

# Accurate Available Bandwidth Estimation Robust Against Traffic Differentiation in Operational MVNO Networks

Takashi Oshiba  
NEC Corporation, Japan

**Abstract**—We uncover the reality of traffic differentiation in mobile virtual network operator (MVNO) networks and shed new light on its negative impact on the accuracy of available bandwidth (i.e., the unused capacity of an end-to-end path) estimation methods that use a probing packet train (i.e., a set of multiple probing packets). To the best of our knowledge, we report, for the first time, that multiple operational LTE MVNO networks in the Asian region are very likely to differentiate traffic according to the time of day. We find, also for the first time, that when a sender transmits a packet train to a receiver during a period of traffic differentiation, the packet loss rate of the packet train jumps up, thus degrading the accuracy of available bandwidth estimation as the packet train is corrupted by the high packet loss rate.

As a countermeasure, we propose a method called PathQuick4 for accurately estimating the available bandwidth while remaining robust against traffic differentiation. The key idea is that we dynamically judge whether traffic differentiation occurs through in-line measurement of packet loss characteristics by using a single packet train itself, without transmitting any extra probing packets. From this judgement, we then choose either packet-train- or packet-pair-based bandwidth estimation as appropriate.

An experimental evaluation over three operational Japanese LTE MVNO networks showed that the overall estimation error of PathQuick4 across the three MVNOs during periods of traffic differentiation was close to an order of magnitude less than the error of a conventional method.

**Keywords**—available bandwidth; traffic differentiation; MVNO

## I. INTRODUCTION

The rise of mobile virtual network operators (MVNOs) has been an emerging trend over the last few years in the mobile communication market. About 1,000 MVNOs are now in operation worldwide [1]. The main market regions are Europe, Asia-Pacific, and North America [1]. MVNOs have recently gained non-negligible market share. For example, Germany, the UK, and Japan have 23, 14, and 16 million MVNO subscribers, respectively. Likewise, the market shares of MVNO subscriptions in the UK and Japan are 15% [2] and 10% [3], respectively. One prominent driver of this market share growth is the low-cost plans (i.e., discount prices) for mobile service via MVNOs [1][2].

In general, MVNOs *lease* access to the infrastructure of a host mobile network operator (MNO), thus avoiding the high costs of building their own base stations and licensing spectrum, enabling MVNOs to provide lower-cost plans than those of their host MNOs. Of course, MVNOs must pay MNOs to lease access to the MNO's networks, i.e., MVNOs *buy* network bandwidth from MNOs. For example, the Ministry of Internal Affairs and Communications (MIC), the regulator of the Japanese telecom industry, discloses that the lease price for each 10 Mbps/month of bandwidth ranges from 6,000 to 9,000 USD [4]. Large-scale investment by MVNOs in bandwidth leases is difficult, because typical MVNOs are relatively smaller companies than MNOs. Given the low margins and highly competitive nature of the mobile market [1], MVNOs try to fully utilize their scarce leased

bandwidth in order to maximize their subscribers' quality of experience (QoE) with minimum investment [5]. Nowadays, however, since ordinary MVNO subscribers use bandwidth-intensive smartphone apps, such as on-demand video streaming (YouTube), live video streaming (Facebook Live), and video chat (Skype), in their daily life, MVNO networks are often congested. Thus, ensuring QoE in an MVNO network remains challenging. This is a possible motive for MVNOs to *differentiate traffic* (as defined precisely in §III-A), because traffic differentiation can mitigate the congestion in an MVNO network. There is a long, ongoing debate about so-called *network neutrality* [6], which is tightly coupled with traffic differentiation. A very recent Federal Communications Commission (FCC) decision to discard network neutrality rules [7] (which had gone into effect in 2015) in the US may further encourage such traffic differentiation by MVNOs in the world.

A potential tool for ensuring QoE even under congested conditions is accurate measurement of available bandwidth (i.e., the physical capacity minus the bandwidth used during a certain time period). For example, in a video streaming system, the system could adapt the initial video bit rate to the currently available bandwidth *just before* video packet transmission. Combining this and adaptive bit-rate control [8] *after* video packet transmission starts would enable video streaming without playback interruption, thus ensuring QoE.

In this paper, we empirically show that MVNOs are very likely to differentiate traffic, and that such differentiation has a negative impact on the accuracy of available bandwidth estimation methods that use a probing packet train (i.e., a set of multiple probing packets). As a countermeasure, we propose a method called PathQuick4 for accurately estimating the available bandwidth despite traffic differentiation. PathQuick4 is the successor to our group's previous methods [9]–[12].

Note that this paper focuses on the downlink direction in mobile networks, because the traffic volume of a downlink can be more than 10 times greater than that of an uplink [13], and thus, downlink traffic is dominant in mobile networks.

The main contributions of this paper are fourfold:

- (1) **Traffic differentiation in MVNO networks (§III)**. To the best of our knowledge, we report, for the first time, that multiple operational LTE MVNO networks in the Asian region are very likely to differentiate traffic based on the time of day.
- (2) **Negative impact of traffic differentiation on accurate estimation of available bandwidth (§IV)**. Also for the first time, we found that during periods of traffic differentiation, the packet loss rates of packet trains for available bandwidth estimation jumped up, thus degrading the accuracy of estimation.
- (3) **Accurate estimation of available bandwidth (§V)**. We developed an available bandwidth estimation method that is accurate while remaining robust against traffic differentiation.

(4) **Experimental evaluation over three operational Japanese LTE MVNO networks (§VI).** We conducted, also for the first time, experimental evaluations of available bandwidth estimation over MVNO networks *in the wild*.

## II. RELATED WORK

Here, we discuss prior work on traffic differentiation detection (§II-A) and available bandwidth estimation (§II-B).

### A. Prior Work on Traffic Differentiation Detection

Although there has been significant effort towards detecting traffic differentiation, only a few of these prior work focus on MVNO networks.

#### 1) Wired and MNO Networks, and MVNO Networks Without Detection

Traffic differentiation detection methods were originally designed for wired networks in the last decade [14]–[21], and they have not yet been applied in detecting traffic differentiation in mobile networks. Although several recent studies [22][23] reported that they detected traffic differentiation in mobile networks, the networks under study were used not by MVNOs but by MNOs. Several other recent studies [24]–[26] conducted performance measurements in MVNO networks but did not detect traffic differentiation.

