
Stream: Internet Engineering Task Force (IETF)
RFC: [8869](#)
Category: Informational
Published: January 2021
ISSN: 2070-1721
Authors: Z. Sarker X. Zhu J. Fu
Ericsson AB Cisco Systems Cisco Systems

RFC 8869

Evaluation Test Cases for Interactive Real-Time Media over Wireless Networks

Abstract

The Real-time Transport Protocol (RTP) is a common transport choice for interactive multimedia communication applications. The performance of these applications typically depends on a well-functioning congestion control algorithm. To ensure a seamless and robust user experience, a well-designed RTP-based congestion control algorithm should work well across all access network types. This document describes test cases for evaluating performances of candidate congestion control algorithms over cellular and Wi-Fi networks.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8869>.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions

with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction
2. Cellular Network Specific Test Cases
 - 2.1. Varying Network Load
 - 2.1.1. Network Connection
 - 2.1.2. Simulation Setup
 - 2.1.3. Expected Behavior
 - 2.2. Bad Radio Coverage
 - 2.2.1. Network Connection
 - 2.2.2. Simulation Setup
 - 2.2.3. Expected Behavior
 - 2.3. Desired Evaluation Metrics for Cellular Test Cases
3. Wi-Fi Networks Specific Test Cases
 - 3.1. Bottleneck in Wired Network
 - 3.1.1. Network Topology
 - 3.1.2. Test/Simulation Setup
 - 3.1.3. Typical Test Scenarios
 - 3.1.4. Expected Behavior
 - 3.2. Bottleneck in Wi-Fi Network
 - 3.2.1. Network Topology
 - 3.2.2. Test/Simulation Setup
 - 3.2.3. Typical Test Scenarios
 - 3.2.4. Expected Behavior
 - 3.3. Other Potential Test Cases
 - 3.3.1. EDCA/WMM usage
 - 3.3.2. Effect of Heterogeneous Link Rates

- [4. IANA Considerations](#)
- [5. Security Considerations](#)
- [6. References](#)
 - [6.1. Normative References](#)
 - [6.2. Informative References](#)

[Contributors](#)

[Acknowledgments](#)

[Authors' Addresses](#)

1. Introduction

Wireless networks (both cellular and Wi-Fi [IEEE802.11]) are an integral and increasingly more significant part of the Internet. Typical application scenarios for interactive multimedia communication over wireless include video conferencing calls in a bus or train as well as live media streaming at home. It is well known that the characteristics and technical challenges for supporting multimedia services over wireless are very different from those of providing the same service over a wired network. Although the basic test cases as defined in [RFC8867] have covered many common effects of network impairments for evaluating RTP-based congestion control schemes, they remain to be tested over characteristics and dynamics unique to a given wireless environment. For example, in cellular networks, the base station maintains individual queues per radio bearer per user hence it leads to a different nature of interactions between traffic flows of different users. This contrasts with a typical wired network setting where traffic flows from all users share the same queue at the bottleneck. Furthermore, user mobility patterns in a cellular network differ from those in a Wi-Fi network. Therefore, it is important to evaluate the performance of proposed candidate RTP-based congestion control solutions over cellular mobile networks and over Wi-Fi networks respectively.

[RFC8868] provides guidelines for evaluating candidate algorithms and recognizes the importance of testing over wireless access networks. However, it does not describe any specific test cases for performance evaluation of candidate algorithms. This document describes test cases specifically targeting cellular and Wi-Fi networks.

2. Cellular Network Specific Test Cases

A cellular environment is more complicated than its wireline counterpart since it seeks to provide services in the context of variable available bandwidth, location dependencies, and user mobilities at different speeds. In a cellular network, the user may reach the cell edge, which may lead to a significant number of retransmissions to deliver the data from the base station to the destination and vice versa. These radio links will often act as a bottleneck for the rest of the

network and will eventually lead to excessive delays or packet drops. An efficient retransmission or link adaptation mechanism can reduce the packet loss probability, but some packet losses and delay variations will remain. Moreover, with increased cell load or handover to a congested cell, congestion in the transport network will become even worse. Besides, there exist certain characteristics that distinguish the cellular network from other wireless access networks such as Wi-Fi. In a cellular network:

- The bottleneck is often a shared link with relatively few users.
 - The cost per bit over the shared link varies over time and is different for different users.
 - Leftover/unused resources can be consumed by other greedy users.
- Queues are always per radio bearer, hence each user can have many such queues.
- Users can experience both inter- and intra-Radio Access Technology (RAT) handovers (see [HO-def-3GPP] for the definition of "handover").
- Handover between cells or change of serving cells (as described in [HO-LTE-3GPP] and [HO-UMTS-3GPP]) might cause user plane interruptions, which can lead to bursts of packet losses, delay, and/or jitter. The exact behavior depends on the type of radio bearer. Typically, the default best-effort bearers do not generate packet loss, instead, packets are queued up and transmitted once the handover is completed.
- The network part decides how much the user can transmit.
- The cellular network has variable link capacity per user.
 - It can vary as fast as a period of milliseconds.
 - It depends on many factors (such as distance, speed, interference, different flows).
 - It uses complex and smart link adaptation, which makes the link behavior ever more dynamic.
 - The scheduling priority depends on the estimated throughput.
- Both Quality of Service (QoS) and non-QoS radio bearers can be used.

