

QUIC Throughput and Fairness over Dual Connectivity

David Hasselquist, Christoffer Lindström, Nikita Korzhitskii,
Niklas Carlsson, and Andrei Gurto

Linköping University, Sweden

Abstract. Dual Connectivity (DC) is an important lower-layer feature accelerating the transition from 4G to 5G that also is expected to play an important role in standalone 5G. However, even though the packet reordering introduced by DC can significantly impact the performance of upper-layer protocols, no prior work has studied the impact of DC on QUIC. In this paper, we present the first such performance study. Using a series of throughput and fairness experiments, we show how QUIC is affected by different DC parameters, network conditions, and whether the DC implementation aims to improve throughput or reliability. Our findings provide insights into the impacts of splitting QUIC traffic in a DC environment. With reasonably selected DC parameters and increased UDP receive buffers, QUIC over DC performs similarly to TCP over DC and achieves optimal fairness under symmetric link conditions when DC is not used for packet duplication.

Keywords: QUIC · Dual Connectivity · Throughput · Fairness · Transport Protocol · Multipath

1 Introduction

The end-to-end performance depends on the interactions between protocols in different network layers. As new features are introduced on the lower layers, it is therefore important to understand the impact that such features and their parameters have on the upper layer protocols [5]. One such feature is Dual Connectivity (DC). DC was introduced in 4G, gained popularity with the introduction of 5G, and currently plays an integral role in accelerating the generational transition from 4G to 5G [7].

With DC, users can transmit and receive data from two base stations concurrently. This allows users to use both 4G and 5G networks in parallel, simplifying the above-mentioned generational transition. However, it has also been argued that DC should be a part of future 5G solutions needed to meet the requirements of Ultra-reliable and Low-Latency Communications (URLLC) [2, 17]. Combined with its increased usage, this has made DC an important 5G feature.

Like multi-path transport protocols [8, 12, 27], DC can be used to combine WiFi with 4G and 5G solutions. Furthermore, like these protocols, DC can be used to achieve improved throughput (by sending different data over different

paths), to increase reliability (by transmitting the same data over the different paths), or both. However, in contrast to the transport-layer multipath solutions, DC is performed within the link layer of the network stack and is therefore in practice invisible to transport layer protocols such as TCP and QUIC. This is an important observation since DC may introduce jitter or reordering of packets that can significantly impact TCP and QUIC performance.

In parallel with the transitioning of different network generations, Google recently introduced QUIC as a next generation transport-layer solution aimed at addressing some shortcomings with TCP [16]. QUIC is implemented in the user-space, on top of UDP, and provides much improved stream multiplexing compared to TCP. This is important to speed up web connections in the presence of packet losses and/or modern HTTP/2 traffic. Initial research shows that QUIC allows performance improvements over TCP in several cases while providing an easy way to achieve fast incremental deployment [16]. Popular services that already today use QUIC include Google search services, Chrome, Chromium, YouTube and Facebook [10, 16].

Due to the increasing use and popularity of both QUIC and DC, combined with the continuous rollout of 5G networks using DC, it is important to understand how QUIC performs over DC under different network conditions, and the impact that different DC parameters have on QUIC performance.

In this paper, we present the first performance evaluation of QUIC over DC. First, a testbed is set up to simulate DC. The testbed captures QUIC and TCP performance under a wide range of network behaviors (based on bandwidth, delay, and loss conditions) and the impact of different DC parameters. Second, using a series of throughput and fairness experiments, we show how QUIC is affected by different DC parameters, network conditions, and whether the DC implementation aims to improve throughput or reliability. For our throughput evaluation, we primarily compare the throughput of QUIC over DC with that of TCP over DC, and for our fairness comparisons we compare the throughput (and calculate a fairness index) of competing flows when using QUIC over DC. We also present results using different QUIC implementations (aioquic, ngtcp2) and congestion control algorithms (NewReno, CUBIC). Our findings provide insights into the impact that DC and its parameters have on QUIC performance. For example, we show the value of increasing the UDP receive buffers when running QUIC over DC, that QUIC over DC can achieve similar throughput as TCP over DC, and that QUIC over DC can achieve optimal fairness under symmetric link conditions, except if DC duplicates packets to increase reliability.

Outline: Sections 2 and 3 introduce DC and present related works, respectively. The following sections present our methodology (Section 4), performance results (Section 5), and conclusions (Section 6).

2 Dual Connectivity

DC, sometimes called inter-node radio resource aggregation, is a multi-connectivity technique introduced in release 12 of the third-generation partnership project

(3GPP) [1]. The aim was to increase reliability, performance, and signaling due to frequent handovers in scenarios where macro and micro cells are connected with a non-ideal backhaul (X2) link. DC tries to achieve this by splitting the traffic over multiple paths.

Figure 1 shows an overview of DC in a Radio Access Network (RAN) environment. With DC, a User Equipment (UE) connects to two different Evolved Node Bs (eNBs) [2]. One of the nodes will serve as Master eNB (MeNB), and the other one will serve as Secondary eNB (SeNB). Each of the MeNB and SeNB contains a separate Radio Link Control (RLC) and Media Access Control (MAC) layer, while sharing the same Packet Data Convergence Protocol (PDCP) layer.

DC is similar to carrier aggregation [26], but is performed in the PDCP layer instead of the MAC layer. Carrier aggregation uses the same scheduler for the separate connections and requires an ideal X2 link. The split connections are therefore often transmitted from the same node. In contrast, DC uses two separate schedulers together with a non-ideal X2 link, and packets are often originating from two different nodes.

