# Poster: Can MPTCP Improve Performance for Dual-Band 60 GHz/5 GHz Clients?

Swetank Kumar Saha[1], Roshan Shyamsunder[1], Naveen Muralidhar Prakash[1], Hany Assasa[2,3], Adrian Loch[2], Dimitrios Koutsonikolas[1], Joerg Widmer[2]
[1]University at Buffalo, SUNY, Buffalo NY, USA, [2]IMDEA Networks Institute, Madrid, Spain, [3]Universidad Carlos III de Madrid, Spain
{swetankk,roshansh,naveenmu,dimitrio}@buffalo.edu,{hany.assasa,adrian.loch,joerg.widmer}@imdea.org

## ABSTRACT

This work conducts one of the first experimental studies of Multipath TCP (MPTCP) in dual-band 60 GHz/5 GHz WLANs using off-the-shelf hardware. We consider both uncoupled and different coupled congestion control algorithms, compare their performance and their potential to improve throughput over single path TCP, and uncover their limitations. In contrast to a recent study that reports reduced throughput with MPTCP compared to single path TCP over 60 GHz, our results show that significant performance improvements are possible, especially in the case of uncoupled congestion control. On the other hand, performance gains with coupled congestion control are lower as these algorithms often fail to fully utilize the capacity of both paths simultaneously. We also observe a pathological case that can lead to significantly reduced throughput with MPTCP regardless of the congestion control algorithm.

## 1 INTRODUCTION

The almost 7 GHz of unlicensed spectrum centered around 60 GHz has attracted ample attention from both academia and industry as a solution for providing multi-gigabit indoor WLAN connectivity. Over the last couple of years, devices compliant with IEEE 802.11ad such as access points and laptops have been released commercially. In the future, 802.11ad (and eventually 802.11ay) devices will likely become as ubiquitous as legacy WiFi WLANs.

To overcome the high attenuation at millimeter wavelengths, 60 GHz radios use directional communication, typically enabled via phased array antennas. Directional communication, however, introduces new challenges as narrow beams are highly fragile under mobility and human blockage, and current beam training algorithms typically take 100s of milliseconds to converge. While several recent works have explored ways to reduce beamforming overhead (e.g., [3, 5, 12, 13]), client mobility and human blockage are still considered major challenges to the realization of 60 GHz WLANs.

The recent advent of tri-band chipsets (supporting 2.4, 5, and 60 GHz) by major chipset manufacturers [8] offers an attractive solution to mitigate these challenges; one can leverage the legacy WiFi interfaces (2.4/5 GHz) to maintain connectivity in cases 60 GHz connectivity is lost. In fact, the 802.11ad standard supports an optional Fast Session Transfer (FST) feature that allows traffic to migrate from 60 GHz to legacy WiFi and vice versa in a way transparent to higher layer protocols. Even more attractive is the possibility of using both interfaces (60 GHz and legacy WiFi) simultaneously when both networks are available, while seamlessly falling back to WiFi when the 60 GHz network becomes unavailable. This is possible via Multipath TCP (MPTCP) [2], a transport layer protocol that allows applications to exchange data over multiple paths, without any modification.

However, caution is needed in using MPTCP in practice. In fact, a large number of studies in datacenters and smartphones (e.g. [10], [1]) have shown that the protocol performs poorly over heterogeneous paths, due to interactions among out of order TCP packets, congestion control, and limited receiver buffer size. Given the large bandwidth disparity between WiFi and 60 GHz interfaces, it is likely that MPTCP will suffer from similar problems in hybrid 60 GHz/5 GHz WLANs. Hence, it is important to study the use of MPTCP in such WLANs, understand its potential to improve performance and robustness, and uncover its pitfalls and limitations.

