Network Working Group
Internet-Draft

T. Alexander
VeriWave, Inc.
S. Bradner
Harvard University

Expires October 2005

April 2005

# Benchmarking Methodology for Wireless LAN Devices

<draft-alexander-wlan-meth-01.txt>

## Status of this Memo



## Copyright Notice



## Abstract

This document provides a framework and methodology for performing stress testing and benchmarking of wireless LAN (WLAN) devices, including clients (i.e., host interfaces) and Access Points. The document defines and discusses a number of tests and associated test conditions that may be used to characterize the performance of such devices, and also supplies the methods used to calculate the results of these tests. This document also describes specific formats for reporting the results of the tests. It extends the methodology defined for benchmarking network interconnecting devices in RFC 2544 and LAN switches in RFC 2889 to IEEE 802.11 WLAN devices.

Table of Contents

## 1. Introduction

This document defines and describes a specific set of tests that may be used by vendors and users of IEEE 802.11 Wireless LAN (WLAN) [802.11] devices to measure and report performance characteristics of such devices. It extends the methodology that was originally defined for benchmarking network interconnecting devices in RFC 2544 [RFC2544], and then subsequently extended to other types of devices (such as LAN switching devices in RFC 2889 [RFC2889]), to cover IEEE 802.11 WLAN devices. Note that only IEEE 802.11 conformant devices are covered by this document; other technologies (e.g., Hiperlan/2) are not considered.

Wireless LANs may be characterized as using a complex, rate-adaptive protocol designed for shared media access and subject to numerous environmental influences, many of which are outside the control of the end-user. The tests in this document are therefore intended to provide a means of comparison and evaluation, rather than absolute measures of performance in an arbitrary end-user environment.

## 2. Existing definitions and requirements

RFC 1242, "Benchmarking Terminology for Network Interconnect Devices" [RFC1242] is relied upon for much of the terminology used in this document, and MUST be consulted before attempting to make use of this document. In addition, RFC 2285, "Benchmarking Terminology for LAN Switching Devices" [RFC2285] SHOULD be reviewed as well. RFC 2544, "Benchmarking Methodology for Network Interconnect Devices" [RFC2544] and RFC 2889, "Benchmarking Methodology for LAN Switching Devices"

[RFC2889], also provide useful background information and context. The WLAN-specific terms and definitions in this document are described in Clauses 3 and 4 of the IEEE 802.11 standard [802.11].

For the sake of clarity and continuity this RFC adopts the general template for benchmarking tests set out in Section 26 of RFC 2544.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

## 3. Tester description and test setups

A common set of test setup and measurement conditions is used across all of the tests described in this document. Exceptions to these conditions are noted, if necessary, in the descriptions of the individual tests.

### 3.1. Functional model of the tester

For the purposes of this document, the tester is defined as a separate device that is used to transmit controlled test traffic to the physical ports of the device under test (DUT) as well as to receive and measure test traffic from the physical ports of the DUT. The tester MUST NOT be a part of the DUT, nor can the DUT provide any portion of the reported test results. An exception is the case of client devices, which MAY be required to host test software in order to support the requirements of one or more tests.

The tester MUST transmit 802.11-conformant traffic to the DUT during the tests described herein, and MUST follow the rules of the 802.11 protocol with respect to media access and frame exchanges. It MAY be configured to transmit non-conformant traffic for special purposes (e.g., for debug), but this is outside the scope of this document. The tester MUST support some means of distinguishing test traffic (either injected into or emitted by the DUT) from normal data, control and management frames that are generated by the DUT itself. The tester SHOULD further support means of unambiguously determining frame loss and frame duplication (e.g., by the use of sequence numbers), as well as time-stamping transmitted and received frames.

No constraints are placed by this document on the specific implementation of the tester or test system, provided that it is capable of measuring DUT responses to the required degree of accuracy, establishing the required test conditions at the physical interfaces of the DUT, and generating test traffic with the relevant parameters. These parameters include frame sizes, offered load, burst sizes and inter-burst gap, signal output level, etc.

## 3.2. Test setup

### 3.2.1. Test setup for Access Points

Many of the tests performed on Access Points require that test traffic be injected into and/or received from both the wired and the wireless ports of these devices. The ideal means of implementing tests on Access Points is therefore to use a tester or test system that can interface to both types of ports on the DUT, as shown in Figure 1.

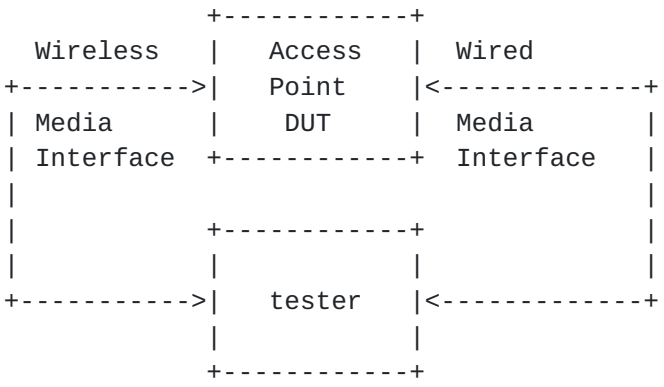


Figure 1

The tests in this document are also generally applicable to DUTs that provide multiple media interfaces, such as an Access Point that supports multiple wireless interfaces that are aggregated into a single wired interface. Section 6 of RFC 2544 [RFC2544], as well as Section 5.2.3 of RFC 2889 [RFC2889], may be consulted for information on extending the tests to multi-port devices. Test results MUST be reported separately for each physical port of the DUT.

### 3.2.2. Test setup for clients

Due to the variety of physical configurations of client devices, a test setup for a client is dependent on the nature of the DUT. If the client is multi-homed, and can be set up to transfer traffic between the wireless interface being tested and some other physically accessible interface, then a test setup similar to Figure 1 can be used, as shown in Figure 2.

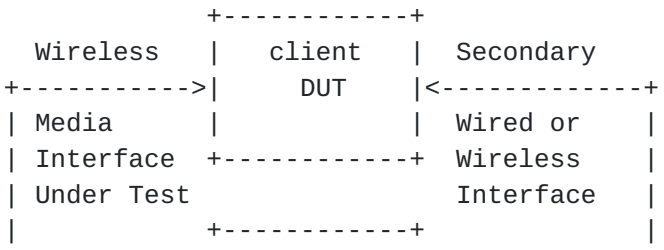

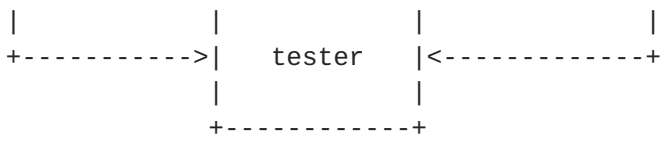


Figure 2

The DUT is set up to route traffic injected into the wireless interface under test to a secondary interface, which may be either wired or wireless. Also, traffic injected into this secondary interface by the tester is transferred to the wireless interface under test. The secondary interface MUST have bandwidth and delay characteristics that are known to be much better than that of the wireless media interface of the DUT that is actually being tested.