Hence, a real-time communication application operating over a cellular network needs to cope with a shared bottleneck link and variable link capacity, events like handover, non-congestion-related loss, and abrupt changes in bandwidth (both short term and long term) due to handover, network load, and bad radio coverage. Even though 3GPP has defined QoS bearers [QoS-3GPP] to ensure high-quality user experience, it is still preferable for real-time applications to behave in an adaptive manner.

Different mobile operators deploy their own cellular networks with their own set of network functionalities and policies. Usually, a mobile operator network includes a range of radio access technologies such as 3G and 4G/LTE. Looking at the specifications of such radio technologies, it is evident that only the more recent radio technologies can support the high bandwidth requirements from real-time interactive video applications. Future real-time interactive applications will impose even greater demand on cellular network performance, which makes 4G (and beyond) radio technologies more suitable for such genre of application.

The key factors in defining test cases for cellular networks are:

- Shared and varying link capacity
- Mobility
- Handover

However, these factors are typically highly correlated in a cellular network. Therefore, instead of devising separate test cases for individual important events, we have divided the test cases into two categories. It should be noted that the goal of the following test cases is to evaluate the performance of candidate algorithms over the radio interface of the cellular network. Hence, it is assumed that the radio interface is the bottleneck link between the communicating peers and that the core network does not introduce any extra congestion along the path. Consequently, this document has left out of scope the combination of multiple access technologies involving both cellular and Wi-Fi users. In this latter case, the shared bottleneck is likely at the wired backhaul link. These test cases further assume a typical real-time telephony scenario where one real-time session consists of one voice stream and one video stream.

Even though it is possible to carry out tests over operational cellular networks (e.g., LTE/5G), and actually such tests are already available today, these tests cannot in general be carried out in a deterministic fashion to ensure repeatability. The main reason is that these networks are controlled by cellular operators, and there exists various amounts of competing traffic in the same cell(s). In practice, it is only in underground mines that one can carry out near deterministic testing. Even there, it is not guaranteed either as workers in the mines may carry with them their personal mobile phones. Furthermore, the underground mining setting may not reflect typical usage patterns in an urban setting. We, therefore, recommend that a cellular network simulator be used for the test cases defined in this document, for example -- the LTE simulator in [NS-3].

2.1. Varying Network Load

The goal of this test is to evaluate the performance of the candidate congestion control algorithm under varying network load. The network load variation is created by adding and removing network users, a.k.a. User Equipment (UE), during the simulation. In this test case, each user/UE in the media session is an endpoint following RTP-based congestion control. User arrivals follow a Poisson distribution proportional to the length of the call, to keep the number of users per cell fairly constant during the evaluation period. At the beginning of the simulation, there should be enough time to warm up the network. This is to avoid running the evaluation in an empty network where network nodes have empty buffers and low interference at the beginning of the simulation. This network initialization period should be excluded from the evaluation period. Typically, the evaluation period starts 30 seconds after test initialization.

This test case also includes user mobility and some competing traffic. The latter includes both the same types of flows (with same adaptation algorithms) and different types of flows (with different services and congestion control schemes).

2.1.1. Network Connection

Each mobile user is connected to a fixed user. The connection between the mobile user and fixed user consists of a cellular radio access, an Evolved Packet Core (EPC), and an Internet connection. The mobile user is connected to the EPC using cellular radio access technology, which is further connected to the Internet. At the other end, the fixed user is connected to the Internet via a wired connection with sufficiently high bandwidth, for instance, 10 Gbps, so that the system bottleneck is on the cellular radio access interface. The wired connection in this setup does not introduce any network impairments to the test; it only adds 10 ms of one-way propagation delay.

The path from the fixed user to the mobile users is defined as "downlink", and the path from the mobile users to the fixed user is defined as "uplink". We assume that only uplink or downlink is congested for mobile users. Hence, we recommend that the uplink and downlink simulations are run separately.