PDCP is a sublayer located inside the link layer, just below the network layer and above RLC and MAC. The main tasks of PDCP are header compression and decompression, ciphering, integrity protection, transfer of data, sequence numbering, reordering and in-order delivery [3]. The PDCP layer can be broken out into a unit called a Packet Processor (PP), which connects to Serving Gateway (SGW), MeNB and SeNB using a GTPU-tunnel. SGW is connected to the Packet Data Network Gateway (PGW), which connects to the public internet. The PP can also be a part of MeNB. In this case, MeNB splits the traffic and the link between MeNB and SeNB becomes the X2 link. In both scenarios, the traffic is split in the PDCP layer.

3 Related Work

Dual connectivity: Unlike TCP, QUIC is relatively new, and there are few studies of it in specific scenarios such as DC. As QUIC shares similarities with TCP, we can obtain initial insights from research about DC that uses TCP as the transport protocol. Polese et al. [22] study the performance of TCP when using DC to perform mobile handovers for an UE and compare the performance with different single connection mobility management solutions. They show that DC can improve TCP goodput by quickly moving the traffic from one of the two DC links to the other.

Other studies have focused on specializations of DC; e.g., LTE-WLAN Aggregation (LWA) [14, 15], which allows for network traffic over LTE and WLAN. Jin et al. [14] show that splitting TCP over LTE and WiFi at the PDCP layer can achieve similar throughput and better fairness than MP-TCP; demonstrating the value of lower-layer traffic splitting. Khadraoui et al. [15] investigate the effect of PDCP reordering when using TCP in LWA over heterogeneous links. Their results show that PDCP reordering can have adverse effects on TCP throughput,

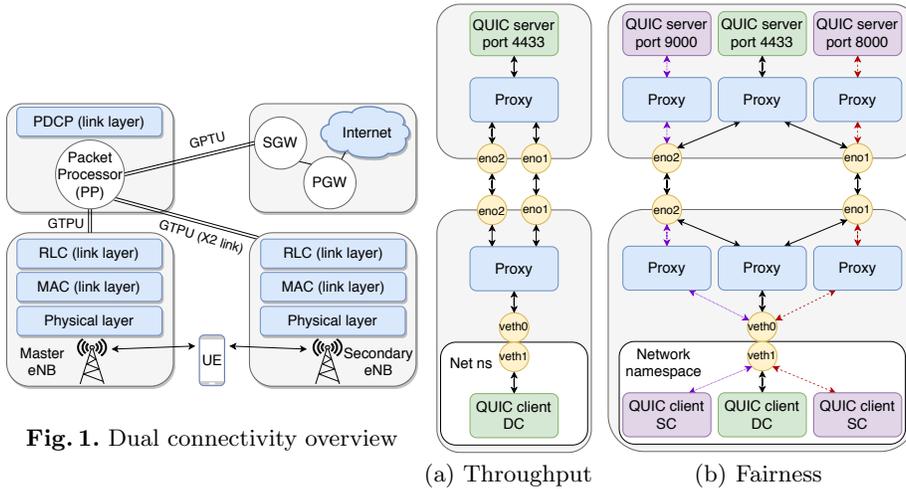


Fig. 1. Dual connectivity overview

(a) Throughput (b) Fairness

Fig. 2. Testbed for throughput and fairness

and that in some cases it is better to use only one link. While some works have looked at TCP with DC, no prior work has studied QUIC performance over DC.

Upper-layer multipathing: Multipathing is similar to DC but performed higher up in the network stack. Most such solutions are implemented in the transport layer, e.g., SCTP [12] and MP-TCP [27], but some are implemented in the network layer [9]. Here, we focus on QUIC-based solutions. De Coninck and Bonaventure [8] implement Multipath QUIC (MP-QUIC) based on quic-go and lessons learned from MP-TCP, and show that serving QUIC over multiple paths is beneficial. Mogensen et al. [19] expands MP-QUIC to Selective Redundant MP-QUIC (SR-MPQUIC). Their solution modifies the congestion control algorithm, the scheduler, and the stream framer. SR-MPQUIC reduces latencies and improves reliability for priority data at a small increase in bandwidth usage and latencies for background data. The results show the importance of proper packet scheduling and the value of packet duplication. While additional cross-layer communication would be required to benefit DC, QUIC also includes some unique attributes to assist packet/flow scheduling [23].

Fairness: Fairness can be difficult to judge when there are multiple paths with different amount of resources. Becke et al. [6] study the fairness of different congestion control algorithms in multipath scenarios, focusing on two fairness types: link-centric and network-centric flow fairness. Raiciu et al. [25] study how MP-TCP can replace single connections and load balancing in data centers. For specific topologies, MP-TCP significantly improved fairness and provided throughput closer to optimal compared to single connectivity using random load balancing. To judge fairness, they and many others [25, 27, 28] evaluate multipathing using Jain’s fairness index (JFI) [13]. Similar to these works, we use JFI here.