#### 2) MVNO Networks with Detection

Two recent studies [27][28] did report detecting traffic differentiation in MVNO networks; thus, these studies are the most closely related to our work. They differ from our work, however, in the following aspects: (1) All the networks in those studies were *light* MVNOs (See §III-B for the definitions of light and full MVNOs. Note that all the MVNOs such as Congstar in [27] and BlackWireless, H2O, and SimpleMobile in [28] are classified as light MVNOs in [25]). (2) The mobile systems in [27] were *EDGE* and *3G*, which are now relatively old technologies. (3) The networks in [27] and [28] were deployed in *Europe* (Germany) and *North America* (the US), respectively. (4) The traffic differentiation in [27] was based on the *type of protocol* (session initiation protocol, or SIP), while that in [28] was based on the *type of application* (YouTube, Netflix, etc.).

In contrast, to the best of our knowledge, this paper is the first to consider (1) *full* MVNOs (more complex network structures than light MVNOs) using (2) *4G LTE* mobile systems (overwhelmingly higher dynamic range of available bandwidth than with 3G) in (3) *Asia* (Japan) (diverse market regions would have diverse network characteristics) with (4) traffic differentiation based on the *time of day* (independent of protocol or application).

### B. Prior Work on Available Bandwidth Estimation

Since the early 2000s, there has been much work on end-to-end available bandwidth estimation techniques that actively send probing packet trains [29]. Only a few prior work, however, dealt with packet loss from packet trains.

#### 1) Representative Methods

The most cited available bandwidth methods in the literature include pathChirp [30] and Pathload [31]. Both, however, are vulnerable to packet loss. The former paper [30] clearly stated that “*pathChirp discards all packet trains with dropped packets*,” and thus, pathChirp cannot output any estimate when packet loss occurs. Another work [32] reported that the estimation accuracy of Pathload drastically deteriorates when the packet loss rate is only 2% (or more).

#### 2) Attempts at Packet Loss Tolerance

Only a few studies on available bandwidth estimation have tried to be tolerant of packet loss. Among these methods, Pathload-P [32] was evaluated only in a wired network with a

small private testbed. Similarly, the method in [33] was evaluated only in a wired network with the ns-2 simulator. While the estimation accuracies of these methods in a wireless setting are unknown, conventional wisdom [34] says that most methods designed for wired networks fail to estimate available bandwidth accurately in a wireless setting even when there is no packet loss. While SB [35] was evaluated in a Wi-Fi network with packet loss rates of up to 10%, its estimation accuracy drastically deteriorated when the packet loss rate approached 10%. As described later in §IV, we observed packet loss rates on the order of tens of percent during our experiments on traffic differentiation.

Thus, to the best of our knowledge, there is no report in the literature of an *accurate* available bandwidth estimation method robust against *high packet loss rates* in *wireless* networks (of course, neither in MVNO networks).

## III. TRAFFIC DIFFERENTIATION IN MVNO NETWORKS

In this section, we first classify the types of traffic differentiation (§III-A) and MVNOs (§III-B). Then, we empirically show that MVNOs are very likely to differentiate traffic according to the time of day (§III-C).

### A. Classification of Traffic Differentiation

This paper focuses on traffic differentiation (also referred to as bandwidth throttling, traffic discrimination, or rate limiting) due to *traffic shaping* or *traffic policing*. Traffic exceeding a pre-configured rate is *buffered* in traffic shaping, or *dropped* in traffic policing [22]. Since a traffic shaper has a finite buffer, packets are *dropped* when the buffer is full. Thus, a traffic policer can be implemented as a special case of a shaper by setting the buffer size to zero (or a few) packets [36].

Note that we do not consider traffic differentiation due to *blocking* (e.g., censorship) or *content modification* (e.g., transcoding).

### B. Classification of MVNOs

MVNOs are classified in one of two ways, as *light* or *full* [25][26]. On the one hand, light MVNOs almost entirely depend on their host MNO’s infrastructure, including the radio access network and IP core network, but they provide their own authentication servers. On the other hand, full MVNOs depend only on their host MNO’s radio access network and part of the IP core network. In an LTE mobile system, full MVNOs have their own packet data network gateways (P-GWs) and authentication servers [37], as shown in Fig. 1. This provides more operational freedom than light MVNOs have. Since full MVNOs can implement deeper, more customized services than can light MVNOs, most Japanese MVNOs are now full MVNOs. In contrast, all of the MVNOs that differentiate traffic in Europe and North America in [27] and [28] are light MVNOs. Since the network structure of a full MVNO is much more complex than that of a light MVNO, the network behavior differs from that of a light MVNO and is poorly understood. This is why we investigated the network behavior of full MVNOs. Note that both light and full MVNOs can deploy traffic differentiators (shapers or policers) in their own networks.

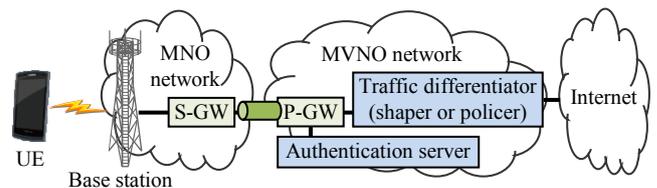


Fig. 1. Network architecture of MNO and MVNO. The network bandwidth (physical capacity) of the interconnect link between an MNO and an MVNO is specified by a contract between them.

### C. Traffic Differentiation in Japanese LTE MVNO Networks

We conducted experiments over three operational *Japanese 4G LTE* MVNO networks to examine whether they differentiate traffic. Here, we anonymize the MVNOs' names as A, B, and C for security and business reasons. Note, however, that all three are primary MVNOs in the Japanese market, and according to their official Web sites, all three have publicly announced that they are *full* MVNOs and depend on the same MNO network. We periodically collected downlink bulk TCP throughput at 10-minute intervals with one of the most well-known *speed test* Android app [38] over each of the three MVNO networks on several days from June to September, 2017.