If the DUT is not multi-homed (or lacks a suitable secondary interface), then some test software MAY be supported on the client itself in order to enable it to be tested, as shown in Figure 3.

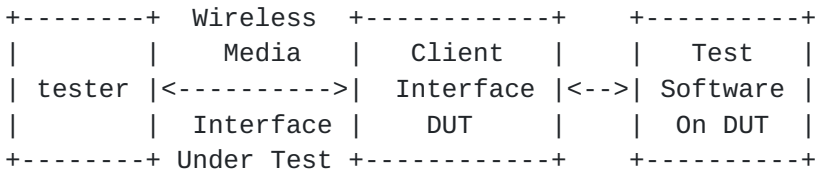


Figure 3

If the approach of Figure 3 is adopted, then it is understood that the DUT comprises only the physical media interface and any device driver and protocol stack software that is actually interposed between the physical interface and the test software. The entire device of which the physical media interface is a part MUST NOT be considered to be tested by this method, as the test software may have an impact on the remainder of the client device that is not characterized by the tests presented in this document.

If a software program is executed on the DUT in order to enable testing, the description and version of this software MUST be included along with the test results.

## 3.3 DUT setup

The general DUT setup MUST follow the requirements described in Section 7 of RFC 2544 [RFC2544].

The specific software or firmware version being used in the DUT, as well as the exact DUT configuration (including any functions that have been disabled) MUST be reported together with the results.

### 3.3.1. Access Point setup

Access Points MUST be configured with both their wired and their wireless interfaces active following the instructions for normal use from the manufacturer. The wired interface MAY be disconnected from the tester for tests that do not involve traffic exchanged between the wired and wireless interfaces.

A System under Test (SUT) comprising one or more Access Points interfaced to a dedicated management and switching device MAY be treated as a single complex DUT and tested as a unit. In this case, the wired interface(s) of the tester MUST be connected to wired interface(s) on the management and switching device, instead of to the wired interfaces of the Access Point(s) directly.

Access Points generally operate in one or more of the following modes:

Intra BSS (Basic Service Set) mode. Here, an Access Point receives data packets from a wireless client and forwards these packets to another wireless client directly, i.e., without traversing the wired media. Both clients must be associated with the same Access Point.

ESS (Extended Service Set) mode. In this case an Access Point receives data packets from a wireless client and forwards these to a wired client (residing on its wired interface) or to another wireless client by means of a second Access Point and the wired network.

WDS (Wireless Distribution System) mode. This is a variant of the ESS mode. Here an Access Point receives data packets from a wireless client associated with it, and forwards these packets to a second Access Point over the wireless medium. The second Access Point then transfers these packets to their final destination. The effect is similar to a wireless repeater network.

A single Access Point is normally capable of performing data transfers using more than one of the above modes simultaneously.

### 3.3.2. Client setup

Devices containing client DUTs SHOULD be set up using the manufacturer's normal instructions to match an normal user configuration; however, user applications and processes that are not part of a vendor-supplied device configuration MAY be removed or suspended during the tests. Vendor-supplied configuration utilities SHOULD be used to configure the various parameters of the wireless interface (e.g., fragmentation thresholds) according to the requirements of the test.

Many client devices offer one or more power-saving modes that can materially impact the test results. Tests SHOULD be run at different power-save settings, but a full suite of tests MUST be run at least one specific setting and the results reported. The power-save mode setting MUST be reported along with the test results. (See Section 3.4.5.)

Clients operate in one of the following modes:

Associated with an Access Point. In this case, the client sends all data traffic to the Access Point, for forwarding to the ultimate destination.

In Independent BSS (IBSS) mode. This is a peer-to-peer mode of operation, where two or more clients form an "ad hoc network" and forward packets amongst each other without the services of an Access Point.

Clients cannot operate in both of the above modes simultaneously.

## 3.4. Wireless configuration parameters

DUT setup considerations specific to WLAN devices are given below.

### 3.4.1. Channel assignment

The DUT MUST be configured to use a wireless channel that a normal user would use at the location where the test was run. For example, if the test is run in the U.S. then a standard U.S. wireless channel is used. The channel used MUST be reported with the test results.

### 3.4.2. Transmit power level

For DUTs with adjustable transmit power levels, tests MAY be run at different transmit power settings, although a full suite of tests MUST be run at each power setting tested. The power setting used in each test MUST be reported with the test results.

### 3.4.3. RTS threshold

The 802.11 protocol supports the use of a Request To Send (RTS) / Clear To Send (CTS) handshake prior to data transfer, as a means for interfaces to seize and reserve the medium before actually transferring data. This is normally expected to be used for larger frame sizes, where contention-based medium access may result in high retransmission overheads. Determination of whether to use an RTS/CTS handshake is based on the size of the frame to be transmitted, and is statically configured as a threshold on the frame size.

For DUTs with adjustable RTS thresholds, tests MAY be run at different RTS thresholds, although a full suite of tests MUST be run at each RTS threshold tested. The RTS threshold used in each test MUST be reported with the test results.

### 3.4.4. Fragmentation threshold

The 802.11 protocol supports fragmentation and reassembly at the link layer, in order to decrease retransmission overhead under high error rates that may prevail in a radio frequency (RF) environment. If a fragment of a frame is lost due to a bit error or a collision, only that fragment is retransmitted. Determination of whether to fragment a frame before transmission is based on the size of the frame, and is statically configured as a threshold on the frame size.

For DUTs with adjustable fragmentation thresholds, tests MAY be run at different fragmentation thresholds, although a full suite of tests MUST be run at each fragmentation threshold tested. The fragmentation threshold used in each test MUST be reported with the test results.

The tests in this document are performed at different fixed frame sizes. The number of fragments produced with a given fragmentation threshold will be known a priori for any given frame size. The tester SHOULD therefore verify that the number of fragments generated by the DUT during a test is correct.

### 3.4.5. Power management mode

The 802.11 protocol supports several power management functions, in order to allow client devices to reduce power by placing their wireless interfaces in a low-power mode during known periods of inactivity. (Access Points do not enter a power-saving mode, but support power-save by clients associated with them.) Transmission to a client in power-save mode is typically bursty; the Access Point accumulates packets destined for the client in internal buffers, notifies the client of these packets by means of special fields in its beacons, and then transfers them when the client and requests its outstanding data. Good time synchronization between Access Points and clients is essential for efficient and low-latency data transfers, as clients need to be awake and listening when Access Points issue notifications via beacons.

Throughput and latency measurements on a client are significantly affected by the power management mode of the client. Therefore, throughput and latency tests SHOULD be run with power management disabled or minimized. If the DUT and tester are capable of supporting multiple power management modes, then tests MAY be run in different modes. The power management mode of the client MUST be

reported with the test results.

### 3.4.6. Service priority

For DUTs with adjustable service priorities (QoS levels), tests MAY be run at different service priorities, although a full suite of tests SHOULD be run at least one service priority. For such DUTs, the service priority used in each test MUST be reported with the test results.