Table 1. Hardware and operating systems

Component	Client	Server
OS	Ubuntu 18.04.3 LTS	Ubuntu 18.04.3 LTS
Kernel	Linux 4.15.0-74-lowlatency	Linux 5.3.0-26-generic
Processor 1 & 2	Intel(R) Xeon(R) CPU E5-2690 v3	Intel(R) Xeon(R) CPU E5-2667 v3
eno1 & eno2	82599ES 10-Gigabit SFI/SFP+	82599ES 10-Gigabit SFI/SFP+

4 Methodology

4.1 Dual connectivity testbed

Figure 2 shows an overview of the testbed used for studying QUIC performance over DC. We used one machine to capture client-side behavior and performance, and one machine to capture server- and network-side effects. The two machines were connected via two network interface pairs, each supporting 10 Gbps full duplex. Hardware specifications are given in Table 1. We next describe the configurations used for our throughput and fairness tests, respectively, and our proxy implementation used to simulate DC and PDCP.

Throughput test configuration: For our throughput tests (Figure 2(a)), we used one client, one server, and studied the performance impact of DC parameters and the network conditions between them. In our baseline tests, both the QUIC server and QUIC client used aioquic [4]. When running comparison tests with TCP, we used a Hypercorn server using HTTP/2 over TLS 1.3, and the client used curl to make HTTP/2 requests. As baseline, both TCP and QUIC used NewReno for congestion control.

Since traffic splitting with DC is implemented in the link layer, QUIC (and TCP) are unaware that the traffic is sent over multiple paths, and therefore do not need to be modified. However, as DC was introduced for radio technology, the link layer functions differ from the Ethernet links used here. To simulate the functionality of DC and PDCP, two proxies were implemented: one at the client and one at the server.

The QUIC client was launched inside a network namespace. Two virtual interfaces were created to forward data to and from the namespace. The server side does not require a network namespace as DC is only studied on the downlink. To simulate different network conditions, tc in Linux was used to add extra delay, jitter, loss, and bandwidth limitations.

Fairness test configuration: For our fairness tests (Figure 2(b)), we used three clients and three servers. One end-to-end connection was performing DC, while the other two used single connectivity (SC) over interface eno1 and eno2, respectively. The server with port 8000 was operating only on eno1, while the server on port 9000 only on eno2. The QUIC server on port 4433 used DC and operated on both interfaces. Each server and client were equipped with its own proxy, simulating the PDCP functionality for each connection independently.

Proxy-based implementation: To capture the PDCP functionality, packets originating from the server are caught by iptables OUTPUT chain and delivered to a NFQUEUE, before being read by the server proxy. The server proxy

then adds a 2-byte PDCP sequence number to each packet and routes the packets to the client over two interfaces. When running DC, the server proxy alternates between the two interfaces.

At the client proxy, packets are caught in the PREROUTING chain and delivered to NFQUEUE. The client proxy can then read from the queue, perform PDCP convergence of the two streams, do PDCP reordering, and remove the sequence numbers that were added by the server proxy.

If a packet is received in order, it is immediately forwarded to the client. However, if a packet is out of order, it is kept until the missing packets are processed or until a PDCP timer of 200ms is reached. If the timer is reached, all packets before the missing packet and all consecutive packets after the missing packets are delivered. The reordering algorithm follows the PDCP standard described in 3GPP [3] and the testbed was developed in close consultation with Ericsson. Our proxy adds around 1ms to the total RTT, assuming that the packets arrive in order. Without PDCP, large reordering occurs, resulting in QUIC having a very low throughput.

4.2 Performance testing

To understand how DC affects QUIC, a series of tests are performed that captures the impact of different DC parameters and network conditions. In our experiments, we vary one parameter at a time, starting with a default configuration, while keeping the others constant (as per the default configuration). In the throughput tests, the client downloads a 100MB file, and in the fairness tests each client downloads a 1GB file and we measure the clients' performance for the first three minutes of the download. For each test configuration, we run ten tests and calculate both average and standard deviation values for the metrics of interest.

DC parameters and default configurations: The primary DC parameters we varied were the DC batch size and DC batch split. These parameters determine how many packets are sent over each interface before the server proxy switches to the other interface. For example, with a DC batch size of 100 and a DC ratio of 9:1 (90% eno1 and 10% eno2), the proxy would send 90 packets over eno1, before switching over to send 10 packets over eno2. In our default experiments, the default DC batch size and DC ratio was configured to 100 and 1:1, respectively.

Network emulation parameters and default configurations: To capture different network conditions, we primarily varied the bandwidths, delays, and loss rates of the links. For both the bandwidths and delays, we present experiments both where we vary the average values and where we vary the ratio between the two links. In the case we vary one of the ratios, we keep the average value of that metric constant. For example, a bandwidth ratio of 3:1 corresponds to 30Mbps and 10Mbps for the downlink interfaces eno1 and eno2, respectively. In our default experiments, each link operates at 20 Mbps and has normally distributed per-packet delays with a mean of 10 ms and a standard deviation of 10%.

QUIC and TCP configurations/versions: Throughput tests for QUIC are performed both with the default UDP receive socket buffer size and a larger receive buffer size. The larger size is used to give a fair comparison to TCP, as the kernel performs buffer autotuning for TCP [21]. When studying fairness, QUIC with modified buffer size is used and the fairness is calculated using JFI.

In our default scenarios we use aioquic with NewReno. However, as discussed by McMillan and Zuck [18], the QUIC RFC is ambiguous, open for interpretation, and differences between QUIC implementations following the RFC have been demonstrated using specification testing. We therefore repeated our experiments with both another QUIC implementation (ngtcp2 [20]) and congestion control algorithm (CUBIC).

Trace-based evaluation: Finally, experiments were repeated using a real LTE bandwidth trace collected by Raca et al. [24]. The specific bandwidth trace (*Static A_2018.02.12.16.14.02*) has an average throughput of 4.5 Mbps for the first 200 seconds.