Figure 2 shows a snapshot of a 1-day monitoring duration. We found that speeds suddenly dropped to a *particular, low, fixed value* during *three particular time periods*. We refer to these periods as *morning, noon, and evening throttling*. During morning and noon throttling, the speeds for all three MVNOs were around 1 Mbps from 7:30 to 8:30 and from 12:00 to 13:00, respectively. During evening throttling, the speeds were around 2 Mbps from 17:00 to 23:00 for MVNOs A and C, and from 17:00 to 22:00 for MVNO B. These *nice round numbers* at fixed low values (1 or 2 Mbps) and the frequent transitions on the hour (e.g., at 12:00, 17:00, etc.) strongly imply that the MVNO network administrators are differentiating traffic *intentionally* according to *time of day*. If a network administrator specifies traffic shaping (or policing) at 1 Mbps per user equipment (UE) device, then the maximum bulk throughput of each UE device can be limited to 1 Mbps or less. This prevents any device from saturating the MVNO's network bandwidth leased from an MNO, and enabling the network administrator to keep the network traffic under control.

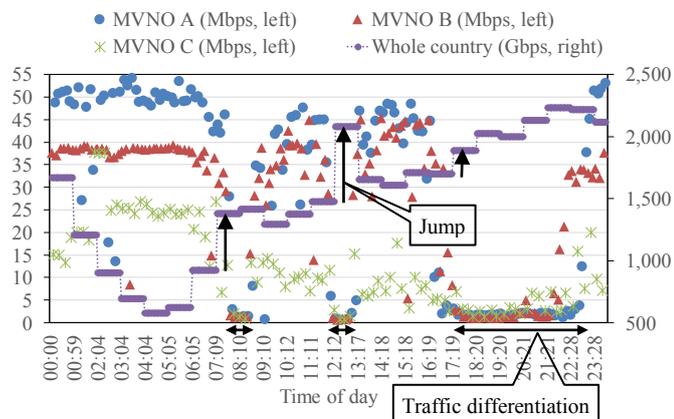


Fig. 2. Speed test results over three Japanese 4G LTE MVNO networks (left y-axis) for a 1-day duration, overlapped with the total downlink throughput among all Japanese MNOs (right y-axis).

We analyzed the data shown in Fig. 2 and found the following three points that support our deduction of traffic differentiation.

First, we also plotted data from MIC summing up the total downlink throughput for one month (June, 2017) among all Japanese MNOs [39], as shown in Fig. 2. We can interpret this data as the nationwide mobile traffic volume. Note that it includes the traffic volume of MVNOs. We found that the traffic volume jumped up when each of the three throttling time periods began. These jumps mean increases in network utilization, providing a motive for MVNOs to differentiate traffic.

Second, if there was no traffic differentiation, then the speed test results during the morning throttling (7:30–8:30) and from 8:30–12:00 should have been almost the same, because the traffic volumes during both time periods were almost the same (about 1,500 Gbps). In reality, however, the speed was around 1

Mbps during the morning throttling but more than 10 Mbps from 8:30–12:00. This indicates deliberate traffic differentiation.

Third, if there was no traffic differentiation, then during the evening throttling the speed test results should have had a decreasing trend, because the traffic volume had an increasing trend. Instead, the speed hovered at 2 Mbps. This also indicates artificial traffic differentiation.

To validate our assumption of traffic differentiation more directly, we are currently conducting experiments with a traffic differentiation detection tool called ShaperProbe [18], and the preliminary results are encouraging.

### IV. NEGATIVE IMPACT OF TRAFFIC DIFFERENTIATION ON AVAILABLE BANDWIDTH ESTIMATION ACCURACY

We conducted another experiments (with the same experimental setup described later in §VI-A) to compare the estimated available bandwidth obtained with our previous method, PathQuick3 [11], and the speed test results over the three MVNO networks. The results confirmed that traffic differentiation has a large impact on the accuracy of estimating the available bandwidth. Figure 3 shows the results for MVNO A. In [11], we confirmed that values estimated by PathQuick3 were close to speed test results over an LTE MNO network. During the non-throttled periods in Fig. 3, we observed similar results. During the three throttling periods, however, PathQuick3 clearly overestimated the available bandwidth, with an error of up to 56 Mbps.

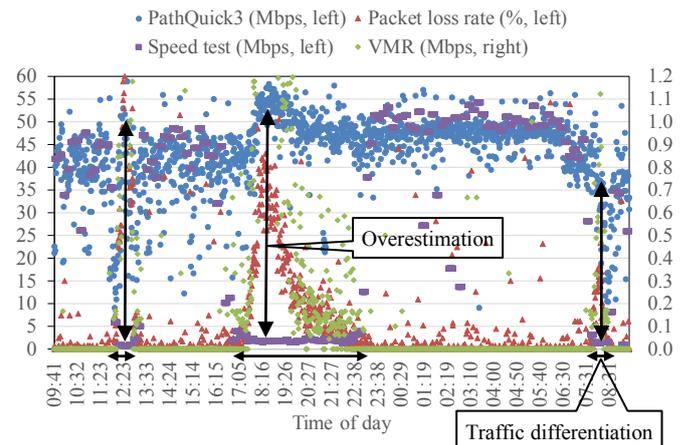


Fig. 3. Available bandwidth estimated by PathQuick3, speed test results, and packet loss rate of packet trains over the network of MVNO A (left y-axis). PathQuick3 clearly overestimated the available bandwidth during periods of traffic differentiation. The variance-to-mean ratio (VMR, see §V-A-2) is also shown (right y-axis).

We investigated the cause of overestimation and found that during the periods of traffic differentiation, the packet loss rate of packet trains jumped up, as shown in Fig. 3. We observed surprisingly high packet loss rates up to 80%. We believe that during these periods a packet train arrived within a short time (574 KB of data in 126 ms with our experimental setup) at a narrow pipe (1 or 2 Mbps), resulting in many dropped packets at a traffic shaper (or policer), as described in §III-A.

In the estimation algorithm of PathQuick3 (see [11] for details), the receiver of a packet train extracts *queuing delays* from it and uses curve fitting between the extracted and theoretical queuing delays by the nonlinear least-squares method. The problem with PathQuick3 is that queuing delays extracted at the receiver from a packet train *corrupted* by a high packet loss rate (i.e., with a much smaller number of packets than at the time the sender transmitted the packet train) confuse the curve fitting algorithm, resulting in overestimation.