Surprisingly, this is an aspect of 60 GHz networking which has remained largely unexplored. To our best knowledge, the only work that has briefly studied MPTCP performance in 60 GHz WLANs is [11]; it reports that using the 60 GHz interface alone achieved 7%-45% higher throughput compared to MPTCP but it does not explore the causes of the observed performance. The only other related work is [7], which explores the use of MPTCP in 5G cellular networks, over 28 GHz and LTE using simulations. The authors find that MPTCP with uncoupled congestion control performs better than single path TCP (SPTCP) by up to 100%; on the other hand, MPTCP with the recently proposed coupled congestion control *balia* algorithm [6] in many cases performs worse than SPTCP. Nonetheless, the use of a simulator instead of experiments with real hardware and the fundamental differences between outdoor cellular networks and indoor WLANs make it unclear if the these findings are still valid in the context of 60 GHz WLANs.

This work fills this gap by conducting an experimental study of the MPTCP performance in dual-band 802.11ac/802.11ad WLANs using off-the-shelf hardware. Our study considers both uncoupled and coupled congestion control and different Tx-Rx distances. In contrast to the results in [11], our study shows that MPTCP with uncoupled congestion control can significantly boost performance, often yielding throughput equal to the sum of the SPTCP throughputs over the two interfaces. On the other hand, performance with recently proposed coupled congestion control algorithms is typically much lower. We also found a few cases where large RTT inflation over both subflows results in significantly reduced throughput,


*MobiCom '17, October 16–20, 2017, Snowbird, UT, USA*

ACM ISBN 978-1-4503-4916-1/17/10.
https://doi.org/10.1145/3117811.3131248

even lower than the throughput of SPTCP over the best path. We are currently investigating the causes of this.

# 2 EXPERIMENTAL METHODOLOGY

## 2.1 Devices

Our setup consists of a Netgear Nighthawk® X10 Smart WiFi Router and an Acer Travelmate P446-M laptop. Both devices have chipset(s) supporting 802.11ac and 802.11ad.

**60 GHz (802.11ad)** The router has the QCA9008-SBD1 module housing the QCA9500 chipset from Qualcomm, supporting all the single-carrier 802.11ad data rates (from 385 Mbps up to 4.6 Gbps). The laptop carries the client-version of the module: QCA9008-TBD1, which houses the 802.11ac, 802.11ad and BT chipsets. It runs a typical Linux OS (kernel 4.x) and uses the open source wil6210 wireless driver to interface with the chipset. Both the router and the laptop use a 32-element phased antenna array on a separate chipset and connected to the main chipset using a MHF4 cable.

The 60 GHz radios on both devices use their own rate adaptation and beamforming algorithms to select MCS and control beam properties, respectively. In case the link is blocked, the radios automatically search for an alternative NLOS path through a reflection to re-establish the connection. On the laptop, the wil6210 driver exports detailed connection parameters, including Tx and Rx MCS, MAC layer throughput, signal quality indicator (SQI), beamforming (BF) status (OK/Failed/Retrying), and sectors in use both by itself and the AP. We log all the parameters every 150 ms.

**WiFi (802.11ac)** Although the router supports up to 4 MIMO spatial streams, our client only supports 2. Hence, our effective link configuration for WiFi is 802.11ac, 2x2 MIMO, 80 MHz, SGI, and default rate adaptation.

**Traffic Generation and Maximum Goodput** A high-end desktop is connected to the router through a 10G LAN SFP+ interface to generate/receive TCP traffic. This setup should ideally allow us to take achieve multi-Gigabit speeds, we found that the maximum goodput is limited to ~2.3 Gbps and ~550 Mbps, with 802.11ad and 802.11ac, respectively.

## 2.2 MPTCP

We use MPTCP version v0.92 with the *fullmesh* path manager that creates a subflow over all available network interfaces for each established TCP connection, and the *default* RTT-based scheduler that sends packets over the available path with the lowest RTT. We experiment with both uncoupled congestion control (TCP Cubic) and three coupled congestion control algorithms, *lia* [9], *olia* [4], and *balia* [6]. We consider both interfaces for the primary subflow and the results are similar (which is expected given that the RTT is similar for both interfaces). In Section 3, we only present results with the primary subflow over the 60 GHz interface.