5 Evaluation Results

5.1 Dual connectivity parameters

DC batch size: When using DC, network operators must select a good DC batch size for each connection. To illustrate the impact of this choice on QUIC performance, Figure 3(a) shows the throughput as a function of the DC batch size. (We omitted the standard deviations from all figures, as they are small; e.g., well within 1Mbps in more than 90% of the cases.) In general, large DC batch sizes result in lower throughput. One reason for this is reduced link utilization. For example, Figures 3(b) and 3(c) show the link utilization of the two links when using a DC batch size of 50 and 500, respectively. With a large DC batch size, we see significant periods during which one of the links is underutilized as almost all packets are being forwarded over the other interface. With smaller DC batch sizes, both links can better be used concurrently. However, there is also a penalty to using too small batch sizes, as this increases the number of re-order events. The best batch sizes are instead typically in the mid-range (e.g., around 100-150), with the sweet spot depending on the protocol being used. Finally, we note that QUIC with modified buffers perform similar to TCP for much of the parameter range.

Figure 4(a) shows summary results for our fairness tests with varying DC batch sizes. Here, we measure fairness using Jain’s fairness index (JFI), shown using purple text, as averaged over 10 full runs. When discussing fairness, it is important to note that the relative throughput of the competing clients can vary significantly over time. This is illustrated in Figure 4(b) where we show example throughput for the three competing clients over a 3-minute long experiment with the default settings.

Similar to the throughput, the fairness is negatively affected by large DC batch sizes. In fact, the user using DC observe a significant throughput reduction

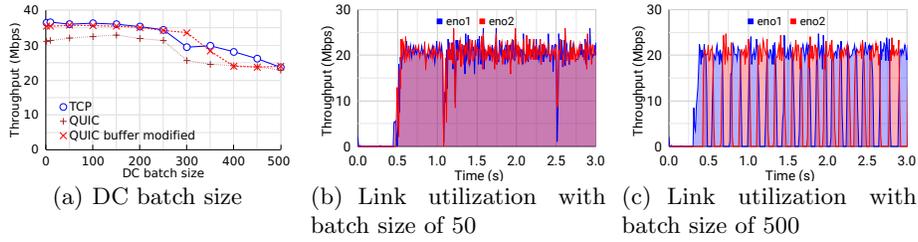


Fig. 3. Throughput and link utilization for different DC batch sizes

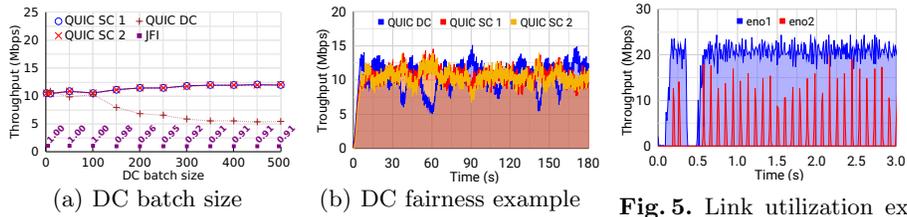


Fig. 4. Fairness for different DC batch sizes

Fig. 5. Link utilization example with unbalanced DC ratio of 9:1

with batch sizes of 150 and above. Here, SC clients can monopolize the links during the DC client’s off periods, while DC is always sharing the link it is currently sending to at every point of time. This allows SC clients to increase their cwnd further than the DC client and to use a larger bandwidth share.

DC batch split: Operators also control the DC ratio. This parameter determines the split over the two links. Figure 5 illustrates an unbalanced example with a DC batch split of 9:1, in which 90% of the packets are sent over the main interface (eno1), and Figure 6(a) shows the throughput as a function of the percent of packets sent over eno1. As per our default case, both links have the same network conditions. In this case, the throughput peaks when using a 50/50 split, and decreases as a convex function as the split becomes more uneven. This decrease is caused by poor link utilization of the less loaded link (eno2 in Figure 5), but also demonstrates the value of DC.

Figure 7(a) shows our corresponding fairness results. When the ratio is significantly skewed (e.g., below 20% or above 80%), the throughput of the SC with the higher throughput increase/decrease at roughly the same rate as the DC’s throughput increase/decrease, whereas the other SC has fairly constant throughput over these skewed splits. In this region, the DC compete (almost) fairly only over the more utilized interface. As the DC split becomes more even, the overall fairness improves, with optimal fairness and all connections having roughly equal throughput when perfectly balanced.

5.2 Network conditions

Bandwidth ratio: Figures 6(b) and 6(c) show the throughput for different bandwidth ratios. Figure 6(b) shows results for the case when the batch split is 50/50, and Figure 6(c) shows results for when the batch split is selected to match