## V. PROPOSAL OF PATHQUICK4

As a countermeasure, we propose a method called PathQuick4 for accurately estimating available bandwidth despite traffic differentiation. PathQuick3 and PathQuick4 use the same packet train structure (see Fig. 4). Each UDP packet is placed at an equal time interval, and the size of each packet increases linearly from the previous one as the packet sequence proceeds.

While the senders in PathQuick3 and PathQuick4 have the same functionality, the receivers differ. Figure 5 depicts a flowchart of the estimation algorithm used by PathQuick4 at a receiver. When the receiver receives a packet train, it judges whether traffic differentiation occurs by using inline measurement of packet loss characteristics, as described in §V-A. From that result, we choose either a packet-train- or packet-pair-based procedure for bandwidth estimation, with the latter procedure described in §V-B. This algorithm is simple but effective, as shown in §VI-B.

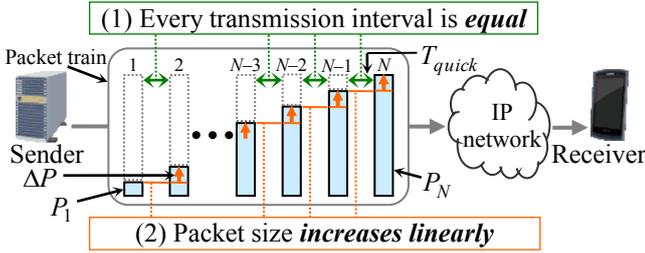


Fig. 4. Packet train structure of PathQuick4 (and PathQuick3).

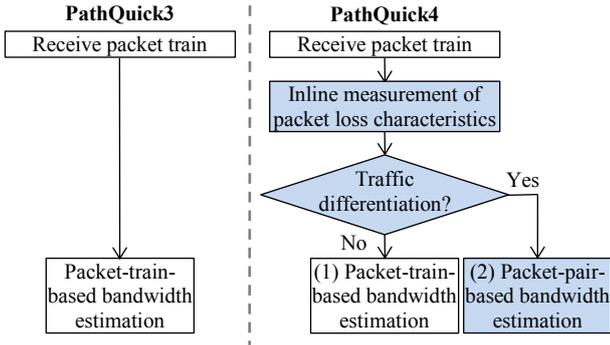


Fig. 5. Flowcharts of the estimation algorithms used by PathQuick3 and PathQuick4 (an extension of PathQuick3) at a receiver.

### A. Inline Measurement of Packet Loss Characteristics

Our idea is to leverage the correlation between traffic differentiation and the packet loss rates of packet trains, observed in Fig. 3, to detect traffic differentiation. To this end, we developed our inline method of measuring packet loss characteristics by using a single packet train already received, *without transmitting any extra packets*. Although it seems technically feasible to detect traffic differentiation by transmitting extra packet trains or using a traffic differentiation detection tool, such an approach would induce more network load and be unwelcome to cost-conscious MVNO subscribers. Note that we apply our inline measurement method with every packet train in order to dynamically capture the current status of traffic differentiation even when an MVNO changes its traffic differentiation approach.

We developed two different inline measurement methods, called PathQuick4-loss and PathQuick4-VMR.

#### 1) PathQuick4-loss

This method simply uses the packet loss rate of a packet train and judges that traffic differentiation exists if the rate exceeds a certain threshold.

#### 2) PathQuick4-VMR

While the PathQuick4-loss method works moderately well, we discovered more room to improve. Essentially, it simply counts the number of lost packets, and thus, information on *how* the packet losses occurred is dropped. In contrast, the PathQuick4-VMR method performs a deep check of a packet train's *packet loss behavior*. To explain the mechanism of this method, we define an integer variable, called the *length of continuous lost segment*,  $L_{cls}$ , as the number of continuous lost packets. For example, suppose that a sender transmits a packet train with 12 packets, so the packet numbers at the sender are  $\{1,2,3,\dots,12\}$ . Suppose also that the packet numbers of the lost packets at the receiver are  $\{2,3,4,6,8,10,11\}$ , meaning that the packets with packet numbers 1, 5, 7, 9, and 12 are successfully received. Then, there are 4 *continuous lost segments*:  $\{2,3,4\}$ ,  $\{6\}$ ,  $\{8\}$ , and  $\{10,11\}$ . Thus, the set of values of the length of continuous lost segment,  $L_{cls}$ , is  $\{L_{cls}\} = \{|\{2,3,4\}|, |\{6\}|, |\{8\}|, |\{10,11\}|\} = \{3,1,1,2\}$ , where the notation  $|\{\cdot\}|$  means the size of a set.

For observation we periodically transmitted 295 packets as a train in MVNO A. Figure 6 shows the observed packet loss behavior.

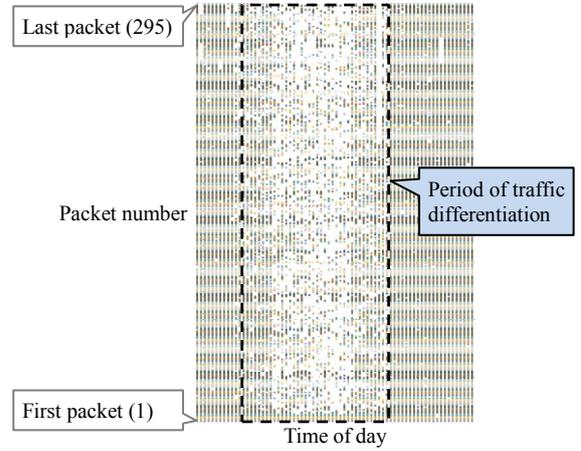


Fig. 6. Packet loss behavior inside and outside the period of traffic differentiation. The x-axis indicates the time of day and corresponds to the packet train number. The y-axis indicates the packet number within a packet train. The colored dots represent successfully received packets, while the white spaces represent lost packets.