Figure 1: Simulation Topology

2.1.2. Simulation Setup

The values enclosed within "[]" for the following simulation attributes follow the same notion as in [RFC8867]. The desired simulation setup is as follows:

Radio environment:

Deployment and propagation model: 3GPP case 1 (see [HO-deploy-3GPP])

Antenna: Multiple-Input and Multiple-Output (MIMO), 2D or 3D antenna pattern

Mobility: [3 km/h, 30 km/h]

Transmission bandwidth: 10 MHz

Number of cells: multi-cell deployment (3 cells per Base Station (BS) * 7 BS) = 21 cells

Cell radius: 166.666 meters

Scheduler: Proportional fair with no priority

Bearer: Default bearer for all traffic

Active Queue Management (AQM) settings: AQM [on, off]

End-to-end Round Trip Time (RTT): [40 ms, 150 ms]

User arrival model: Poisson arrival model

User intensity:

Downlink user intensity: {0.7, 1.4, 2.1, 2.8, 3.5, 4.2, 4.9, 5.6, 6.3, 7.0, 7.7, 8.4, 9.1, 9.8, 10.5}

Uplink user intensity: {0.7, 1.4, 2.1, 2.8, 3.5, 4.2, 4.9, 5.6, 6.3, 7.0}

Simulation duration: 91 s

Evaluation period: 30 s - 60 s

Media traffic:

Media type: Video

Media direction: [uplink, downlink]

Number of media sources per user: One (1)

Media duration per user: 30 s

Media source: same as defined in [Section 4.3](#) of [\[RFC8867\]](#)

Media type: Audio

Media direction: [uplink, downlink]

Number of media sources per user: One (1)

Media duration per user: 30 s

Media codec: Constant Bit Rate (CBR)

Media bitrate: 20 Kbps

Adaptation: off

Other traffic models:

Downlink simulation: Maximum of 4 Mbps/cell (web browsing or FTP traffic following default TCP congestion control [\[RFC5681\]](#))

Uplink simulation: Maximum of 2 Mbps/cell (web browsing or FTP traffic following default TCP congestion control [\[RFC5681\]](#))

2.1.3. Expected Behavior

The investigated congestion control algorithms should result in maximum possible network utilization and stability in terms of rate variations, lowest possible end-to-end frame latency, network latency, and Packet Loss Rate (PLR) at different cell load levels.

2.2. Bad Radio Coverage

The goal of this test is to evaluate the performance of the candidate congestion control algorithm when users visit part of the network with bad radio coverage. The scenario is created by using a larger cell radius than that in the previous test case. In this test case, each user/UE in the media session is an endpoint following RTP-based congestion control. User arrivals follow a Poisson distribution proportional to the length of the call, to keep the number of users per cell fairly constant during the evaluation period. At the beginning of the simulation, there should be enough time to warm up the network. This is to avoid running the evaluation in an empty network where network nodes have empty buffers and low interference at the beginning of the simulation. This network initialization period should be excluded from the evaluation period. Typically, the evaluation period starts 30 seconds after test initialization.

This test case also includes user mobility and some competing traffic. The latter includes the same kind of flows (with same adaptation algorithms).

2.2.1. Network Connection

Same as defined in [Section 2.1.1](#).

2.2.2. Simulation Setup

The desired simulation setup is the same as the Varying Network Load test case defined in [Section 2.1](#) except for the following changes:

Radio environment: Same as defined in [Section 2.1.2](#) except for the following:

Deployment and propagation model: 3GPP case 3 (see [[HO-deploy-3GPP](#)])

Cell radius: 577.3333 meters

Mobility: 3 km/h

User intensity: {0.7, 1.4, 2.1, 2.8, 3.5, 4.2, 4.9, 5.6, 6.3, 7.0}

Media traffic model: Same as defined in [Section 2.1.2](#)

Other traffic models:

Downlink simulation: Maximum of 2 Mbps/cell (web browsing or FTP traffic following default TCP congestion control [[RFC5681](#)])

Uplink simulation: Maximum of 1 Mbps/cell (web browsing or FTP traffic following default TCP congestion control [[RFC5681](#)])

2.2.3. Expected Behavior

The investigated congestion control algorithms should result in maximum possible network utilization and stability in terms of rate variations, lowest possible end-to-end frame latency, network latency, and Packet Loss Rate (PLR) at different cell load levels.

2.3. Desired Evaluation Metrics for Cellular Test Cases

The evaluation criteria document [RFC8868] defines the metrics to be used to evaluate candidate algorithms. Considering the nature and distinction of cellular networks, we recommend that at least the following metrics be used to evaluate the performance of the candidate algorithms:

- Average cell throughput (for all cells), shows cell utilization.
- Application sending and receiving bitrate, goodput.
- Packet Loss Rate (PLR).
- End-to-end media frame delay. For video, this means the delay from capture to display.
- Transport delay.
- Algorithm stability in terms of rate variation.