Throughput and latency tests on Access Points involving traffic traversing wired interfaces can be affected by QoS settings on these wired interfaces. In such situations, the QoS settings assigned to the wired interfaces of Access Points MUST be reported with the test results.

## 3.5. Test conditions

Test conditions for measurements on WLAN devices are covered in this section. The complexity of the wireless LAN media and protocol necessitate special attention to specifying and setting up these conditions in order to obtain repeatable results.

### 3.5.1. Test environment

Wireless LAN test environments may be divided into two general categories: shielded environments and open-air environments.

Shielded environments use cabling and/or RF shielding techniques to significantly attenuate signals and noise unrelated to the test. Typically, the DUT is enclosed within an RF-tight chamber and cabled to the tester (which is also placed in an RF-tight chamber); alternatively, both the DUT and the tester may be placed in the same chamber, such as a Faraday cage. Unless the chamber is fairly large, coupling between the DUT and tester is conducted (i.e., via cables) and not radiated (via antennas).

Open-air environments mimic the actual use model of a WLAN DUT. In this case, the DUT is placed at a specific location within some moderately controlled (or at least characterized) indoor or outdoor environment. The tester is also placed within the same environment at a specific position relative to the DUT and other known signal sources (if any). Antennas are used on both the DUT and the tester to enable signal transfer; coupling is hence radiated.

The tests described in this document can be carried out in either of the above environments, provided that the remainder of the test conditions specified in this section (particularly the signal-to-

noise and signal-to-interference ratios, described in Section 3.5.9) are met. The test environment used in the tests MUST be described along with the results.

### 3.5.2. Frame sizes

All of the described tests SHOULD be performed using several fixed sizes of test data frames. Frame sizes are calculated from the first byte of the MAC DA to the last byte of the FCS (i.e., all of the 802.11 MAC header and trailer fields are included, but the PLCP header is not included). The various mode-specific encapsulations supported by the 802.11 protocol makes frame size calculations somewhat challenging. The test results MUST list the frame sizes used for test data frames.

When using test data encoded using the 3-address 802.11 frame format without encryption, the data frame sizes that SHOULD be used during the tests are:

28, 64, 128, 256, 512, 1024, 1528, 2048, 2332

The payload length may be obtained by subtracting 28 from the frame size. A frame size of 28 bytes corresponds to a null frame (i.e., an 802.11 data frame with a zero-length payload). A frame size of 1528 bytes results in a payload length of exactly 1500 bytes, which is the maximum sized frame that can be bridged on to an Ethernet LAN. A frame size of 2332 bytes corresponds to the maximum defined 802.11 MAC Service Data Unit (upper-layer payload) of 2304 bytes.

Test data that uses a 3-address 802.11 frame format with Wired Equivalent Priority (WEP) encryption SHOULD use the following frame sizes:

28, 64, 128, 256, 512, 1024, 1536, 2048, 2340

With the exception of null frames, the payload length may be obtained by subtracting 36 from the frame size. Null frames contain no WEP header and thus remain at 28 bytes.

When using a 4-address 802.11 frame format without encryption, the 802.11 header/trailer overhead increases to 34 bytes, and the test data frame sizes that SHOULD be used are:

34, 64, 128, 256, 512, 1024, 1534, 2048, 2346

In this case, a frame size of 34 bytes corresponds to a null frame, 1534 bytes to a payload length of 1500 bytes, and 2346 bytes to the maximum defined 802.11 payload length of 2312 bytes. Note that the

usable 802.11 payload sizes for the other frame size values are smaller by 6 bytes over the 3-address case.

Finally, when using a 4-address 802.11 test data frame format with WEP, the 802.11 header/trailer overhead increases to 42 bytes, and the frame sizes that SHOULD be used are:

34, 64, 128, 256, 512, 1024, 1542, 2048, 2354

As before, null frames do not contain a WEP header.

Inclusion of additional 802.11-specific header and trailer fields for extended capabilities (e.g., advanced encryption, QoS support) can cause the frame size to change. Test data traffic that includes such additional header and trailer fields SHOULD use frame sizes consistent with those given above.

In some cases, such as tests on clients, or ESS testing on APs (see Subclause 3.25 of the IEEE 802.11 standard [802.11] for a description of ESS), not all of the above frame sizes are practicable. For example, a null frame will not be processed by clients. Test procedures for these specific situations detail exceptions to the frame sizes to be used.

Frame sizes of 802.11 management and control frames generated during the test MUST conform to those required by the 802.11 standard [802.11].

### 3.5.3. Frame formats and verification

The frame formats used for test data frames (with the exception of null 802.11 frames) SHOULD follow the recommendations in Appendix C of RFC 2544 [RFC2544]. LLC/SNAP encapsulation as per RFC 1042 [RFC1042] MUST be used to contain higher-layer payloads. Note that when the added overhead of the 8-byte LLC/SNAP header is taken into account, 64 byte WLAN frame sizes can only support link layer payloads ranging from 28 bytes to as little as 14 bytes, and hence may not be able to contain an IP header.

With the exception of null frames, the test frame format MUST contain some means (such as a unique signature field, as described in Section 4 of RFC 2544 [RFC2544]) that will enable the tester to filter out frames that are not part of the offered load, or are duplicated by the DUT. In tests on Access Points involving multiple virtual or physical test clients, the test frame format MAY also support means for distinguishing between frames originating from different clients.

The provisions for verifying received frames in Section 10 of RFC

2544 [RFC2544] SHOULD be followed as well. This is particularly significant for 802.11, which implements retransmission at the link layer. The verification of received frames SHOULD be independent of the facilities provided by the 802.11 protocol.

### 3.5.4. Half-duplex effects on calculating offered load

WLAN Access Points and clients perform medium access in half-duplex mode, implementing deference and backoff to minimize the probability of collisions. Further, packet transmission is bidirectional, in that data transmission from source to destination is immediately followed by an acknowledgement from destination to source. This leads to a number of issues when attempting to impose or calculate an offered load during testing.

The relevant characteristics of 802.11 media access are:

Carrier sensing before access. Stations are required to defer to all ongoing transmissions before attempting to transmit. Small changes in relative access timing between stations may result in large variations in the actual access pattern.

Random backoff after all transmission attempts. The 802.11 protocol attempts to avoid collisions by forcing all stations to delay for a random interval after both successful or unsuccessful transmission attempts. The random backoff must be implemented even if only one station is attempting to transmit.

Suspension of backoff timers when the medium is busy. To preserve station priority in a busy environment, the 802.11 protocol requires backoff timers for stations contending for the medium to be suspended (rather than reset) during deference periods. Hence stations contending for media access over more than one deference period will have a higher probability of access than a station that is contending for the first time.

