 700th second and the 1050th second on HSR. We choose these two time-spots because the driving road is perfectly in parallel with the high-speed rail, as shown in Fig. 1. In addition, we can also see from Fig. 5 (a) that there are the locations with the consistent signal.

Fig. 6(a) shows the signal strength of high-speed train and Fig. 6(b) shows the signal strength during driving. Note that since the speed of high-speed train is three times faster than that of a car, the total duration of Fig. 6(a) and Fig. 6(b) is 350 seconds and 1200 seconds, respectively. We can make two observations from Fig. 6. First, the signal strength in driving scenarios is mostly above Good, usually between -80 dBm and -75 dBm, which is much better than high-speed train. Second, the variance in driving is also consistently smaller than that of the high-speed train. These two observations show that the high-speed mobility does negatively affect the performance on signal strength.

After understanding the signal strength, we then investigate the impact of signal strength on throughput under high-speed mobility. Fig. 7 illustrates the correlation between TCP downlink throughput and signal strength on HSR, where the three lines represent the Good, Average, Weak signal conditions. We find that a Good signal condition leads to good TCP throughput and a Weak signal condition leads to much worse TCP throughput. In addition, we observe that the throughput in the Average and Good signal conditions are reasonably close, yet the Weak signal condition leads to a significantly worse performance. This shows that it is necessary to increase the signal strength to an Average condition.

VI. Root Cause II: Handoff

A good wireless signal strength alone does not guarantee good performance for the throughput of upper-layer protocols. We next consider another key factor, handoff.

Current mobile networks are based on the cellular architecture and each base station can only cover a certain range.

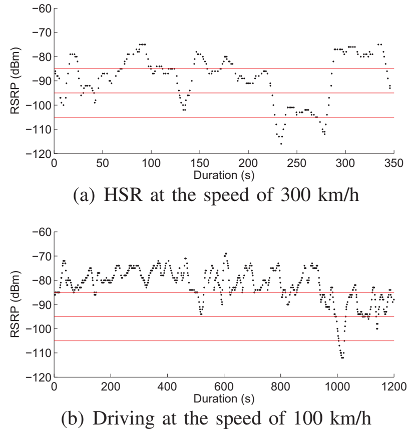

Fig. 6. Signal strength comparison at different speed.

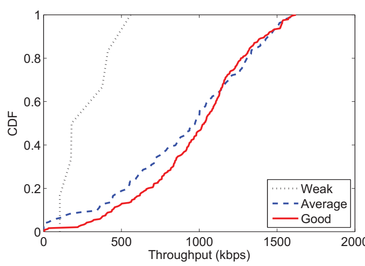

Fig. 7. Signal strength impact on TCP performance.

When a mobile device (MD) goes through the coverage and the signal strength is too weak, a handoff is triggered by the MD or the serving cell. The MD would select a better cell and switch connection from one base station to another.

With the increasing user mobility, handoff frequency will increase. Since the link is temporarily lost during handoffs, packets will be lost or delayed. If looking from the TCP point of view, either case may cause an increase in its retransmission timer and a decrease in its transmission window size, which finally leads to throughput reduction.

A. *Handoff Types*

In our experiments, we observe multiple types of handoffs. Intuitively, there are handoffs between base stations of the same type of network, e.g., 4G to 4G, 3G to 3G, and there are handoffs of different types of network, e.g., 4G to 3G, 3G to 2G. As a matter of fact, 3G networks alone have many different subtypes, e.g., UMTS, HSPA, HSPA+. Nevertheless, these 3G networks are using the same radio access network. Therefore, in this paper, we classify two types of handoffs, inter-system handoff and intra-system handoff. More specifically, for both inter and intra handoff, the MD will choose a new base station, yet for an inter-system handoff, the MD uses different radio access networks [10].

The total number of handoffs we observe is 71, including 65 intra-system handoffs and 6 inter-system handoffs. This translates to a handoff in every 25.4 seconds on average. For comparison, driving at the speed of 100 km/h, the handoff occurs in every 248.2 seconds on average.

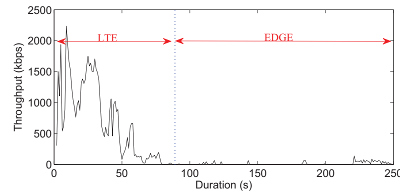

Fig. 8. Inter-hanoff impact on TCP throughput.