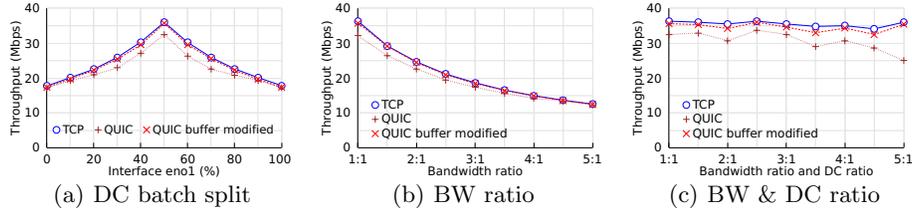


Fig. 6. Throughput for different batch splits and bandwidth ratios

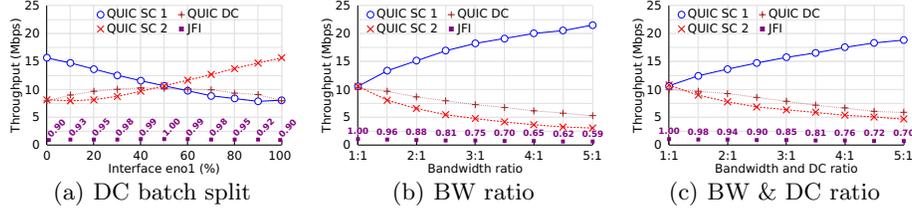


Fig. 7. Fairness for different batch splits and bandwidth ratios

the bandwidth ratio. Figure 6(b) illustrates the importance of matching ratios, as the highest throughputs are achieved with a ratio of 1:1. As the ratio increases, a 50/50 batch split underutilizes the link with higher bandwidth. In contrast, when the DC batch split is selected to match the bandwidth ratio (Figure 6(c)), a much better overall throughput is achieved. With QUIC buffer modified and TCP, the impact is very small. The reason for the worse performance of QUIC with default buffers is the higher burstiness caused by increased reordering. Despite PDCP mitigating reordering, it results in increasing RTTs.

The fairness results for the cases when we vary the bandwidth ratio of the two links are shown in Figures 7(b) and 7(c). Similar to the throughput results, higher fairness is achieved when the DC split is selected based on the capacity of the two links. For example, even when the bandwidth ratio is 5:1, the scenario in which the DC split matches the bandwidth ratio achieves a JFI of 0.70, compared to 0.59 in the case a 50/50 split is used. In both cases, the bandwidth usage is dominated by the SC user with higher bandwidth and the DC user relies heavily on the throughput achieved via the weaker link. However, the fairness improves as DC moves more traffic to the link with the higher bandwidth.

Delay ratio: Both the throughput and fairness are negatively affected by increasing delays, and in the case of a high average delay, these metrics are also negatively affected by an increasing delay ratio. This is illustrated by comparing the throughput Figures 8(a) and 8(b) or fairness Figures 9(a) and 9(b). For both types of experiments, the two figures show results for low-delay and high-delay scenarios, respectively. In the low-delay scenarios, the sum of the delays over the two links is 20ms, and in the high-delay scenario the sum is 200ms.

The throughput decrease is mostly due to increased packet reordering caused by the higher delays. In these cases, the PDCP layer will buffer more packets before performing a batch delivery to the QUIC client, causing packet bursts as

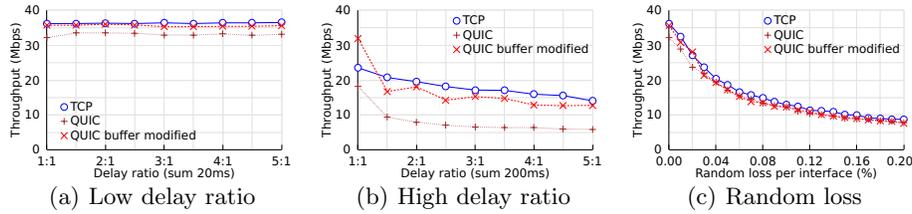


Fig. 8. Throughput for different delay and loss ratios

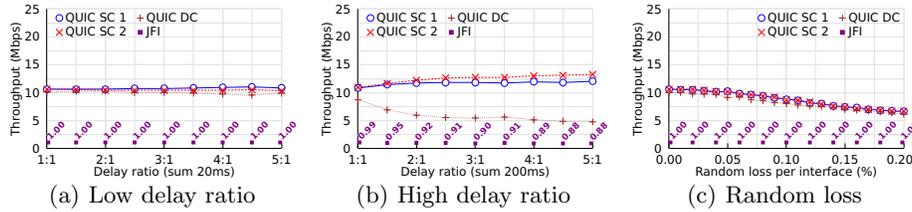


Fig. 9. Fairness for different delay and loss ratios

well as a higher RTT. Furthermore, after receiving a batch delivery, the clients will send a cumulative ACK for many packets, which will, for a short time, largely decrease the number of packets in flight when received at the server. The draft for QUIC [11] recommends a pacer, which helps the QUIC server recover from an ACK-burst by sending new packets at steadier pace. The advantage of more even pacing can be seen by the higher values observed with a delay ratio of 1:1 in Figure 8(b).

The increasing delays and delay ratios also negatively impact fairness. For example, in the low-delay case (Figure 9(a)), JFI reduces from 0.9996 to 0.9986 as the delay ratio increases from 1:1 to 5:1, whereas JFI drops from 0.9903 to 0.8758 for the high-delay case (Figure 9(b)). The higher throughput of SC eno2 compared to that of SC eno1 is due to its lower RTT.

Loss rates: While increased packet losses negatively impact the throughput (Figure 8(c)), small packet losses have very limited impact on the fairness index (Figure 9(c)).