## 2.3 Environment and Methodology

Our experiments are conducted in an open Lobby thinly populated by some desks and chairs. The ceiling height is rather high and thus it does not serve as a viable reflector. We select this location to emulate near-free space propagation. Each experiment consists of a 10-second backlogged iperf3 TCP session from the desktop behind the router to the laptop. All the results are the average of 5 sessions. We further capture all the packets on the receiver side using tcpdump and calculate the throughput of each subflow over each 100-ms interval in order to study the evolution of throughput over time for each run and the interaction between the two subflows.

All our measurements are done during night time to make sure there is no WiFi (802.11ac) interference from other users in the building. Further, we ensure that the entire 80 MHz bandwidth is indeed available exclusively to our link by continuously checking, in between our actual measurement runs, that the maximum throughput (with 2x2 MIMO, 80 MHz, SGI, and high SNR) is achievable when using single-path TCP.

# 3 RESULTS

Figure 1a compares the performance of different versions of MPTCP against SPTCP for two Tx-Rx distances. In contrast to the observations in [11], we observe that *MPTCP has the potential to improve performance compared to SPTCP over 60 GHz.* The average throughput improvement of MPTCP with *cubic*, *lia*, *olia*, and *balia* compared to SPTCP over 60 GHz is 29%, 9%, 9%, and 2%, respectively at 10 ft, and 32%, 27%, 5%, and 14%, respectively, at 70 ft. Note that these performance differences are not due to different channel conditions or due to different MAC layer decisions over different runs. We observed that the same sector was used 100% of the time at each experiment at both distances. Further, MCS 8 was used at least 99% of the time in all experiments at 10 ft and MCS 6 was used 88-90% of the time in all experiments at 70 ft. Hence, any performance differences can be attributed to decisions made at the transport layer (congestion control and scheduler).

Interestingly, MPTCP with uncoupled congestion control performs the best at both distances, achieving a throughput roughly equal to the sum of the SPTCP throughputs over WiFi and 60 GHz alone. Figures 1b, 1c show two examples of the evolution of throughput over time separately for each interface in the case of MPTCP with uncoupled congestion control, at 10 ft and 70 ft, respectively. This *ideal* behavior is very similar for 4 out of 5 runs with uncoupled congestion control at each of the two distances. In spite of this promising result (which is similar to the result in [7] for 5G cellular networks over LTE and 28 GHz), we note that previous studies [4, 6, 9] showed that MPTCP with uncoupled congestion control can be unfriendly to SPTCP flows and proposed coupled congestion control algorithms to address this issue.

Nonetheless, Figure 1a shows that the performance improvements with coupled congestion control are in most cases much lower than with uncoupled congestion control; in fact, the most recently proposed *balia* algorithm performs the worst among the three algorithms at 10 ft. The *ideal* behavior is observed only 1, 0, 0 out of 5 times with *lia*, *olia*, and *balia*, respectively, at 10 ft, and 3, 0, and 1 out of 5 times, respectively, at 70 ft. In contrast, the most common behavior with all three algorithms is the one shown in Figures 1d and 1e. At 10 ft (Figure 1d), the throughput of the 60 GHz subflow is initially lower than its maximum value (below 1.5 Gbps); after some time (between 0.5 and 4 s for different runs), it obtains its maximum value but at the same time, the throughput of the WiFi subflow drops to a lower level (100 Mbps in Figure 1d, anywhere between 50-400 Mbps in different runs). At 70 ft, this transition