The last two characteristics above are peculiar to 802.11, and not usually found in other half-duplex media access methods such as Ethernet. Further, 802.11 media access uses an acknowledged transfer (thus leading to inherently bidirectional packet flow, even though the intended test data traffic is unidirectional); imposes further delays when acknowledgement frames are lost; and cause stations to emit a much higher proportion of management and control packets during data transfer. All of these conspire to render data traffic in a WLAN inherently bursty, as interfaces are forced to insert gaps in otherwise steady flows when packets are being received or when backoffs occur. Care must therefore be taken when measuring throughput or computing offered load, especially for bidirectional data transfers.

In particular, as noted in RFC 2285 [RFC2285], special attention must be paid towards ensuring that the actual offered load on the DUT is measured, instead of (or as well as) the intended load. The offered load may be less than the intended load due to contention effects for the wireless medium. The tester MUST adjust the inter-frame spacing according to the target intended load (i.e., to achieve the desired rate of frame transmission), and then MUST measure and report the actual offered load at the end of the trial.

See Appendix A of this document for notes about generating the intended load for these tests. Either the frame-based or the time-based method described in Appendix B of RFC 2889 [RFC2889] MAY be used for these tests but, in either case, the method used MUST be reported with the results. Most of the tests in this document use a constant (non-bursty) load, and the Iload calculations in Section A.2 apply. Burst loads use the calculations in Section A.3.

The tester MUST follow the normal 802.11 protocol requirements when transferring test data frames to the DUT, unless the DUT is to be oversubscribed, or the burst capacity of the DUT is to be measured. The tester SHOULD also note attempts by the DUT to violate the timing requirements of the 802.11 protocol by not conforming to the backoff rules, or by reducing its inter-frame spacing to less than the legal minimum.

For bidirectional test data traffic, any offered load beyond 50% of the theoretical maximum (computed as per Appendix A) MUST be considered to oversubscribe the DUT. Oversubscription of the DUT MUST be recorded as part of the test results.

### 3.5.5. Retry of unacknowledged frames

Recovery from frame errors and collisions is performed by 802.11 stations using a simple stop-and-wait acknowledgement handshake. If a sending station does not receive a valid acknowledgement frame within a short time delay after it transmits a unicast data or management frame, it is required to perform a backoff and retry. Sequence numbers and flag bits are used to distinguish between retries and new frames.

The tester MUST follow the backoff and retry rules of the 802.11 protocol. When performing tests involving multiple virtual stations sending to a single DUT, the tester SHOULD maintain a separate backoff timer for each individual virtual station. In the case of throughput tests, however, the tester MAY perform retransmissions with a reduced or zero backoff period, in order to maintain a reasonably high offered load in spite of the high error rates found in WLAN media. If this is done, the actual backoff period used MUST

be reported with the test results.

The number of retries performed by the tester when conducting throughput and forwarding rate tests SHOULD be at least 7.

Note that tests involving reduced (non-standard) backoff periods - and the consequently high offered loads - could expose catastrophic behavior on the part of the DUT, which in turn could be indicative of mysterious field failures. Also, an aggregate of stations sending packets to a single DUT, such as an Access Point in a LAN, could result in the DUT receiving bursts of packets at the minimum interframe gap, even though each client individually conforms to the prescribed backoff rules.

### 3.5.6. Physical layer (PHY) data rates

The physical layer of the 802.11 WLAN protocol supports data transfer at a number of different data rates, such as 1, 2, 5.5, 6, 9, 11, 12, etc. Mb/s. Further, the receiver and transmitter in a data exchange may use different PHY data rates. These data rates are achieved with different modulation formats, generally resulting in different packet error ratios and/or signal-to-noise ratios at the receiver. Tests performed at one data rate may not correlate with tests performed at another rate. The test results MUST therefore record the data rate used by both the DUT and the tester.

Rate adaptation is supported by 802.11 devices in order to maintain data transfer under changing noise and interference environments (although at different throughputs). Devices may automatically fall back to lower PHY data rates in order to cope with decreasing signal strength or increasing noise levels, as the lower PHY data rates provide better signal-to-noise ratios. This can make test results impossible to reproduce or interpret in the case of measurements needing constant PHY data rates. A given trial MUST maintain a constant PHY data rate for all test data packets, unless otherwise specified.

### 3.5.7. Management and control frames

Unlike most other LAN link layer protocols, 802.11 involves the exchange of a substantial number of management and control frames during and in between data transfers. Control frames are typically used for acknowledgement, medium reservation, and polling. Management frames are required for discovery, connection setup and teardown, and other signaling. Note that every unicast data frame that is successfully transferred requires an acknowledgement control frame to be transmitted by the receiver.

To emulate the conditions occurring in a real 802.11 WLAN, therefore, tests SHOULD include some (small) amount of management and control frames in addition to acknowledgement frames. The proportion of management and control frames, including beacons, MAY be kept under 5% of all frames transferred.

### 3.5.8. Authentication and association

Transfer of data between stations on an 802.11 WLAN is permitted to begin only after a successful connection setup, as indicated by an authentication handshake followed by an association handshake. (See Clause 8 and Subclause 11.3, respectively, of the 802.11 standard [802.11].) Also, it is illegal to transfer data after a connection has been terminated.

The tester MUST be authenticated and associated with the DUT prior to the start of a trial. It SHOULD perform an authentication and association handshake with the DUT at the start of each trial, and a deauthentication handshake at the conclusion; however, it MAY elect to remain authenticated with the DUT across multiple trials. In the case of throughput, forwarding rate, latency and burst capacity tests, disassociation by the DUT during the test data transfer portion of a trial MUST be reported and SHOULD cause the trial to be terminated.

The tester MUST NOT accept any unicast data frames from the DUT until after the successful completion of the authentication and association handshake, and MUST NOT accept any data frames, whether valid or not, after receiving an acknowledgement for a disassociation or deauthentication frame transmitted to the DUT.

### 3.5.9. Signal level and signal-to-noise ratios

The 802.11 WLAN protocol is fundamentally based on a shared-medium RF physical layer, and is therefore significantly affected by:

- The absolute signal level received by the PHY;
- The signal-to-noise ratio, as measured at the input of the PHY; and
- The carrier-to-interference (C/I) or signal-to-interference ratio, also as measured at the input of the PHY. This parameter includes both co-channel interference (CCI) and adjacent-channel interference (ACI).

The above parameters impact key attributes of the link between two stations, such as packet error rate, retries and backoffs, deference

to ongoing transmissions, etc. These parameters MUST be measured and maintained within known limits during every trial. Measurements are usually performed in one of three ways:

Directly from the DUT, assuming that the DUT is capable of accurately measuring these parameters and providing them to the tester upon request.

Directly from the tester, assuming that the tester is provided with the requisite measurement capabilities and the path between the DUT and tester is controlled and well characterized.

Using a passive listening device (which MAY be part of the tester or test system) that is co-located with the DUT.

Unless otherwise specified, the tester MUST ensure that the signal-to-noise and carrier-to-interference ratios are at least 10 dB above the minimum and 10 dB below the maximum specified levels for a 10% Packet Error Ratio (PER) as measured at the DUT. Further, the tester MUST ensure that the absolute signal level received by the DUT is held constant to within +/- 5 dB over the duration of the trial. The signal level MUST be reported with the test results.