For most packet trains during the periods without traffic differentiation (left and right sides of Fig. 6), we observed three cases: (a) no packet loss, (b) a few random losses, or (c) a single, modest burst loss in the last part of a packet train. Thus, examples of the set of  $L_{cls}$  values would be  $\{L_{cls}\} = \{0\}$  in case (a),  $\{L_{cls}\} = \{1\}$  or  $\{L_{cls}\} = \{1,1,1\}$  in case (b), and  $\{L_{cls}\} = \{5\}$  in case (c). Thus, the elements in each  $\{L_{cls}\}$  have *no variation* when there is no traffic differentiation. For most packet trains during the period of traffic differentiation (center of Fig. 6), however, we observed only case (d): a mix of many random losses and many burst losses throughout each packet train. In this case, an example of the set of  $L_{cls}$  values for each packet train would be  $\{L_{cls}\} = \{5,1,3,82,1,37,\dots,4,25\}$ . In other words, the elements in each  $\{L_{cls}\}$  have *large variation*.

We can leverage the difference in variation of  $\{L_{cls}\}$  to judge whether traffic differentiation occurs. The *variance-to-mean ratio* (VMR, also known as the *index of dispersion*) [40] is naturally suited for our purpose, because it is a normalized measure used to quantify the degree of variation of a set of observed occurrences. The VMR is defined as  $s^2/m$ , where  $m$  and  $s^2$  are the sample mean and variance of a set, respectively. If  $m = 0$ , the VMR is undefined mathematically, but for our convenience we treat the VMR as 0 in this case. Thus, the VMR for case (a) is treated as 0. Next, because  $s^2$  is not an *unbiased*

sample variance but a *sample* variance, the  $s^2$  of a set with a single element is 0. Therefore, the VMR is also 0 for a set like  $\{L_{cls}\} = \{1\}$  in case (b) and for case (c). Furthermore, because the  $s^2$  of a set whose elements all have the same value is 0, the VMR for a set like  $\{L_{cls}\} = \{1,1,1\}$  in case (b) is also 0. Thus, as shown in Fig. 3, for most packet trains during periods without traffic differentiation, the VMR remains 0 even though the packet loss rate is 1–5%. In contrast, the VMR in case (d) is large because  $s^2$  is large. During a period of traffic differentiation, the VMR has a tendency similar to that of the packet loss rate shown in Fig. 3. Hence, as with the PathQuick4-loss method, if the VMR exceeds a certain threshold, the PathQuick4-VMR method judges that traffic differentiation exists.

### B. Packet-Pair-Based Bandwidth Estimation

From the judgement obtained by the methods described in §V-A, we choose the subsequent procedure, as shown in Fig. 5: (1) If we judge that no traffic differentiation occurs, we delegate the subsequent procedure to PathQuick3’s packet-train-based algorithm to estimate the available bandwidth. (2) On the other hand, if we judge that traffic differentiation does occur, we use our novel *packet-pair*-based bandwidth estimation algorithm, in order to avoid PathQuick3’s overestimation shown in Fig. 3. For both cases (1) and (2) above, we *do not transmit any extra packets*. Note that a packet pair means precisely two packets.

The key idea of our packet pair method is our finding that the physical capacity of an end-to-end path is very likely to be close to the available bandwidth during a period of traffic differentiation. We can obtain the currently available bandwidth not by estimating itself but instead by estimating the physical capacity. Conventional wisdom [29] already suggests that we can estimate physical capacity by using a packet pair method. The originality of our packet pair method is twofold.

First, recall that if an MVNO specifies traffic shaping or policing at 1 Mbps per UE device, we assume that the traffic differentiator’s link becomes narrow, and then the physical capacity of an end-to-end path is limited to 1 Mbps. We also assume that a UE device receives no traffic other than a packet train during the period of packet train transmission through the traffic differentiator’s link. Note that it is possible that cross-traffic to other UE devices can pass through the traffic differentiator’s link during the *same* period of packet train transmission. If the two assumptions above are true, then because the definition of available bandwidth is “*physical capacity minus bandwidth being used during a certain time period*” [29], the available bandwidth is *equal* to the physical capacity *just before* the period of packet train transmission. We can thus leverage this relationship and estimate the physical capacity to obtain the available bandwidth.

A second point of originality is that we *generate* multiple packet pairs *dynamically from a received packet train*. For the same example described in §V-A-2, if a receiver successfully receives packets with packet numbers 1, 5, 7, 9, and 12, then the receiver generates four packet pairs:  $\{1,5\}$ ,  $\{5,7\}$ ,  $\{7,9\}$ , and  $\{9,12\}$ . In our method, while a sender *actually* transmits a *packet train*, a receiver behaves as if the sender *virtually* transmitted *multiple packet pairs* in the first place. With these *virtually generated* packet pairs, the receiver estimates physical capacity similarly to conventional packet pair methods (see [29] for details), and it outputs the resulting value as the estimated available bandwidth. In contrast, in the conventional methods [29], if a sender transmits six packet pairs (i.e., 12 packets for the example in §V-A-2), such as  $\{1,2\}$ ,  $\{3,4\}$ ,  $\{5,6\}$ ,  $\{7,8\}$ ,  $\{9,10\}$ , and  $\{11,12\}$  (note that in those methods, a sender transmits packet pairs every few seconds (or less frequently), and thus the sender cannot transmit packet pairs such as  $\{1,2\}$ ,  $\{2,3\}$ ,  $\{3,4\}$ , etc.), and if a receiver successfully receives the

packets with packet numbers 1, 5, 7, 9, and 12, then all of these packets’ *partners* (i.e., packet numbers 2, 6, 8, 10, and 11) have been lost. Because the partner relationships of the packets are invariant at both the sender and the receiver in the conventional methods, those methods recognize no successfully received packet pair and are thus *unable* to estimate physical capacity.

## VI. EXPERIMENTAL EVALUATION OVER THREE OPERATIONAL JAPANESE LTE MVNO NETWORKS

We conducted an experimental evaluation to compare PathQuick4 and PathQuick3 over three operational Japanese LTE MVNO networks for 13 weeks from June to September, 2017.

### A. Experimental Setup

#### 1) Ground Truth of Available Bandwidth

Since we cannot directly access the network nodes of mobile operators, the *ground truth* of the available bandwidth is unknown to us. Instead, although the available bandwidth and bulk TCP throughput are different network metrics [41], we follow [34] as our precedent and treat the bulk TCP throughput as a *reference* to the ground truth, or a *best effort* ground truth [34]. We consider this treatment reasonable because a bulk TCP flow aggressively utilizes the unused capacity of a narrow link. On the one hand, if no traffic differentiation occurs, then the wireless link between a base station and a UE device usually becomes a narrow link, and a bulk TCP flow aggressively exhausts radio resources. On the other hand, if traffic differentiation does occur, then the traffic differentiator’s link usually becomes narrow, and a bulk TCP flow aggressively exhausts the limited bandwidth (e.g., 1 Mbps per UE device). By using ns-3 to simulate an LTE environment with various network scenarios, including traffic shaping and policing, we validated that the available bandwidth is always reasonably close to the bulk TCP throughput. Because of a space limitation, however, the simulation results cannot be shown here.