3. Wi-Fi Networks Specific Test Cases

Given the prevalence of Internet access links over Wi-Fi, it is important to evaluate candidate RTP-based congestion control solutions over test cases that include Wi-Fi access links. Such evaluations should highlight the inherently different characteristics of Wi-Fi networks in contrast to their wired counterparts:

- The wireless radio channel is subject to interference from nearby transmitters, multipath fading, and shadowing. These effects lead to fluctuations in the link throughput and sometimes an error-prone communication environment.
- Available network bandwidth is not only shared over the air between concurrent users but also between uplink and downlink traffic due to the half-duplex nature of the wireless transmission medium.
- Packet transmissions over Wi-Fi are susceptible to contentions and collisions over the air. Consequently, traffic load beyond a certain utilization level over a Wi-Fi network can introduce frequent collisions over the air and significant network overhead, as well as packet drops due to buffer overflow at the transmitters. This, in turn, leads to excessive delay, retransmissions, packet losses, and lower effective bandwidth for applications. Note further that the collision-induced delay and loss patterns are qualitatively different from those caused by congestion over a wired connection.
- The IEEE 802.11 standard (i.e., Wi-Fi) supports multi-rate transmission capabilities by dynamically choosing the most appropriate modulation and coding scheme (MCS) for the given received signal strength. A different choice in the MCS Index leads to different physical-layer (PHY-layer) link rates and consequently different application-layer throughput.
- The presence of legacy devices (e.g., ones operating only in IEEE 802.11b) at a much lower PHY-layer link rate can significantly slow down the rest of a modern Wi-Fi network. As discussed in [Heusse2003], the main reason for such anomaly is that it takes much longer to transmit the same packet over a slower link than over a faster link, thereby consuming a substantial portion of air time.

- Handover from one Wi-Fi Access Point (AP) to another may lead to excessive packet delays and losses during the process.
- IEEE 802.11e has introduced the Enhanced Distributed Channel Access (EDCA) mechanism to allow different traffic categories to contend for channel access using different random back-off parameters. This mechanism is a mandatory requirement for the Wi-Fi Multimedia (WMM) certification in Wi-Fi Alliance. It allows for prioritization of real-time application traffic such as voice and video over non-urgent data transmissions (e.g., file transfer).

In summary, the presence of Wi-Fi access links in different network topologies can exert different impacts on the network performance in terms of application-layer effective throughput, packet loss rate, and packet delivery delay. These, in turn, will influence the behavior of end-to-end real-time multimedia congestion control.

Unless otherwise mentioned, the test cases in this section choose the PHY- and MAC-layer parameters based on the IEEE 802.11n standard. Statistics collected from enterprise Wi-Fi networks show that the two dominant physical modes are 802.11n and 802.11ac, accounting for 41% and 58% of connected devices, respectively. As Wi-Fi standards evolve over time -- for instance, with the introduction of the emerging Wi-Fi 6 (based on IEEE 802.11ax) products -- the PHY- and MAC-layer test case specifications need to be updated accordingly to reflect such changes.

Typically, a Wi-Fi access network connects to a wired infrastructure. Either the wired or the Wi-Fi segment of the network can be the bottleneck. The following sections describe basic test cases for both scenarios separately. The same set of performance metrics as in [RFC8867] should be collected for each test case.

We recommend carrying out the test cases as defined in this document using a simulator, such as [NS-2] or [NS-3]. When feasible, it is encouraged to perform testbed-based evaluations using Wi-Fi access points and endpoints running up-to-date IEEE 802.11 protocols, such as 802.11ac and the emerging Wi-Fi 6, so as to verify the viability of the candidate schemes.

3.1. Bottleneck in Wired Network

The test scenarios below are intended to mimic the setup of video conferencing over Wi-Fi connections from the home. Typically, the Wi-Fi home network is not congested, and the bottleneck is present over the wired home access link. Although it is expected that test evaluation results from this section are similar to those in [RFC8867], it is still worthwhile to run through these tests as sanity checks.

3.1.1. Network Topology

Figure 2 shows the network topology of Wi-Fi test cases. The test contains multiple mobile nodes (MNs) connected to a common Wi-Fi AP and their corresponding wired clients on fixed nodes (FNs). Each connection carries either an RTP-based media flow or a TCP traffic flow. Directions of the flows can be uplink (i.e., from mobile nodes to fixed nodes), downlink (i.e., from fixed nodes to mobile nodes), or bidirectional. The total number of uplink/downlink/bidirectional flows for RTP-based media traffic and TCP traffic are denoted as N and M , respectively.