In order to evaluate the performance of the DUT at various signal levels, the tests described in this document MAY be run with various values of power output by the tester. If this is done, the power output SHOULD be varied in steps of 10 dB.

Note that the specific method used to control the received signal and the signal-to-interference ratios is implementation-specific and is outside the scope of this document.

### 3.5.10. Beacons and PCF access method settings

IEEE 802.11 networks use special management frames, referred to as beacons, to enable discovery of Access Points by clients and synchronize protocol timers within a group of stations. Beacons are typically emitted every 100 Time Units (a Time Unit is 1024 microseconds, so the usual inter-beacon period is 102.4 milliseconds).

As beacons are key elements of the 802.11 discovery and connection maintenance protocol, the tester MUST generate beacons with the correct contents and at accurate intervals when performing tests on clients. Beacons generated by the tester MUST support a timing accuracy of at least 1% relative to the advertised beacon period. In the case of tests on Access Points, the tester MUST follow the standard 802.11 deference rules in order to allow the Access Point to

transmit its beacons. The DUT MUST be configured with the correct parameters necessary to generate beacons with the required contents.

Beacons are also used to define the starting points of contention-free periods (CFPs), where the 802.11 Point Coordination Function (PCF) access rules apply and a polling mode of data transfer is performed under the control of an Access Point. This document does not address tests performed during CFPs. In the case of tests on Access Points, PCF mode SHOULD be disabled or turned off if possible. If PCF mode cannot be disabled, it MUST be minimized as far as possible, and traffic SHOULD NOT be transmitted to the DUT by the tester using PCF mode.

### 3.5.11. Multiple clients

In the case of tests on Access Points, measurements of frame loss, throughput, forwarding rate, latency and burst capacity SHOULD be performed with multiple clients concurrently transmitting to the same DUT. This simulates the situation in an actual network, where multiple clients connect to a single Access Point and contribute to the total offered load. Each client is represented by a different MAC address. The MAC addresses SHOULD be chosen randomly to exercise the DUT's ability to perform lookups.

The number of clients used in the test MUST be reported. For throughput, frame loss, forwarding rate and burst capacity tests, the aggregate results for all of the clients combined MUST be reported. For latency tests, the worst-case latency and latency variation among all of the clients MUST be reported.

### 3.5.12. Trial duration

The duration of each trial SHOULD be selected using the guidelines of Section 24 of RFC 2544 [RFC2544]. Further, it SHOULD be long enough to minimize any connection setup and startup effects, such as authentication and association, that can affect the test results. It SHOULD also be long enough to make the random fluctuations of the CSMA/CA access method statistically insignificant.

The recommended duration of each trial is 30 seconds. The trial duration SHOULD be adjustable between 1 second and 300 seconds. The tester MUST transmit all test data frames within the trial duration. To eliminate the case where a device possessing large frame buffers can appear to be faster than it actually is, the tester MUST NOT accept received test data frames for more than 1% beyond the end of the trial duration, or 1 second, whichever is larger. (Thus a 30 second trial duration causes the tester to receive frames for no more than 31 seconds, starting from the beginning of the trial.) Frames

received outside of these limits are not counted as part of the results.

### 3.5.13. Configuration combinations

WLAN DUTs typically offer a number of orthogonal configuration parameters. For instance, the setting of fragmentation is independent of RTS/CTS usage, and both are independent of the security mode employed. This leads to a potentially enormous number of combinations of configuration parameters, and consequently a very large number of possible tests.

To reduce the amount of test time and effort, each test defines a baseline configuration that MUST be set up and tested. The baseline configuration is then modified by altering a single parameter at a time (holding all others at the baseline), and the test repeated. This process reduces the total number of tests to be performed, while still providing useful information regarding the influence of user-configurable parameters on DUT performance.

### 3.5.14. Basic test parameters

Unless otherwise specified, the following are the defaults for test parameters. The specific parameters to be used are listed for each test.

Burst size - Tests that measure the burst capacity of the DUT SHOULD be run with burst sizes between 1 (constant load) and 930 frames.

Fragmentation - Tests MAY be performed with fragmentation enabled as well as disabled. The results MUST be reported separately. If fragmentation is enabled, the DUT and tester MUST be configured to produce at least two fragments per data frame.

Frame sizes - The frame sizes listed in Section 3.2.3 above MUST be used for data traffic in the tests described in this document, with the exception of tests on clients, where the 28 byte frame size MAY be omitted, and tests on Access Points requiring transfers between wired and wireless media, in which case frame sizes greater than 1542 bytes MAY be omitted. The actual frame sizes used MUST be reported.

Frame spacing - For tests involving constant loads, the spacing between frames MUST be computed from the intended load according to the calculations in Section A.2. When performing tests involving bursts of frames generated by the tester, the spacing between frames within a burst MUST be set to the DIFS value for

the PHY data rate and type for the test, and the spacing between bursts MUST be computed from the intended load according to the calculations in Section A.3. The tester MAY additionally perform tests with an IFS within the burst being larger than the DIFS value; these results MUST be reported separately. A maximum offered load exceeding 50% of the theoretical capacity of the wireless medium will oversubscribe the DUT, resulting in frame loss.

Nominal beacon interval - In the case of tests on Access Points, if the DUT enables the beacon period to be adjusted, it SHOULD be set to 102.4 milliseconds. (See Section 3.5.10 above.) The nominal beacon period of the DUT MUST be reported.

Number of STAs - In the case of tests on Access Points, the minimum number of test clients is 2; however, the test SHOULD be carried out with 64 or more concurrent clients. Each client SHOULD make up an equal fraction of the total offered load. The number of virtual or physical clients that are used in the test MUST be reported.

RTS/CTS usage - Tests MAY be performed with the RTS/CTS handshake enabled as well as disabled. The results MUST be reported separately.

Security usage - If the DUT can be configured to implement one or more security modes, such as WEP [802.11], tests SHOULD be performed using each of the security modes as well as with security disabled. The results MUST be reported separately.

Signal level - Tests SHOULD be performed using multiple signal levels as generated by the tester and measured at the DUT. The average signal level recorded at the DUT, as well as the signal level being generated by the tester, MUST be reported for each trial.

Test Access Point beacon period - The beacon period of the Access Points used to test a client MUST be the same (to within 1%) and MUST be reported with the test results.

Uni-directional transfer - Each trial causes test data to flow in only one direction; i.e., from the wired to the wireless media, or vice versa, but not both. No test data frames are to be directed between clients on the wireless medium during tests involving uni-directional transfers. As explained in Sections 3.5.4 and 3.5.5, acknowledgement frames MUST NOT cause test data traffic to be interpreted as bidirectional.

Bi-directional transfer - Each trial causes test data frames to flow in both directions (for example, from the wired to the wireless media and from the wireless to the wired media). This also covers the case where test data traffic is directed between clients on the wireless medium.