#### 2) Experimental Environment

Figure 7 illustrates the experimental environment. We implemented the PathQuick4 receiver as an Android app in a Samsung Galaxy S7 Edge smartphone. A Linux server with a 1-Gbps FTTH connection was deployed for the PathQuick4 sender. The experiment was performed at diverse locations in Kanagawa Prefecture, Japan (population: 9 million). We obtained the bulk TCP throughput with one of the most well-known speed test Android apps in Japan [38]. For each measurement, we (1) ran a downlink speed test once, which took about 10 seconds, and (2) received 10 of PathQuick4’s probing packet trains, which also took about 10 seconds.

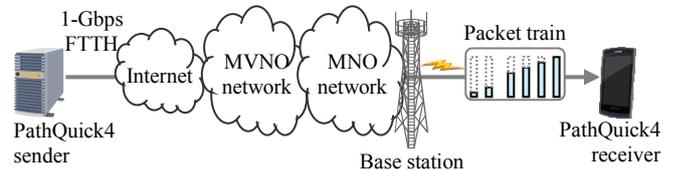


Fig. 7. Experimental environment, including operational networks.

#### 3) Estimation Parameter Choice

We chose the probable bandwidth range for PathQuick4 as follows. The current fastest average LTE downlink speed among the world’s nations is 47 Mbps [42], and typical speeds have continued to increase in recent years. With an eye to the future, we multiplied the above rate by 1.5, obtaining 70.5 Mbps as the maximum probable bandwidth for PathQuick4.

Figure 4 illustrates the parameters involved in reaching the 70.5-Mbps target. We set the size of the first packet as  $P_1 = 36$  bytes, the increase in packet size as  $\Delta P = 13$  bytes, the number

of packets in a packet train as  $N = 295$ , and the time interval as  $T_{quick} = 0.43$  ms. Therefore, the size of the last packet was  $P_N = 36 + 13 \times (295 - 1) = 3,858$  bytes, giving a maximum probable bandwidth of  $P_N / T_{quick} = 8 \times 3,858 / (0.43 \times 10^{-3}) = 71.8$  Mbps.

## B. Experimental Results

### 1) Estimated Available Bandwidth

Figure 8 shows the mean estimated available bandwidth from 10 packet trains for both PathQuick4 and PathQuick3, together with the corresponding speed test results over the three MVNO networks. Figures 8(a) and (b) show snapshots of MVNOs A and B, respectively, for a 1-day duration, while Fig. 8(c) shows a snapshot of MVNO C for a half-day duration.

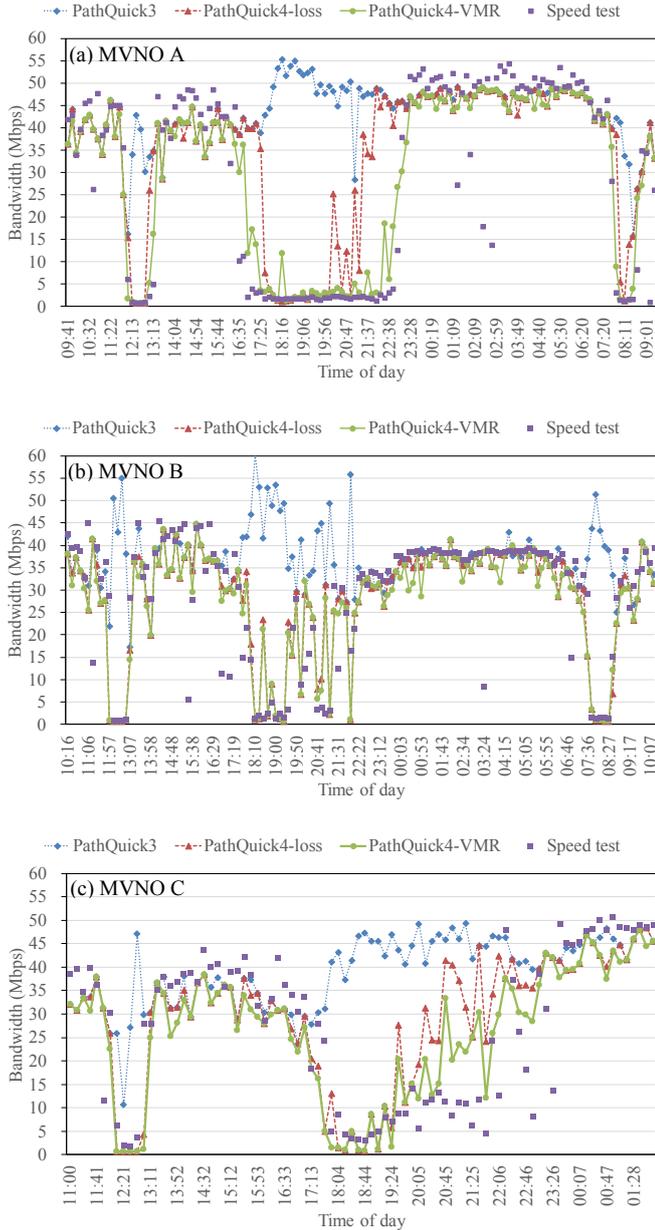


Fig. 8. Mean estimated available bandwidth obtained by each method and speed test results over the three MVNO networks.

Note that we obtained three estimated values (via PathQuick3, PathQuick4-loss, and PathQuick4-VMR) from a single packet train, as shown in Fig. 5. We empirically chose the thresholds for PathQuick4-loss and PathQuick4-VMR as 6% and 0.05 Mbps, respectively.