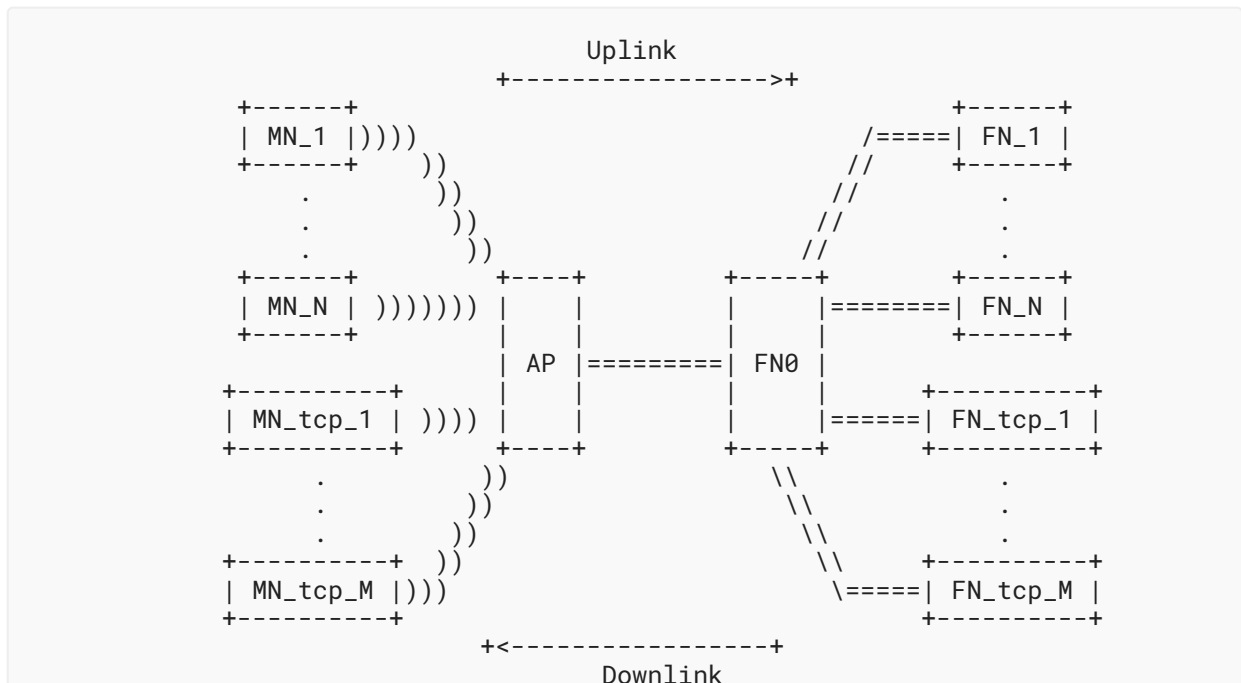


Figure 2: Network Topology for Wi-Fi Test Cases

3.1.2. Test/Simulation Setup

Test duration: 120 s

Wi-Fi network characteristics:

Radio propagation model: Log-distance path loss propagation model (see [NS3WiFi])

PHY- and MAC-layer configuration: IEEE 802.11n

MCS Index at 11: Raw data rate at 52 Mbps, 16-QAM (Quadrature amplitude modulation) and 1/2 coding rate

Wired path characteristics:

Path capacity: 1 Mbps

One-way propagation delay: 50 ms

Maximum end-to-end jitter: 30 ms

Bottleneck queue type: Drop tail

Bottleneck queue size: 300 ms

Path loss ratio: 0%

Application characteristics:

Media traffic:

Media type: Video

Media direction: See [Section 3.1.3](#)

Number of media sources (N): See [Section 3.1.3](#)

Media timeline:

Start time: 0 s

End time: 119 s

Competing traffic:

Type of sources: Long-lived TCP or CBR over UDP

Traffic direction: See [Section 3.1.3](#)

Number of sources (M): See [Section 3.1.3](#)

Congestion control: Default TCP congestion control [[RFC5681](#)] or CBR traffic over UDP

Traffic timeline: See [Section 3.1.3](#)

3.1.3. Typical Test Scenarios

Single uplink RTP-based media flow: N=1 with uplink direction and M=0.

One pair of bidirectional RTP-based media flows: N=2 (i.e., one uplink flow and one downlink flow); M=0.

One pair of bidirectional RTP-based media flows: N=2; one uplink on-off CBR flow over UDP: M=1 (uplink). The CBR flow has ON time at t=0s-60s and OFF time at t=60s-119s.

One pair of bidirectional RTP-based media flows: N=2; one uplink off-on CBR flow over UDP: M=1 (uplink). The CBR flow has OFF time at t=0s-60s and ON time at t=60s-119s.