5.3 Use of duplicate packets

Besides improving throughput, DC can also be used to increase connection reliability. However, DC with packet duplication negatively effects fairness. For example, in fairness tests with loss rates of 0-to-5% (Figure 10(b)) JFI is in the range from 0.39 to 0.48. For DC in Figure 10(b), we show both the combined interface throughput (B) and the goodput (X), which under an independence model with retransmissions (after simplification) can be related as $X = B(1 + p)/2$, where p is the loss rate. The low fairness stems from DC having a much higher end-to-end packet delivery probability (i.e., $1 - p^2$ vs $1 - p$ under independence assumptions) and lower end-to-end packet loss probability (i.e., p^2 vs p) compared to SC. This results in DC obtaining a larger share of the link bandwidths. These results show

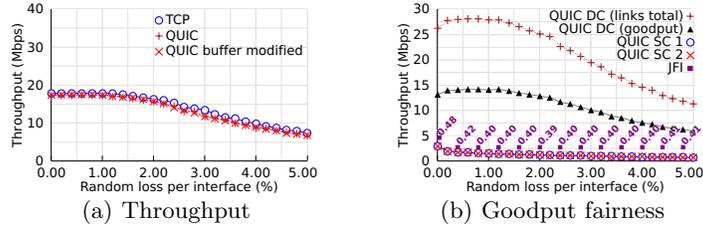


Fig. 10. DC throughput and fairness with duplicate packets

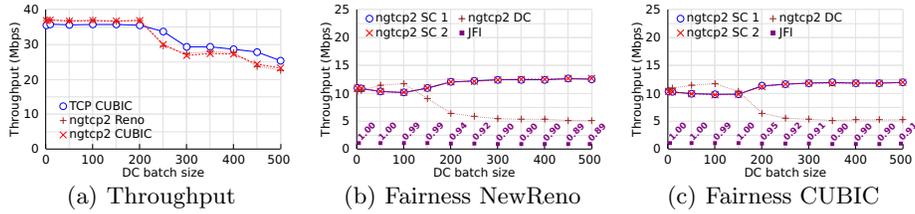


Fig. 11. Impact of congestion control algorithm: Performance examples with ngtcp2 as QUIC version for different batch sizes

that duplication can provide much higher reliability at the cost of fairness and goodput.

5.4 QUIC implementation and congestion control algorithm

To explore the impact of other QUIC implementations and congestion control algorithms, experiments were repeated using ngtcp2 and CUBIC. Figure 11(a) shows little to no differences in the results between different congestion control algorithms when varying DC batch size. However, when compared to Figure 3(a), differences can be observed between the QUIC implementations. While the results follow the same patterns for both implementations, the considerable throughput drop occurs at different batch sizes. Another noticeable difference is that ngtcp2 has a slightly higher throughput than aioquic at smaller batch sizes, exceeding the throughput for TCP.

Figures 11(b) and 11(c) show the corresponding fairness results. Again, only small differences between the NewReno and CUBIC results are observed. For example, the JFI differ by at most 0.02 (DC batch size of 500) between the algorithms. DC using CUBIC is initially slightly more resilient to performance drops occurring with a larger batch size. In contrast, NewReno allows the SC connections to achieve slightly higher throughput at larger batch sizes while the DC throughput is similar to CUBIC at higher batches. When comparing Figures 11(b) and 11(c) to 4(a), some differences can be seen between the QUIC implementations. Ngtcp2 is more aggressive, leading to the DC connection having slightly higher throughput than the SC connections at smaller batch sizes and a drastic reduction in throughput when the batch size increases. Aioquic has a more balanced sharing of the bandwidths at smaller batch sizes and see a smaller reduction in throughput at larger batch sizes.

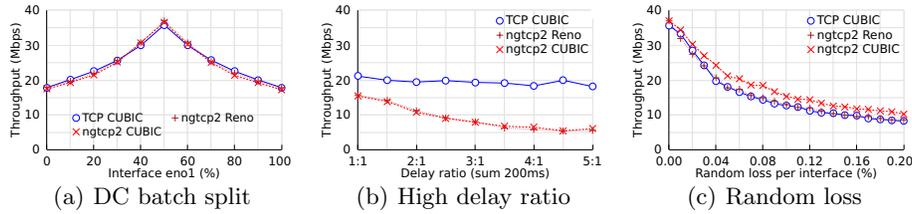


Fig. 12. Throughput when using ngtcp2 (as QUIC implementation) with NewReno and CUBIC for different parameters

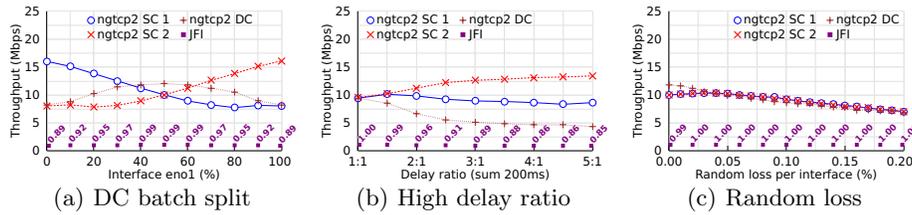


Fig. 13. Fairness when using ngtcp2 with NewReno for different parameters (CUBIC results very similar)

When studying the DC batch split using ngtcp2 and different congestion control algorithms (Figure 12(a)), minimal difference in the overall throughput is observed. Compared to aioquic in Figure 6(a), little differences are observed at more uneven ratios. Ngtcp2 achieves a higher throughput than aioquic and TCP at more balanced ratios. When comparing the corresponding fairness results in Figure 13(a) to 7(a), larger differences can be seen between the QUIC implementations. While the two implementations exhibit similar behavior at the most uneven split, the DC connection using ngtcp2 grows more aggressively than the aioquic counterpart when the ratio becomes more balanced. This growth lead to optimal fairness for aioquic, but results in a slightly unfair bandwidth allocation for ngtcp2.