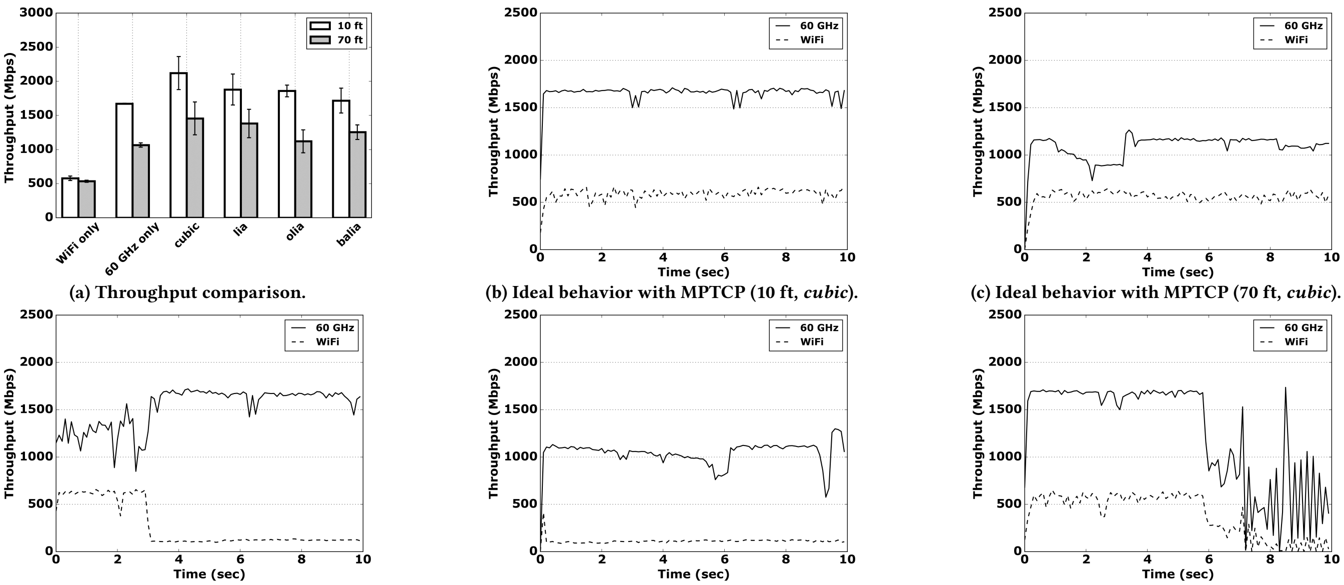


(a) Throughput comparison. (b) Ideal behavior with MPTCP (10 ft, *cubic*). (c) Ideal behavior with MPTCP (70 ft, *cubic*). (d) Reduced WiFi throughput (10 ft, *balia*). (e) Reduced WiFi throughput (70 ft, *lia*). (f) Throughput oscillation (10 ft, *cubic*).

**Figure 1: Experimental results.**

happens much faster (within 0.2 s), as shown in Figure 1c. In other words, MPTCP with coupled congestion control fails to fully utilize the capacity of both paths simultaneously in most runs. The varying WiFi throughput over different runs and the different times at which the 60 GHz subflow achieves its maximum throughput are among the reasons for the large standard deviations observed for *lia*, *olia*, and *balia* in Figure 1a.

We are currently investigating the causes of this behavior. Note that the work in [7] also found *balia* to perform worse than *cubic* in 5G cellular networks because it perceives losses over the mmWave link as congestion, and consequently shifts all the traffic to the low-rate LTE link. In contrast, we do not observe any TCP timeouts or duplicate ACKs and the RTTs are similar over both subflows in our experiments exhibiting the behavior shown in Figures 1d, 1e.