Wireless Distribution System (WDS) - Tests MAY be performed on Access Points supporting the Wireless Distribution System (WDS) mode of operation [802.11]. (See Section 3.3.1 above.) The results of tests performed in this mode MUST be reported separately.

Power management mode - As noted in Section 3.4.5, tests on clients SHOULD be run with a baseline configuration that disables power management.

## 4. Interpreting and reporting test results

Test results SHOULD be reported in a common format to aid the reader in interpreting results and comparing them across DUTs. Results from a set of trials involving the variation of one or more test parameters described in Section 3.5.14 above SHOULD be presented as graphs, with the x coordinate being the parameter value and the y coordinate being the test results.

The following test conditions MUST be reported with the results of all trials:

PHY type (e.g., 802.11g), bit rate and channel used

Signal power level, signal-to-noise ratio and signal-to-interference ratio of signals from tester, measured at or near the DUT

PLCP layer options (short slot time, short preamble, etc.) configured for the DUT and the tester

PCF mode settings if PCF mode is not disabled

## 5. Benchmarking tests

The following tests are divided into three categories. Throughput related tests deal with the rates at which the DUT can perform various functions, such as forwarding data traffic or performing authentications and associations. Latency and timing tests measure the time taken by the DUT to carry out specific tasks, such as forwarding a frame, adapting to a new rate, or recovering from a reset. Finally, capacity tests quantify the storage capacity of a DUT for different functions, such as dealing with bursts of data traffic.

Objectives, test parameters, procedures and reporting formats are described for each test.

## 5.1. Throughput related tests

Throughput related tests measure the rates at which the DUT can perform its tasks. For an Access Point, this includes the rates at which it can forward data within the BSS or between the BSS and the DS, as well as the rate at which new clients can establish associations with it. For a client, this measures the rate at which it can accept and transmit data.

For throughput and forwarding rate tests, either the Frame Based or Time Based modes of testing may be used, as described in Appendix B of RFC 2889 [RFC2889]. The DUT is initially set up according to the baseline configuration, using a starting combination of test parameters. Packets are then sent to the DUT by the tester at a specific offered load for the duration of the trial, and the number of frames received from the DUT are counted. The process MUST be iterated at different offered loads, using a search algorithm, until the desired measurement (throughput or maximum forwarding rate) has been made. Additional trials are then performed in the same manner using different DUT configurations until all configurations have been exhausted.

The tester MUST count, as valid received test frames, those which it receives without error and with the proper signature (i.e., the right combination of source and destination addresses, frame length and frame payload). In addition, on the wireless medium the tester must only count as valid those unique data frames for which it sent an 802.11 ACK frame to the DUT in response. It MUST NOT count duplicate frames, frames originating from the DUT, data frames that it did not acknowledge, or management and control frames as part of the measurements. Such frames MAY be counted separately, or not counted at all.

For the purposes of computing the actual offered load, the tester MUST count as valid transmitted frames only those test frames that were acknowledged by the DUT on the wireless medium (i.e., with an 802.11 ACK frame), or transferred to the DUT without a locally detected error on the wired medium. All other frames MUST NOT be counted as part of the offered load.

### 5.1.1. Unicast intra-BSS throughput, forwarding rate and frame loss

#### 5.1.1.1. Objective

To determine the throughput of the DUT when handling unicast WLAN

data frames that are confined to the wireless medium (and, in the case of clients, internally to the client). This test is applicable to both clients and Access Points. In the case of clients, this test provides the basic measure of their ability to transmit and receive frames without loss across their wireless interface. In the case of Access Points, this test measures their ability to forward frames from one wireless client to another on the wireless medium.

This test is applicable to IBSS (Independent BSS) as well as infrastructure BSS client configurations. If an IBSS client is being tested, the results determine the ability of the client to exchange data traffic with another IBSS client. In infrastructure mode, the results determine the ability of the client to exchange data with an Access Point.

#### 5.1.1.2. Test parameters

The following parameters are relevant to this test. See Section 3.5.14.

> Frame sizes, Signal level, Number of STAs, Fragmentation, RTS/CTS usage, Security usage and Wireless Distribution System (WDS).

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, security and WDS not used.

#### 5.1.1.3. Procedure

The DUT is initially set up according to the baseline configuration, using a starting combination of frame size and tester signal level. The frame spacing required for a given offered load is computed per Appendix A. The tester MUST follow the half-duplex test conditions described in Section 3.5.4. The throughput, maximum forwarding rate and frame loss rate are then found as described below.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 3 plus the number of security modes.

When testing Access Points with multiple physical or virtual clients, consecutive frames transmitted by the tester to the DUT MUST have different combinations of source and destination addresses, and all possible such combinations of addresses MUST be represented equally within each trial. This distributes the load of transmission and reception uniformly among the clients. Failure to ensure this can lead to inconsistent results.

#### 5.1.1.4. Analysis and reporting

The throughput of the DUT is computed and reported (per Section 26.1 of RFC 2544 [2544]) as the maximum offered load, in frames per second, resulting in zero frame loss rate [RFC1242].

The maximum forwarding rate of the DUT is computed and reported as the maximum number of test frames per second that the DUT is observed to successfully forward, irrespective of frame loss, at some value of offered load. The offered load applied to the DUT at the maximum forwarding rate MUST be reported as well.

The frame loss rate MUST be reported with the maximum forwarding rate. In this context, "rate" refers to the percentage of frames that were successfully injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as graphs of throughput and maximum forwarding rate versus each of: frame size and signal level. Separate results MUST be reported per configuration.

### 5.1.2. Unicast ESS throughput, forwarding rate and frame loss

#### 5.1.2.1. Objective

To determine the throughput of the DUT when forwarding unicast data frames between the wireless and the wired media (i.e., between the BSS and the DS, as described in 5.2.2 of [802.11]). This test is only applicable to Access Points. The results of this test can be used to determine the ability of an Access Point to support multiple wireless clients transferring data to a wired LAN segment.

The general setup for the test comprises two or more virtual or physical clients on the wireless side of the DUT that transfer data to or from two or more virtual clients on the wired side.

#### 5.1.2.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

> Frame sizes, Frame Spacing, Signal level, Number of STAs, Fragmentation, RTS/CTS usage, Security usage and Wireless Distribution System (WDS).

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, and security not used.

#### 5.1.2.3. Procedure

The DUT is first set up according to the baseline configuration, using the initial combination of transfer direction, frame size and tester signal level. The required frame spacing is computed per Appendix A. For bidirectional tests, the tester MUST follow the half-duplex test conditions described in Section 3.5.4 on the wireless medium. The throughput, maximum forwarding rate and frame loss rate are then measured as described below. The measurements are repeated for each combination of transfer direction, frame size and tester signal level.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 2 plus the number of security modes.

Consecutive frames transmitted by the tester to the DUT MUST have different combinations of source and destination addresses, representing conversations between different sets of clients on the wired and wireless media. All possible such combinations of addresses MUST be represented equally within each trial. This distributes the load of transmission and reception uniformly among the clients. Failure to ensure this can lead to inconsistent results.