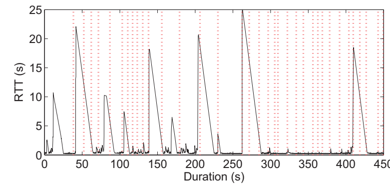

Fig. 9. Intra-system handoff impact on RTT.

In our experiments, all inter-system handoffs occur during the coverage hole, i.e., from the 462th second to the 591th second. As discussed in the previous section, there is no 4G LTE network in this area, but 2G network service only. We can see from Fig. 8 that the throughput after the handoff drops significantly. At the 89th second, MD is triggered by inter-system handoff from LTE to EDGE with TCP throughput dropping to 20-30 kbps.

We believe that for high-speed train, inter-system handoff will exist for certain time because infrastructure deployment may take long period of time. Especially, for regions, e.g., China, where different provinces are responsible to their own network development, coordination to develop a uniform network infrastructure is even more difficult. Nevertheless, with the improvement of infrastructure, such coverage hole can be ultimately fixed. As such, we focus on intra-system handoff in the remaining part of this section.

B. *Intra-system Handoff*

We present intra-system handoff between the 750th second to the 1200th second. We choose this period of time because this is the largest continuous period that the train is moving consistently at the speed of 300 km/h.

Fig. 9 shows the RTT during this 450 seconds. The red lines show the time when intra-system handoffs occur. In total, we observe 35 handoffs during this 450 seconds. The RTT varies from 12 milliseconds to 25 seconds. Such long RTT will easily affect TCP retransmission mechanisms and RTO setting.

We next look into the handoff frequency. We define *handoff interval* to describe the interval between two adjacent handoffs. Recall that the average handoff time is 25.4 seconds. Fig. 10 shows the CDF of handoff interval. We see that 45.1% of handoff intervals are within 10 seconds. If we think that the RTT can be as long as 25 seconds, the TCP retransmission may be seriously affected.

As a matter of fact, during handoffs, the MDs change the current connection, and there is an interval when no packets are received or sent, namely a *suspension duration*. Fig. 11 shows the CDF of suspension duration. We see that 27.8% has a suspension duration of more than 10 seconds. This also

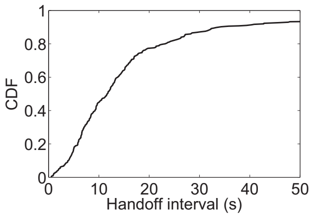

Fig. 10. CDF of handoff interval under high-speed mobility.

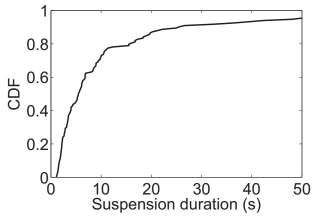

Fig. 11. CDF of suspension duration under high-speed mobility.

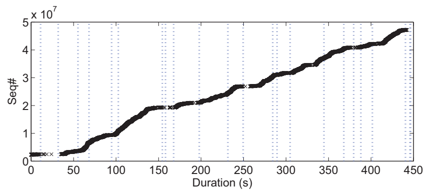

(a) Sequence number

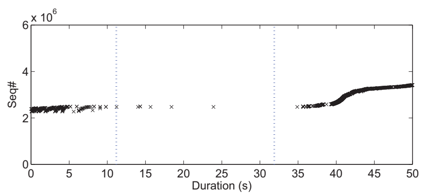

(b) Zoom in within 50 seconds

Fig. 12. Handoff frequency impact on data suspension.

explains the huge variation and barely no throughput as shown in Fig. 2.

With the increase of handoff frequency, TCP data suspension becomes more serious. In Fig. 12, we plot the increasing pattern of sequence numbers under high-speed mobility. Fig. 12(a) shows the results during a period of 450 seconds. TCP packet-sending becomes discontinuous with many blank intervals, especially for the one from the 9.8th second to the 34.9th second, there are few packets received or sent, as shown in Fig. 12(b). We observe suspension duration sometimes can exceed 25 seconds, even 10 minutes in other traces. And TCP data suspension time (71.08 seconds) is 15.8% of the total 450 seconds.