Figure 1a shows that throughput with *cubic* also features large standard deviations, similar to *lia*, *olia*, and *balia*, even though *cubic* never exhibited the behavior shown in Figures 1d, 1e. On the other hand, note that the standard deviations for SPTCP performance over both interfaces are negligible, suggesting again a stable wireless channel. The reason for the large standard deviations with *cubic* is that, in one of the 5 runs for each distance, both subflows with *cubic* experienced significantly reduced throughput and very large oscillations over time for a fraction of the 10-second interval. One such example is shown in Figure 1f. We observed a similar behavior in one of the 5 runs with *lia* and *balia* at 10 ft and one run with *olia* at 70 ft. In all these runs, RTT is significantly inflated during the throughput oscillation intervals – up to 40-60 ms for the 60 GHz subflow and up to 150 ms for the WiFi subflow. In spite of the RTT inflation, in most of these runs there are no TCP retransmissions. Further, the RTT inflation is not caused due to ACKs being lost in the channel or delayed due to congestion. The receiver simply delays to send the ACKs, albeit without following the standard rule of one ACK for every two segments. In several cases, it receives multiple segments before sending an ACK.

In our future work, we plan to investigate the causes of the observed performance and look into the effectiveness and responsiveness of MPTCP in mobility and blockage scenarios. We also plan to study the impact of simultaneously using both interfaces on power consumption which will be important as smartphones become the next target of 802.11ad.

## REFERENCES


[1] Simone Ferlin, Thomas Dreibholz, and Özgu Alay. 2014. Multi-Path Transport Over Heterogeneous Wireless Networks: Does It Really Pay Off?. In *Proc. of IEEE Global Telecommunications Conference (GLOBECOM)*.

[2] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure. 2013. TCP Extensions for Multipath Operation with Multiple Addresses. (2013). RFC 1546.

[3] Muhammad Kumail Haider and Edward W. Knightly. 2016. Mobility Resilience and Overhead Constrained Adaptation in Directional 60 GHz WLANs: Protocol Design and System Implementation. In *Proc. of ACM MobiHoc*.

[4] Ramin Khalili, Nicolas Gast, Miroslav Popovic, and Jean-Yves Le Boudec. 2013. MPTCP Is Not Pareto-Optimal: Performance Issues and a Possible Solution. *IEEE/ACM Transactions on Networking* 21, 5 (2013).

[5] Thomas Nitsche, Adriana B. Flores, Edward W. Knightly, and Joerg Widmer. 2015. Steering with Eyes Closed: mm-Wave Beam Steering without In-Band Measurement. In *Proc. of IEEE INFOCOM*.

[6] Qiuyu Peng, Anwar Walid, Jaehyun Hwang, and Steven H. Low. 2016. Multipath TCP: Analysis, Design, and Implementation. *IEEE/ACM Transactions on Networking* 24, 1 (2016).

[7] Michele Polese, Rittwik Jana, and Michele Zorzi. 2017. TCP in 5G mmWave Networks: Link Level Retransmissions and MP-TCP. In *Proc. of IEEE Workshop on 5G New Radio Technologies*.

[8] Qualcomm. 2017. Qualcomm Technologies' Tri-Band Solution. https://www.qualcomm.com/products/features/80211ad. (2017).

[9] Costin Raiciu, Mark Handley, and Damon Wischik. 2011. Coupled Congestion Control for Multipath Transport Protocols. (2011). RFC 6356.

[10] Costin Raiciu, Christoph Paasch, Sebastien Barre, Alan Ford, Michio Honda, Fabien Duchene, Olivier Bonaventure, and Mark Handley. 2012. How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP. In *Proc. of USENIX NSDI*.

[11] Sanjib Sur, Ioannis Pefkianakis, Xinyu Zhang, and Kyu-Han Kim. 2017. WiFi-Assisted 60 GHz Wireless Networks. In *Proc. of ACM MobiCom*.

[12] Sanjib Sur, Xinyu Zhang, Parameswaran Ramanathan, and Ranveer Chandra. 2016. BeamSpy: Enabling Robust 60 GHz Links Under Blockage. In *Proc. of USENIX NSDI*.

[13] Anfu Zhou, Xinyu Zhang, and Huadong Ma. 2017. Beam-forecast: Facilitating Mobile 60 GHz Networks via Model-driven Beam Steering. In *Proc. of IEEE INFOCOM*.