#### 5.1.2.4. Analysis and reporting

The throughput of the DUT is computed and reported (per Section 26.1 of RFC 2544 [RFC2544]) as the maximum offered load, in frames per second, resulting in zero frame loss rate [RFC1242].

The maximum forwarding rate of the DUT is computed and reported as the maximum number of test frames per second that the DUT is observed to successfully forward, irrespective of frame loss, at some value of offered load. The offered load applied to the DUT at the maximum forwarding rate MUST be reported as well.

The frame loss rate MUST be reported with the maximum forwarding rate. In this context, "rate" refers to the percentage of frames that were successfully injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as graphs of throughput and maximum forwarding rate versus each of: frame size and signal level. Separate results MUST be reported per configuration.

Note that the wired interfaces of Access Points are often capable of much higher link rates than the wireless interfaces, potentially leading to extremely high frame loss rates when transferring frames from the wired to the wireless media. Care should be taken to allow enough time for the DUT to recover and return to a normal state between trials.

### 5.1.3. Multicast forwarding rate

#### 5.1.3.1. Objective

To determine the maximum rate at which the DUT can forward multicast data frames between the wireless and the wired media (i.e., between the BSS and the DS, as described in 5.2.2 of [802.11]). This test is only applicable to Access Points. As multicast or broadcast traffic is dealt with differently from unicast traffic by the 802.11 protocol, this test therefore determines the ability of an Access Point to handle such traffic.

The general setup for the test comprises one virtual or physical client on the wireless side of the DUT that injects multicast data destined for the wireless side, as well as one virtual or physical client on the wired side that injects multicast data in the reverse direction.

Note that the 802.11 protocol does not make special provisions for multicast versus broadcast traffic. A single test is hence used to measure the ability of DUTs to handle both.

#### 5.1.3.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Frame sizes, Frame Spacing, Signal level, Number of STAs, RTS/CTS usage, Security usage, and Uni-directional transfer.

The baseline DUT configuration for performing this test consists of: RTS/CTS disabled and security not used.

#### 5.1.3.3. Procedure

The DUT is first set up according to the baseline configuration, using an initial combination of transfer direction, frame size and tester signal level. The required frame spacing is computed per Appendix A. The throughput, maximum forwarding rate and frame loss rate are then measured as described below. The measurements are repeated for each combination of transfer direction, frame size and

tester signal level.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 1 plus the number of security modes.

Either broadcast addresses, random multicast addresses, or a mixture of the two MAY be used as destination addresses for the test data. The addresses used MUST be reported with the results.

#### 5.1.3.4. Analysis and reporting

The maximum multicast forwarding rate of the DUT is computed and reported as the maximum number of test frames per second that the DUT is observed to successfully forward, irrespective of frame loss, at some value of offered load. The offered load applied to the DUT at the maximum forwarding rate MUST be reported as well.

The frame loss rate MUST be reported with the maximum forwarding rate. In this context, "rate" refers to the percentage of frames that were successfully injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as a graph of maximum forwarding rate versus each of: frame size and signal level. Separate results MUST be reported per configuration.

Note that the wired interfaces of Access Points are often capable of much higher link rates than the wireless interfaces, potentially leading to extremely high frame loss rates when transferring multicast frames to the wireless media. Care should be taken to allow enough time for the DUT to recover and return to a normal state between trials.

### 5.1.4. Forward pressure

This test measures forward pressure, as defined in Section 3.7.2 of RFC 2285 [RFC2285] and described in Section 5.6 of RFC 2889 [RFC2889].

#### 5.1.4.1. Objective

The objective of the forward pressure test is to determine if a DUT has been configured to transmit on the wireless medium at less than the minimum interframe spacing, or to transmit without adhering to the backoff rules of the 802.11 protocol. Such DUT configurations

will enable the DUT to obtain an unfair share of the available capacity of the medium. The test overloads the wireless port of the DUT, by injecting traffic into its wired interface, and measures the minimum and average interframe spacing used by the DUT to access the wireless medium.

As this test requires a minimum of two ports on the DUT, it is performed only on Access Points. Further, it cannot be performed if the aggregate bandwidth of the wired interface(s) of the DUT does not exceed that of the wireless interface.

#### 5.1.4.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Frame sizes, Frame Spacing, Number of STAs, RTS/CTS usage and Uni-directional transfer.

The baseline DUT configuration for performing this test consists of RTS/CTS disabled.

#### 5.1.4.3. Procedure

The DUT is set up according to the baseline configuration, using an initial value for frame size. A continuous stream of test frames is injected into the wired port(s) of the DUT, directed at a test client located on the wireless side. Initially, the IFS between injected frames is set according to the following equation:

IFS1 = ACKtime + SIFS + DIFS + 0.5 * CWmin * slotTime

where
IFS1 = starting IFS
ACKtime = time required to transmit 1 ACK frame on the wireless medium
SIFS = short interframe spacing for 802.11 PHY type
DIFS = DCF interframe spacing for 802.11 PHY type
CWmin = minimum value of contention window parameter for PHY
slotTime = slot time for PHY

All times are measured in microseconds. This initial value of IFS is selected to enable the DUT to forward all frames on to the wireless medium without forward pressure. The tester MUST acknowledge all frames transmitted by the DUT on the wireless medium strictly according to the 802.11 medium access rules.

The IFS between frames injected into the wired interface is then

iteratively reduced, in steps of 5 microseconds, until it reaches:

IFS2 = ACKtime + SIFS + 0.5 * DIFS

where the notation is as above.

At each iteration, the minimum and average spacing used by the DUT between the end of an ACK frame and the start of a data frame is measured by the tester.

The above process MAY be repeated with the RTS/CTS handshake enabled on the wireless side of the DUT. In this case, IFS1 and IFS2 are computed as follows:

IFS1 = RTStime + CTStime + ACKtime + 3 * SIFS + DIFS + 0.5 * CWmin * slotTime IFS2 = RTStime + CTStime + ACKtime + 3 * SIFS + 0.5 * DIFS

where
RTStime = time required to transmit 1 RTS frame on the wireless medium
CTStime = time required to transmit 1 CTS frame on the wireless medium

and the rest of the notation is as before.

#### 5.1.4.4. Analysis and reporting

Forward pressure is occurring if:

- the minimum spacing between an ACK frame sent to the DUT by the tester and an immediately succeeding RTS or data frame is less than one DIFS time for the 802.11 PHY type, or

- the average spacing between the ACK frame and the immediately succeeding RTS or data frame is less than (DIFS + 0.5 * CWmin * slotTime) for the 802.11 PHY type.

If this condition is detected, the test results MUST indicate that forward pressure is occurring. The test results MAY also be reported as a graph of minimum and average interframe spacing (between an ACK frame and the following RTS or data frame) as a function of offered load on the wired interface of the DUT.