A significant throughput difference between the QUIC implementations can be seen when comparing the high delay ratio experiments in Figures 12(b) and 8(b). Ngtcp2 achieves a significantly lower throughput than aioquic throughout the experiment. This is most likely due to differences in the pacer implementations. Differences can also be seen when comparing corresponding fairness tests (Figures 13(b) and 9(b)). Here, the DC connection using ngtcp2 achieves more fair throughput than the aioquic counterpart at balanced ratios and sees a slower drop in throughput when the ratio gets skewed. However, after the 2:1 ratio point, the DC connections' throughputs become the same for the two implementations. The ngtcp2 SC connection with a higher delay has a much worse performance than the aioquic counterpart at more skewed ratios.

Finally, when studying the impact of loss rates using ngtcp2 and CUBIC, only small differences are observed (Figures 12(c) and 13(c) compared to Figures 8(c) and 9(c)). However, in contrast to the other experiments, ngtcp2 using CUBIC shows a noticeable better performance compared to TCP CUBIC. Looking

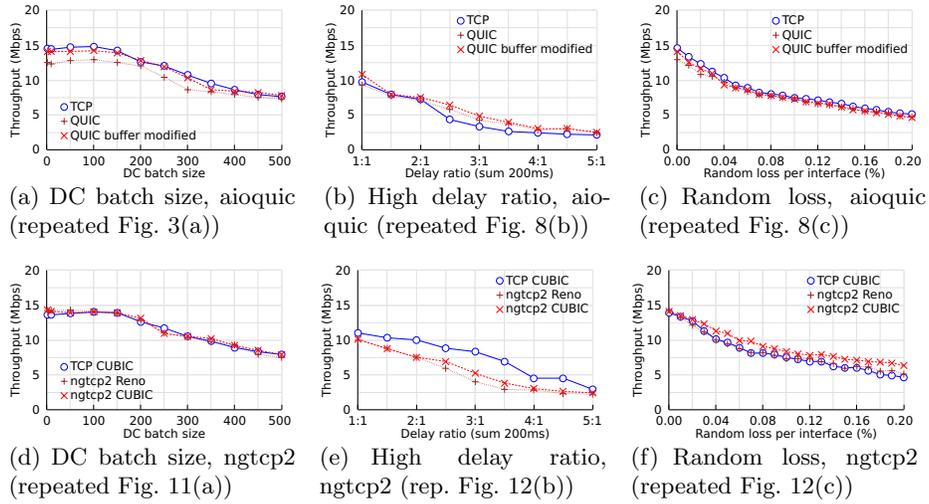


Fig. 14. Trace-driven bandwidth variability tests using aioquic and ngtcp2 as QUIC implementations

closer at the 0.08% loss case, we have observed that the TCP implementation more often stays in CUBIC’s TCP mode (used when detecting growth slower than Reno). This also explains why TCP CUBIC, TCP NewReno, and ngtcp2 NewReno perform similarly here.

In general, ngtcp2 achieves higher throughput than aioquic, even though both follow the same IETF recommendations. As discussed, differences can occur due to the RFC being open for interpretation. The execution speed and resources required by the two implementations also differ. Ngtcp2 is implemented in C and aioquic in Python. With ngtcp2, a larger receive buffer did not impact throughput, as the client buffer was quickly emptied. Ngtcp2 is also noted to be greedier than aioquic over DC, often introducing some unfairness to scenarios that were fair for aioquic. One potential reason is the difference in pacer implementation, as the IETF only recommends a pacer but does not specify it in detail. The difference in pacer implementation is also clearly shown in high delay ratio tests.

5.5 Bandwidth variability scenario

To capture a more realistic bandwidth user scenario, Figures 14(a) to 14(c) show repeated experiments with aioquic for DC batch size, high delay ratio and loss rates performed over a LTE sampled bandwidth trace. Figures 14(d) to 14(f) show these results but using ngtcp2 with CUBIC. For DC batch size and loss rates, similar trends are observed as when using a fixed bandwidth capability. Similar trends are also observed in the case of delay ratio, but with the effect of the pacer more clearly shown. Lastly, we note that ngtcp2 is more aggressive than aioquic and that CUBIC achieves higher throughput.

6 Conclusions

In this paper, we present the first performance study of QUIC over DC. Key insights are given for network operators to understand how different DC parameters and network conditions affect QUIC performance. QUIC's throughput is found to be similar to that of TCP in general cases, provided that the UDP receive buffer (when using aioquic) has been increased to similar size as the corresponding TCP buffer. We show that QUIC can take advantage of DC when the links share similar properties, and the DC batch size is small. When the properties of the links are too far apart, QUIC performance suffers to the degree that the performance would be better if DC was turned off. Furthermore, we show that QUIC can achieve system-wide fairness, provided that the link properties are similar. We also show that packet duplication allows QUIC to improve throughput for lossy environments at the cost of substantially increased unfairness.

With aioquic, the QUIC throughput is considerably lower if the UDP receive buffer remains at default values for Linux, as PDCP introduces packet bursts, causing packet drops due to full buffers. This occurs especially often in asymmetric link scenarios with high throughput. With the increased use of QUIC, we emphasize the importance of studying and optimizing the resources provided by the kernel to QUIC.

Acknowledgement

This work was funded in part by the Swedish Research Council (VR).