One RTP-based media flow competing against one long-lived TCP flow in the uplink direction: N=1 (uplink) and M=1 (uplink). The TCP flow has start time at t=0s and end time at t=119s.

3.1.4. Expected Behavior

Single uplink RTP-based media flow: The candidate algorithm is expected to detect the path capacity constraint, to converge to the bottleneck link capacity, and to adapt the flow to avoid unwanted oscillations when the sending bit rate is approaching the bottleneck link capacity. No excessive oscillations in the media rate should be present.

Bidirectional RTP-based media flows: The candidate algorithm is expected to converge to the bottleneck capacity of the wired path in both directions despite the presence of measurement noise over the Wi-Fi connection. In the presence of background TCP or CBR over UDP traffic, the rate of RTP-based media flows should adapt promptly to the arrival and departure of background traffic flows.

One RTP-based media flow competing with long-lived TCP flow in the uplink direction: The candidate algorithm is expected to avoid congestion collapse and to stabilize at a fair share of the bottleneck link capacity.

3.2. Bottleneck in Wi-Fi Network

The test cases in this section assume that the wired segment along the media path is well-provisioned, whereas the bottleneck exists over the Wi-Fi access network. This is to mimic the application scenarios typically encountered by users in an enterprise environment or at a coffee house.

3.2.1. Network Topology

Same as defined in [Section 3.1.1](#).

3.2.2. Test/Simulation Setup

Test duration: 120 s

Wi-Fi network characteristics:

Radio propagation model: Log-distance path loss propagation model (see [[NS3WiFi](#)])

PHY- and MAC-layer configuration: IEEE 802.11n

MCS Index at 11: Raw data rate at 52 Mbps, 16-QAM (Quadrature amplitude modulation) and 1/2 coding rate

Wired path characteristics:

Path capacity: 100 Mbps

One-Way propagation delay: 50 ms

Maximum end-to-end jitter: 30 ms

Bottleneck queue type: Drop tail

Bottleneck queue size: 300 ms

Path loss ratio: 0%

Application characteristics

Media traffic:

Media type: Video

Media direction: See [Section 3.2.3](#)

Number of media sources (N): See [Section 3.2.3](#)

Media timeline:

Start time: 0 s

End time: 119 s

Competing traffic:

Type of sources: long-lived TCP or CBR over UDP

Number of sources (M): See [Section 3.2.3](#)

Traffic direction: See [Section 3.2.3](#)

Congestion control: Default TCP congestion control [[RFC5681](#)] or CBR traffic over UDP

Traffic timeline: See [Section 3.2.3](#)

3.2.3. Typical Test Scenarios

This section describes a few test scenarios that are deemed as important for understanding the behavior of a candidate RTP-based congestion control scheme over a Wi-Fi network.

Multiple RTP-based media flows sharing the wireless downlink: N=16 (all downlink); M=0. This test case is for studying the impact of contention on the multiple concurrent media flows. For an 802.11n network, given the MCS Index of 11 and the corresponding link rate of 52 Mbps, the total application-layer throughput (assuming reasonable distance, low interference, and infrequent contentions caused by competing streams) is around 20 Mbps. A total of N=16 RTP-based media flows (with a maximum rate of 1.5 Mbps each) are expected to saturate the wireless interface in this experiment. Evaluation of a given candidate scheme should focus on whether the downlink media flows can stabilize at a fair share of the total application-layer throughput.

Multiple RTP-based media flows sharing the wireless uplink: N=16 (all uplink); M=0. When multiple clients attempt to transmit media packets uplink over the Wi-Fi network, they introduce more frequent contentions and potential collisions. Per-flow throughput is expected to be lower than that in the previous downlink-only scenario. Evaluation of a given candidate scheme should focus on whether the uplink flows can stabilize at a fair share of the total application-layer throughput.

Multiple bidirectional RTP-based media flows: N=16 (8 uplink and 8 downlink); M=0. The goal of this test is to evaluate the performance of the candidate scheme in terms of bandwidth fairness between uplink and downlink flows.

Multiple bidirectional RTP-based media flows with on-off CBR traffic over UDP: N=16 (8 uplink and 8 downlink); M=5 (uplink). The goal of this test is to evaluate the adaptation behavior of the candidate scheme when its available bandwidth changes due to the departure of background traffic. The background traffic consists of several (e.g., M=5) CBR flows transported over UDP. These background flows are ON at time t=0-60s and OFF at time t=61-120s.