Overall, PathQuick4 clearly outperformed PathQuick3 in terms of estimation accuracy. Specifically, during the period of

traffic differentiation, while PathQuick3 consistently overestimated the available bandwidth, the estimated values from PathQuick4-loss and PathQuick4-VMR were closer to the speed test results. This result means that our packet pair method worked well. As we showed in §V-A-2, for certain cases the VMR is constantly 0 even though the packet loss rate is non-zero (e.g., 1–5%). As a result, the judgements of PathQuick4-VMR were more often correct than those of PathQuick4-loss, and thus, PathQuick4-VMR could output more accurate estimates.

### 2) Estimation Error

To analyze the estimation accuracy of each method in detail, we obtained the cumulative distribution function (CDF) of the mean absolute error (MAE) of each method, as shown in Fig. 9. The MAE is defined as

$$MAE = \frac{1}{n} \sum_{i=1}^n |\hat{A} - A|,$$

where  $\hat{A}$  is the estimated available bandwidth,  $A$  is the measured speed test value, and  $n$  is the sample size of  $\hat{A}$  ( $n = 10$  packet trains in this case). Note that Fig. 9 includes the results for all 13 weeks of experiments, while Fig. 8 shows snapshot results for 1 day or less. The colored circles in Fig. 9 show the 95th percentile of the MAE for each method. Thus, PathQuick4-VMR and PathQuick4-loss were consistently more accurate than PathQuick3. While PathQuick4-VMR and PathQuick4-loss were comparable for MVNO B, PathQuick4-VMR was more accurate than PathQuick4-loss for MVNOs A and C.

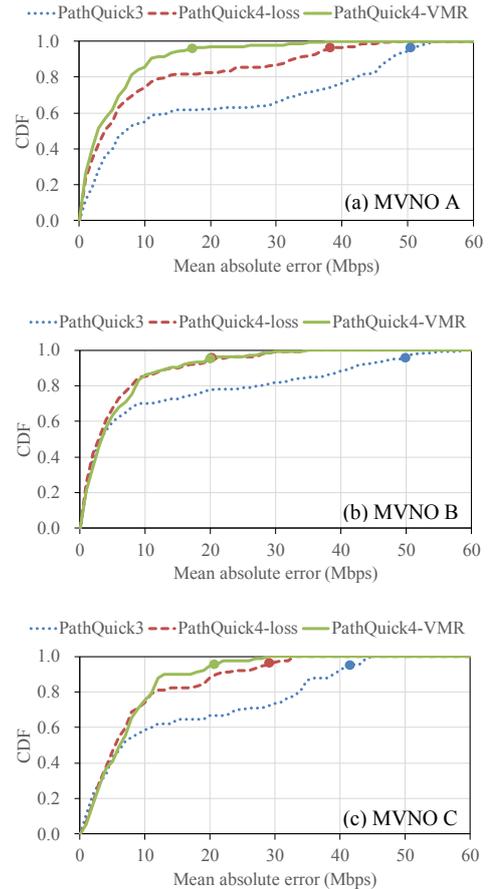


Fig. 9. CDFs (cumulative distribution functions) of the MAE (mean absolute error) for each method over the three MVNO networks. The colored circles indicate the 95th percentile.

### 3) During Periods of Traffic Differentiation Only

To analyze the impact of traffic differentiation, we excerpted all periods in which it occurred from the complete results for all

13 weeks of experiments (excluding the rest). Figure 10 shows the averaged MAE of each method for each MVNO, as well as the ratio of the averaged MAE for PathQuick4 to that for PathQuick3, expressed as a percentage. We found that this ratio across all three MVNOs (rightmost data in Fig. 10) was only 12%, i.e., close to an order of magnitude better.

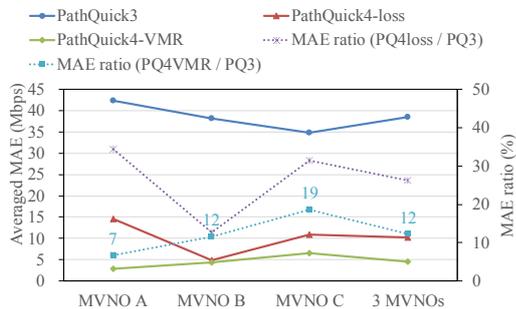


Fig. 10. Averaged MAE during periods of traffic differentiation only.

## VII. CONCLUSION AND FUTURE WORK

The proposed PathQuick4 method can accurately estimate the available bandwidth despite traffic differentiation. Our experimental evaluation showed that the overall estimation error of PathQuick4 during periods of traffic differentiation was only 12% of a conventional method's error.

In our future work, we plan to compare pathChirp [30], Pathload [31], and PathQuick4 over various operational MVNO networks. We also plan to conduct larger-scale experiments of longer duration than the one reported in this paper, in a diverse range of nations.