VII. Conclusion

In this paper, we present a measurement study on the performance of mobile data networks under high-speed mobility. For comparison, we also measure the stationary and driving scenarios in the same network environment. We observe that the TCP throughput and RTT are worse than that of the stationary and driving scenarios; yet a reasonable throughput and RTT is possible. We find that the real killer is the large variance. We investigate the root causes for such performance, namely, the wireless channel and handoff. The signal strength is worse and coverage holes may exist. We also show frequent handoffs can lead to long RTT and in many occasions suspend the transmission for a long period of time. We believe our main observations provide valuable insights into high-speed mobility scenarios. In our future work, we plan to further investigate other mobile data networks, such as HSPA+ and EVDO. Besides Beijing-Tianjin Intercity Railway, we also plan to study other railway lines. We believe these understanding can help design better upper-layer protocols in the future.

Acknowledgment

This work has been supported in part by NSFC Project (61170292, 61161140454, 61272464), National Science and Technology Major Project (2012ZX03005001), 973 Project of China (2012CB315803), 863 Project of China (2013AA013302), EU MARIE CURIE ACTIONS EVANS (PIRSES-GA-2013-610524), RGC/GRF PolyU 5264/13E, HK PolyU G-UB72, 1-ZVC2, and 1-ZE9B.

We would like to appreciate all the anonymous reviewers and Dr. Chunyi Peng for shepherding the paper. We also thank all the volunteers who help us collect data traces on HSR.

References

[1] W. L. Tan, F. Lam, and W. C. Lau, "An empirical study on 3G network capacity and performance," in *Proc. IEEE INFOCOM'07, Anchorage, USA*. IEEE, 2007, pp. 1514–1522.

[2] 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 *Proc. ACM MOBIHOC'10, Chicago, USA*. ACM, 2010, pp. 81–90.

[3] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4G LTE networks," in *Proc. ACM MOBISYS'12, Low Wood Bay, Lake District, United Kingdom*. ACM, 2012, pp. 225–238.

[4] X. Zhang, Y. Xu, H. Hu, Y. Liu, Z. Guo, and Y. Wang, "Profiling skype video calls: Rate control and video quality," in *Proc. IEEE INFOCOM'12, Orlando, USA*. IEEE, 2012, pp. 621–629.

[5] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl, "Anatomizing application performance differences on smartphones," in *Proc. ACM MOBISYS'10, San Francisco, USA*. ACM, 2010, pp. 165–178.

[6] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck, "An in-depth study of LTE: Effect of network protocol and application behavior on performance," in *Proc. ACM SIGCOMM'13, Hong Kong, China*. ACM, 2013, pp. 363–374.

[7] "MobiPerf website," Available: http://www.mobiperf.com/.

[8] A. Schulman, V. Navda, R. Ramjee, N. Spring, P. Deshpande, C. Grunewald, K. Jain, and V. N. Padmanabhan, "Bartendr: a practical approach to energy-aware cellular data scheduling," in *Proc. ACM MOBICOM'10, Chicago, USA*. ACM, 2010, pp. 85–96.

[9] P. Deshpande, X. Hou, and S. R. Das, "Performance comparison of 3G and metro-scale WiFi for vehicular network access," in *Proc. ACM IMC'10, Melbourne, Australia*. ACM, 2010, pp. 301–307.

[10] G.-H. Tu, C. Peng, C.-Y. Li, X. Ma, H. Wang, T. Wang, and S. Lu, "Accounting for roaming users on mobile data access: issues and root causes," in *Proc. ACM MOBISYS'13, Taipei, Taiwan*. ACM, 2013, pp. 305–318.

[11] D. Pacifico, M. Pacifico, C. Fischione, H. Hjalrmasson, and K. H. Johansson, "Improving TCP performance during the intra LTE handover," in *Proc. IEEE GLOBECOM'09, Honolulu, USA*. IEEE, 2009, pp. 1–8.

[12] C. Paasch, G. Detal, F. Duchene, C. Raiciu, and O. Bonaventure, "Exploring mobile/WiFi handover with multipath TCP," in *Proc. ACM SIGCOMM Workshop on Cellular Networks: Operations, Challenges, and Future Design (CELLNET'12)*. ACM, 2012, pp. 31–36.

[13] 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. ACM Workshop on Mobile Internet through Cellular Networks (MICNET'09)*. ACM, 2009, pp. 19–24.

[14] "Iperf website," Available: http://iperf.fr/.

[15] T. Jansen, I. Balan, J. Turk, I. Moerman, and T. Kurner, "Handover parameter optimization in LTE self-organizing networks," in *Proc. IEEE VTC'10 Fall, Ottawa, Canada*. IEEE, 2010, pp. 1–5.