Multiple bidirectional RTP-based media flows with off-on CBR traffic over UDP: N=16 (8 uplink and 8 downlink); M=5 (uplink). The goal of this test is to evaluate the adaptation behavior of the candidate scheme when its available bandwidth changes due to the arrival of background traffic. The background traffic consists of several (e.g., M=5) parallel CBR flows transported over UDP. These background flows are OFF at time t=0-60s and ON at times t=61-120s.

Multiple bidirectional RTP-based media flows in the presence of background TCP traffic: N=16 (8 uplink and 8 downlink); M=5 (uplink). The goal of this test is to evaluate how RTP-based media flows compete against TCP over a congested Wi-Fi network for a given candidate scheme. TCP flows have start time at t=40s and end time at t=80s.

Varying number of RTP-based media flows: A series of tests can be carried out for the above test cases with different values of N, e.g., N=[4, 8, 12, 16, 20]. The goal of this test is to evaluate how a candidate scheme responds to varying traffic load/demand over a congested Wi-Fi network. The start times of the media flows are randomly distributed within a window of t=0-10s; their end times are randomly distributed within a window of t=110-120s.

3.2.4. Expected Behavior

Multiple downlink RTP-based media flows: Each media flow is expected to get its fair share of the total bottleneck link bandwidth. Overall bandwidth usage should not be significantly lower than that experienced by the same number of concurrent downlink TCP flows. In other words, the behavior of multiple concurrent TCP flows will be used as a performance benchmark for this test scenario. The end-to-end delay and packet loss ratio experienced by each flow should be within an acceptable range for real-time multimedia applications.

Multiple uplink RTP-based media flows: Overall bandwidth usage by all media flows should not be significantly lower than that experienced by the same number of concurrent uplink TCP flows. In other words, the behavior of multiple concurrent TCP flows will be used as a performance benchmark for this test scenario.

Multiple bidirectional RTP-based media flows with dynamic background traffic carrying CBR flows over UDP:

The media flows are expected to adapt in a timely fashion to the changes in available bandwidth introduced by the arrival/departure of background traffic.

Multiple bidirectional RTP-based media flows with dynamic background traffic over TCP:

During the presence of TCP background flows, the overall bandwidth usage by all media flows should not be significantly lower than those achieved by the same number of bidirectional TCP flows. In other words, the behavior of multiple concurrent TCP flows will be used as a performance benchmark for this test scenario. All downlink media flows are expected to

obtain similar bandwidth as each other. The throughput of each media flow is expected to decrease upon the arrival of TCP background traffic and, conversely, increase upon their departure. Both reactions should occur in a timely fashion, for example, within 10s of seconds.

Varying number of bidirectional RTP-based media flows: The test results for varying values of N -- while keeping all other parameters constant -- is expected to show steady and stable per-flow throughput for each value of N . The average throughput of all media flows is expected to stay constant around the maximum rate when N is small, then gradually decrease with increasing value of N till it reaches the minimum allowed rate, beyond which the offered load to the Wi-Fi network exceeds its capacity (i.e., with a very large value of N).

3.3. Other Potential Test Cases

3.3.1. EDCA/WMM usage

The EDCA/WMM mechanism defines prioritized QoS for four traffic classes (or Access Categories). RTP-based real-time media flows should achieve better performance in terms of lower delay and fewer packet losses with EDCA/WMM enabled when competing against non-interactive background traffic such as file transfers. When most of the traffic over Wi-Fi is dominated by media, however, turning on WMM may degrade performance since all media flows now attempt to access the wireless transmission medium more aggressively, thereby causing more frequent collisions and collision-induced losses. This is a topic worthy of further investigation.

3.3.2. Effect of Heterogeneous Link Rates

As discussed in [Heusse2003], the presence of clients operating over slow PHY-layer link rates (e.g., a legacy 802.11b device) connected to a modern network may adversely impact the overall performance of the network. Additional test cases can be devised to evaluate the effect of clients with heterogeneous link rates on the performance of the candidate congestion control algorithm. Such test cases, for instance, can specify that the PHY-layer link rates for all clients span over a wide range (e.g., 2 Mbps to 54 Mbps) for investigating its effect on the congestion control behavior of the real-time interactive applications.

4. IANA Considerations

This document has no IANA actions.

5. Security Considerations

The security considerations in [RFC8868] and the relevant congestion control algorithms apply. The principles for congestion control are described in [RFC2914], and in particular, any new method must implement safeguards to avoid congestion collapse of the Internet.

Given the difficulty of deterministic wireless testing, it is recommended and expected that the tests described in this document would be done via simulations. However, in the case where these test cases are carried out in a testbed setting, the evaluation should take place in a controlled lab environment. In the testbed, the applications, simulators, and network nodes ought to be well-behaved and should not impact the desired results. It is important to take appropriate caution to avoid leaking nonresponsive traffic with unproven congestion avoidance behavior onto the open Internet.