References

1. 3GPP: Study on Small Cell enhancements for E-UTRA and E-UTRAN; Higher layer aspects. Tech. Rep. 36.842 Release 12 (2013)
2. 3GPP: Summary of Rel-15. Tech. Rep. 21.915 Release 15 (2019)
3. 3GPP: Evolved Universal Terrestrial Radio Access; Packet Data Convergence Protocol specification. Tech. Rep. 36.323 Release 16 (2020)
4. aioquic: aioquic (2020), <https://github.com/aiortc/aioquic>
5. Alfredsson, S., Brunstrom, A., Sternad, M.: Cross-layer analysis of TCP performance in a 4G system. In: Proc. of SoftCOM (2007)
6. Becke, M., Dreibholz, T., Adhari, H., Rathgeb, E.P.: On the fairness of transport protocols in a multi-path environment. In: Proc. of IEEE ICC (2012)
7. Da Silva, I., Mildh, G., Rune, J., Wallentin, P., Vikberg, J., Schliwa-Bertling, P., Fan, R.: Tight integration of new 5G air interface and LTE to fulfill 5G requirements. In: Proc. of VTC Spring (2015)
8. De Coninck, Q., Bonaventure, O.: Multipath QUIC: Design and Evaluation. In: Proc. of ACM CoNEXT (2017)
9. Gurtov, A., Polishchuk, T.: Secure multipath transport for legacy Internet applications. In: Proc. of IEEE Broadnets (2009)
10. IETF 106 Singapore: Some updates on QUIC deployment numbers (2019), <https://datatracker.ietf.org/meeting/106/materials/slides-106-maprg-quic-deployment-update>

11. Iyengar, J., Swett, I.: QUIC Loss Detection and Congestion Control. Internet-Draft draft-ietf-quic-recovery-29, IETF (2020)
12. Iyengar, J.R., Amer, P.D., Stewart, R.: Concurrent multipath transfer using SCTP multihoming over independent end-to-end paths. *IEEE/ACM Trans. on Networking* (2006)
13. Jain, R.K., Chiu, D.M.W., Hawe, W.R.: A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Tech. Rep. DEC-TR-301, Eastern Research Lab, Digital Equipment Corporation (1984)
14. Jin, B., Kim, S., Yun, D., Lee, H., Kim, W., Yi, Y.: Aggregating LTE and Wi-Fi: Toward Intra-Cell Fairness and High TCP Performance. *IEEE Trans. on Wireless Communications* (2017)
15. Khadraoui, Y., Lagrange, X., Gravey, A.: TCP Performance for Practical Implementation of Very Tight Coupling between LTE and WiFi. In: *Proc. of IEEE VTC Fall* (2016)
16. Langley, A., Riddoch, A., Wilk, A., Vicente, A., Krasic, C., Zhang, D., Yang, F., Kouranov, F., Swett, I., Iyengar, J., Bailey, J., Dorfman, J., Roskind, J., Kulik, J., Westin, P., Tenneti, R., Shade, R., Hamilton, R., Vasiliev, V., Chang, W.T., Shi, Z.: The QUIC Transport Protocol: Design and Internet-Scale Deployment. In: *Proc. of ACM SIGCOMM* (2017)
17. Mahmood, N.H., Lopez, M., Laselva, D., Pedersen, K., Berardinelli, G.: Reliability Oriented Dual Connectivity for URLLC services in 5G New Radio. In: *Proc. of ISWCS* (2018)
18. McMillan, K.L., Zuck, L.D.: Formal Specification and Testing of QUIC. In: *Proc. of ACM SIGCOMM* (2019)
19. Mogensen, R.S., Markmoller, C., Madsen, T.K., Kolding, T., Pocovi, G., Lauridsen, M.: Selective Redundant MP-QUIC for 5G Mission Critical Wireless Applications. In: *Proc. of IEEE VTC Spring* (2019)
20. ngtcp2: ngtcp2 (2020), <https://github.com/ngtcp2/ngtcp2>
21. Paasch, C., Khalili, R., Bonaventure, O.: On the Benefits of Applying Experimental Design to Improve Multipath TCP. In: *Proc. of ACM CoNEXT* (2013)
22. Polese, M., Mezzavilla, M., Rangan, S., Zorzi, M.: Mobility Management for TCP in MmWave Networks. In: *Proc. ACM mmNets* (2017)
23. Rabitsch, A., Hurtig, P., Brunstrom, A.: A Stream-Aware Multipath QUIC Scheduler for Heterogeneous Paths. In: *Proc. of ACM SIGCOMM workshop EPIQ* (2018)
24. Raca, D., Quinlan, J.J., Zahran, A.H., Sreenan, C.J.: Beyond Throughput: A 4G LTE Dataset with Channel and Context Metrics. In: *Proc. of ACM MMSys* (2018)
25. Raiciu, C., Pluntke, C., Barre, S., Greenhalgh, A., Wischik, D., Handley, M.: Data center networking with multipath TCP. In: *Proc. of ACM SIGCOMM workshop HotNets* (2010)
26. Ravanshid, A., Rost, P., Michalopoulos, D.S., Phan, V.V., Bakker, H., Aziz, D., Tayade, S., Schotten, H.D., Wong, S., Holland, O.: Multi-connectivity functional architectures in 5G. In: *Proc. of IEEE ICC* (2016)
27. Wischik, D., Raiciu, C., Greenhalgh, A., Handley, M.: Design, Implementation and Evaluation of Congestion Control for Multipath TCP. In: *Proc. of USENIX Symposium on NSDI* (2011)
28. Zhang, X., Li, B.: Dice: A Game Theoretic Framework for Wireless Multipath Network Coding. In: *Proc. of ACM MobiHoc* (2008)