## REFERENCES

- [1] GSMA Intelligence, "The global MVNO footprint: A changing environment," February 20, 2015. [online] <https://www.gsmainelligence.com/research/2015/02/the-global-mvno-footprint-a-changing-environment/490/>
- [2] GlobalData Plc, "MVNOs in Europe: Data and M2M segments will present the highest growth opportunities to MVNOs in Europe," March, 2016. [online] <https://www.globaldata.com/store/report/tc0050ir--mvnos-in-europe-data-and-m2m-segments-will-present-the-highest-growth-opportunities-to-mvnos-in-europe/>
- [3] Ministry of Internal Affairs and Communications (MIC), "Quarterly data on the number and share of telecommunications service subscribers," September 22, 2017 [in Japanese]. [online] [http://www.soumu.go.jp/menu\\_news/s-news/01kiban04\\_02000126.html](http://www.soumu.go.jp/menu_news/s-news/01kiban04_02000126.html)
- [4] —, "White paper on information and communications in Japan," Figure 7-2-1-1, July 28, 2017 [in Japanese]. [online] <http://www.soumu.go.jp/johotsusintokei/whitepaper/ja/h29/html/nc272110.html>
- [5] M. Raiss-El-Fenni et al., "Towards an optimal pricing for mobile virtual network operators," *Networked Systems*, pp. 507–512, Springer, 2015.
- [6] T. Wu, "Network neutrality, broadband discrimination," *Journal on Telecommunications and High Technology Law*, pp. 141–176, 2003.
- [7] The New York Times, "F.C.C. repeals net neutrality rules," December 14, 2017. [online] <https://www.nytimes.com/2017/12/14/technology/net-neutrality-repeal-vote.html>
- [8] S. Akhshabi, A. C. Begen, and C. Dovrolis, "An experimental evaluation of rate-adaptation algorithms in adaptive streaming over HTTP," *ACM MMSys*, pp. 157–168, 2011.
- [9] T. Oshiba and K. Nakajima, "Quick end-to-end available bandwidth estimation for QoS of real-time multimedia communication," *IEEE ISCC*, pp. 162–167, 2010.
- [10] T. Oshiba and K. Nakajima, "Quick and simultaneous estimation of available bandwidth and effective UDP throughput for real-time communication," *IEEE ISCC*, pp. 1123–1130, 2011.
- [11] T. Oshiba, K. Nogami, K. Nihei, and K. Satoda, "Robust available bandwidth estimation against dynamic behavior of packet scheduler in operational LTE networks," *IEEE ISCC*, pp. 1276–1283, 2016.
- [12] N. Sato, T. Oshiba, K. Nogami, A. Sawabe, and K. Satoda, "Experimental comparison of machine learning-based available bandwidth estimation methods over operational LTE networks," *IEEE ISCC*, pp. 339–346, 2017.
- [13] H. Holma et al., "LTE small cell optimization: 3GPP evolution to release 13," p. 380, John Wiley & Sons, Inc., 2016.
- [14] Y. Zhang, Z. M. Mao, and M. Zhang, "Detecting traffic differentiation in backbone ISPs with NetPolice," *ACM IMC*, pp. 103–115, 2009.
- [15] M. B. Tariq et al., "Detecting network neutrality violations with causal inference," *ACM CoNEXT*, pp. 289–300, 2009.
- [16] M. Dischinger et al., "Glasnost: Enabling end users to detect traffic differentiation," *USENIX NSDI*, pp. 405–418, 2010.
- [17] P. Kanuparth and C. Dovrolis, "DiffProbe: Detecting ISP service discrimination," *IEEE INFOCOM*, pp. 1649–1657, 2010.
- [18] —, "ShaperProbe: End-to-end detection of ISP traffic shaping using active methods," *ACM IMC*, pp. 473–482, 2011.
- [19] S. Sundaresan et al., "Broadband internet performance: A view from the gateway," *ACM SIGCOMM*, pp. 134–145, 2011.
- [20] S. Basso, A. Servetti, and J. C. D. Martin, "The network neutrality bot architecture: A preliminary approach for self-monitoring of Internet access QoS," *IEEE ISCC*, pp. 1131–1136, 2011.
- [21] Z. Zhang, O. Mara, and K. Argyraki, "Network neutrality inference," *ACM SIGCOMM*, pp. 63–74, 2014.
- [22] T. Flach et al., "An Internet-wide analysis of traffic policing," *ACM SIGCOMM*, pp. 468–482, 2016.
- [23] A. M. Kakhki et al., "BingeOn under the microscope: Understanding T-Mobile's zero-rating implementation," *ACM Internet-QoE*, pp. 43–48, 2016.
- [24] F. Zarinni et al., "A first look at performance in mobile virtual network operators," *ACM IMC*, pp. 165–172, 2014.
- [25] N. Vallina-Rodriguez et al., "Beyond the radio: Illuminating the higher layers of mobile networks," *ACM MobiSys*, pp. 375–387, 2015.
- [26] P. Schmitt, M. Vigil, and E. Belding, "A study of MVNO data paths and performance," *PAM*, pp. 83–94, Springer, 2016.
- [27] V. Bashko et al., "BonaFide: A traffic shaping detection tool for mobile networks," *IFIP/IEEE IM*, pp. 328–335, 2013.
- [28] A. M. Kakhki et al., "Identifying traffic differentiation in mobile networks," *ACM IMC*, pp. 239–251, 2015.
- [29] R. Prasad, C. Dovrolis, M. Murray, and K. Claffy, "Bandwidth estimation: Metrics, measurement techniques, and tools," *IEEE Network*, vol. 17, issue 6, pp. 27–35, 2003.
- [30] V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell, "pathChirp: Efficient available bandwidth estimation for network paths," *PAM Workshop*, 2003.
- [31] M. Jain and C. Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput," *ACM SIGCOMM*, pp. 295–308, 2002.
- [32] Y. Takano et al., "Estimating available bandwidth in mobile networks by correlation coefficient," *IEICE/KICS APNOMS poster*, pp. 1–4, 2016.
- [33] Y. Yang et al., "A novel available bandwidth estimation scheme in probe packets losing condition," *IEEE IECC*, pp. 125–129, 2009.
- [34] D. Koutsonikolas and Y. C. Hu, "On the feasibility of bandwidth estimation in wireless access networks," *Wireless Networks*, vol. 17, issue 6, pp. 1561–1580, Springer, 2011.
- [35] V. Fajardo et al., "An efficient and loss tolerant method for measuring available bandwidth," *IEEE GLOBECOM*, pp. 3407–3412, 2009.
- [36] S. Blake et al., "An architecture for differentiated services," IETF RFC 2475, §2.3.3.3 Shapers and §2.3.3.4 Droppers, 1998.
- [37] S. Miyamoto, "MVNO LTE access," IJ Technology, February 27, 2012. [online] [https://www.ij.ad.jp/en/company/development/tech/lte\\_mvno/](https://www.ij.ad.jp/en/company/development/tech/lte_mvno/)
- [38] IID, Inc., "RBB SPEED TEST," *Google Play*. [online] <https://play.google.com/store/apps/details?id=com.rbbtoday.speedtest>
- [39] Ministry of Internal Affairs and Communications (MIC), "Current status of mobile communication traffic in Japan," 2017 [in Japanese]. [online] <http://www.soumu.go.jp/johotsusintokei/field/data/gt010601.xls>
- [40] D. J. Wilkinson, "Stochastic modelling for systems biology, second edition," p. 97, CRC Press, 2011.
- [41] M. Jain and C. Dovrolis, "Ten fallacies and pitfalls on end-to-end available bandwidth estimation," *ACM IMC*, pp. 272–277, 2004.
- [42] OpenSignal, "The state of LTE (November 2017)," 2017. [online] <https://opensignal.com/reports/2017/11/state-of-lte>