6. References

6.1. Normative References

- [HO-deploy-3GPP]** 3GPP, "Physical layer aspects for evolved Universal Terrestrial Radio Access (UTRA)", TS 25.814, October 2006, <http://www.3gpp.org/ftp/specs/archive/25_series/25.814/25814-710.zip>.
- [IEEE802.11]** IEEE, "Standard for Information technology--Telecommunications and information exchange between systems Local and metropolitan area networks--Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", IEEE 802.11-2012, <<https://ieeexplore.ieee.org/document/7786995>>.
- [NS3WiFi]** "ns3::YansWifiChannel Class Reference", <https://www.nsnam.org/doxygen/classns3_1_1_yans_wifi_channel.html>.
- [RFC5681]** Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC8867]** Sarker, Z., Singh, V., Zhu, X., and M. Ramalho, "Test Cases for Evaluating Congestion Control for Interactive Real-Time Media", RFC 8867, DOI 10.17487/RFC8867, January 2021, <<https://www.rfc-editor.org/info/rfc8867>>.
- [RFC8868]** Singh, V., Ott, J., and S. Holmer, "Evaluating Congestion Control for Interactive Real-Time Media", RFC 8868, DOI 10.17487/RFC8868, January 2021, <<https://www.rfc-editor.org/info/rfc8868>>.

6.2. Informative References

- [Heusse2003]** Heusse, M., Rousseau, F., Berger-Sabbatel, G., and A. Duda, "Performance anomaly of 802.11b", IEEE INFOCOM 2003, Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies, DOI 10.1109/INFCOM.2003.1208921, March 2003, <<https://ieeexplore.ieee.org/document/1208921>>.
- [HO-def-3GPP]** 3GPP, "Vocabulary for 3GPP Specifications", 3GPP TS 21.905, December 2009, <http://www.3gpp.org/ftp/specs/archive/21_series/21.905/21905-940.zip>.

- [HO-LTE-3GPP]** 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol specification", 3GPP TS 36.331, December 2011, <http://www.3gpp.org/ftp/specs/archive/36_series/36.331/36331-990.zip>.
- [HO-UMTS-3GPP]** 3GPP, "Radio Resource Control (RRC); Protocol specification", 3GPP TS 25.331, December 2011, <http://www.3gpp.org/ftp/specs/archive/25_series/25.331/25331-990.zip>.
- [NS-2]** "ns-2", December 2014, <http://nslam.sourceforge.net/wiki/index.php/Main_Page>.
- [NS-3]** "ns-3 Network Simulator", <<https://www.nslam.org/>>.
- [QoS-3GPP]** 3GPP, "Policy and charging control architecture", 3GPP TS 23.203, June 2011, <http://www.3gpp.org/ftp/specs/archive/23_series/23.203/23203-990.zip>.
- [RFC2914]** Floyd, S., "Congestion Control Principles", BCP 41, RFC 2914, DOI 10.17487/RFC2914, September 2000, <<https://www.rfc-editor.org/info/rfc2914>>.

Contributors

The following individuals contributed to the design, implementation, and verification of the proposed test cases during earlier stages of this work. They have helped to validate and substantially improve this specification.

Ingemar Johansson <ingemar.s.johansson@ericsson.com> of Ericsson AB contributed to the description and validation of cellular test cases during the earlier stage of this document.

Wei-Tian Tan <dtan2@cisco.com> of Cisco Systems designed and set up a Wi-Fi testbed for evaluating parallel video conferencing streams, based upon which proposed Wi-Fi test cases are described. He also recommended additional test cases to consider, such as the impact of EDCA/WMM usage.

Michael A. Ramalho <mar42@cornell.edu> of AcousticComms Consulting (previously at Cisco Systems) applied lessons from Cisco's internal experimentation to the draft versions of the document. He also worked on validating the proposed test cases in a virtual-machine-based lab setting.

Acknowledgments

The authors would like to thank Tomas Frankkila, Magnus Westerlund, Kristofer Sandlund, Sergio Mena de la Cruz, and Mirja Kühlewind for their valuable inputs and review comments regarding this document.

Authors' Addresses

Zaheduzzaman Sarker

Ericsson AB
Torshamnsgatan 23
SE-164 83 Stockholm
Sweden
Phone: +46 10 717 37 43
Email: zaheduzzaman.sarker@ericsson.com

Xiaoqing Zhu

Cisco Systems
Building 4
12515 Research Blvd
Austin, TX 78759
United States of America
Email: xiaoqzhu@cisco.com

Jiantao Fu

Cisco Systems
771 Alder Drive
Milpitas, CA 95035
United States of America
Email: jianfu@cisco.com