### 5.1.5. Authentication and association rate

#### 5.1.5.1. Objective

The 802.11 protocol requires that a station wishing to transmit data to another station must authenticate itself with that station, and also establish a connection (i.e., associate with that station). The rate at which these functions can be carried out directly impacts the time taken for a wireless LAN to recover from faults and transient conditions, such as an Access Point being reset, a group of clients being turned on concurrently, or a client roaming from one Access Point to another. The objective of this test is hence to determine the rate at which a DUT can perform the authentication and association functions. This test is only applicable to Access Points.

#### 5.1.5.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Frame sizes, Signal level, Number of STAs, and Security usage.

In addition, the following test parameters MUST be configured to be the same for all trials:

Association Timeout - The tester MUST wait a predetermined amount of time for the DUT to respond to an association request with an association response. If the DUT fails to respond within this time, the association attempt MUST be considered to have failed, and a timeout error SHOULD be reported for that client. The minimum value of the association timeout MUST be at least 10 milliseconds, and MUST NOT exceed 100 milliseconds.

The association timeout used by the tester MUST be reported with the test results.

The baseline DUT configuration for performing this test consists of: security not used.

#### 5.1.5.3. Procedure

The DUT is first set up according to the baseline configuration. The tester then causes the required number of virtual or physical test clients to authenticate and associate themselves with the DUT, and measures the rate at which the DUT successfully completes authentications and associations. Each client is authenticated and associated in turn; the tester MUST NOT present a new client to the DUT until the authentication and association for the previous client has been completed. The DUT therefore determines the maximum rate at which clients can associate.

It is recommended that the authentication and association database capacity test in Section 5.3.2 be performed first to determine the maximum number of clients that can successfully associate with the DUT. The number of virtual clients presented to the DUT SHOULD be kept below this number. An authentication failure for a client SHOULD keep the tester from attempting an association for that client. Failure of the DUT to respond to an association request within the specified timeout MUST be counted as an association failure.

After the test clients successfully authenticate and associate with the DUT, the tester MUST verify that these clients have indeed been associated by causing the test clients to transmit data frames to one another, and ensure that these data frames are correctly forwarded by the DUT. The rate at which verification data frames are transmitted to the DUT MUST be well below the intra-BSS throughput supported by the DUT. The tester MUST ensure that at least one data frame originating from each client is forwarded.

If the DUT deauthenticates one or more clients during the data transfer phase, these MUST be counted as authentication failures. If the DUT disassociates one or more clients during this phase, these MUST be counted as association failures. If none of the test data frames transmitted by a physical or virtual client are forwarded successfully, this MUST be treated as a verification failure. If failures do occur, the tester MAY attempt to find a lower rate of authentication and association for which no verification failures are found for all of the test clients.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is equal to the number of security modes. Additional tests MAY be performed to determine authentication and association rates with different numbers of STAs, but the number MUST remain constant for each test run.

After each trial has been completed, the tester MUST remove the test client authentications and associations from the DUT database by performing the 802.11 deauthentication procedure for each client.

#### 5.1.5.4. Analysis and reporting

The authentication and association rate of the DUT is computed and reported as the maximum number of authentications and associations that can be successfully performed per second. Authentication, association and verification failures MUST be reported along with the test results.

If the test is performed at different signal levels, the test results SHOULD be reported as a table of signal level versus the authentication and association rate for each DUT configuration. If the test is done for different numbers of STAs, the results MAY be presented as graphs of authentication rate versus number of test clients.

### 5.1.6 Power management mode throughput, forwarding rate and frame loss

#### 5.1.6.1. Objective

To determine the throughput of the DUT when operating in power management mode. This test is applicable to clients only, and measures their ability to receive and transmit frames without loss while they attempt to conserve power.

This test is applicable to IBSS (Independent BSS) as well as infrastructure BSS client configurations. It is generally conducted using a process similar to that for the unicast intra-BSS throughput, forwarding rate and frame loss test described in Section 5.1.1 above.

#### 5.1.6.2. Test parameters

The following parameters are relevant to this test. See Section 3.5.14.

> Power Management Mode, Frame sizes, Frame Spacing, Signal level, Fragmentation, RTS/CTS usage and Security usage.

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, and security not used.

#### 5.1.6.3. Procedure

The DUT is initially set up according to the baseline configuration, using a starting combination of frame size, frame spacing and tester signal level. In addition, the DUT is placed into power-save or power management mode, such that it attempts to conserve power by shutting down its 802.11 interface when it is not transferring data. (If necessary the DUT MAY be run on batteries in order to ensure that power management mode is enabled.)

The test traffic used consists of unicast WLAN data frames that are confined to the wireless medium (and internally to the client). Test traffic is then exchanged with the DUT by the tester. The required frame spacing is computed per Appendix A. The tester MUST follow the half-duplex test conditions described in Section 3.5.4, and the client setup and test conditions described in Section 3.2.2 and

3.3.2. The throughput, forwarding rate and frame loss rate are then found as described below. The test SHOULD be repeated with different frame sizes and signal levels.

The tester MUST verify that the client enters power management mode, and SHOULD also verify that the client uses the PS Poll mechanism specified in 11.2 of the IEEE 802.11 standard [802.11] to transfer data. Failure to enter power management mode MUST be reported along with the test results.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 2 plus the number of security modes.

#### 5.1.6.4. Analysis and reporting

The throughput of the DUT is computed and reported (per Section 26.1 of RFC 2544 [2544]) as the maximum offered load, in frames per second, resulting in zero frame loss rate [RFC1242].

The maximum forwarding rate of the DUT is computed and reported as the maximum number of test frames per second that the DUT is observed to successfully forward, irrespective of frame loss, at some value of offered load. The offered load applied to the DUT at the maximum forwarding rate MUST be reported as well.

The frame loss rate MUST be reported with the maximum forwarding rate. In this context, "rate" refers to the percentage of frames that were injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as graphs of throughput and maximum forwarding rate versus each of: frame size and signal level. Separate results MUST be reported per configuration.

### 5.2. Latency and timing tests

Latency and timing tests measure the delays encountered when the DUT forwards traffic, or performs other essential tasks. These delays have significant impact on delay-sensitive protocols (such as those handling voice and video), as well as system responsiveness and network stability.

RFC 1242 [RFC1242] provides two possible definitions of latency (either the delay from last bit in to first bit out, or the delay from first bit in to first bit out). The test results MUST indicate

which definition is applicable.

### 5.2.1. Intra-BSS latency and latency variation

#### 5.2.1.1. Objective

To determine the latency and latency variation (a.k.a. jitter) exhibited by the DUT when forwarding unicast WLAN data frames that are confined to the wireless medium. This test is only applicable to Access Points.

#### 5.2.1.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

> Frame sizes, Frame Spacing, Signal level, Number of STAs, Fragmentation, RTS/CTS usage, Security usage and Wireless Distribution System (WDS).

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, security and WDS not used.

#### 5.2.1.3. Procedure

The DUT is initially set up according to the baseline configuration. The tester MUST follow the half-duplex test conditions described in Section 3.5.4. The latency and latency variation of the DUT are measured over a 1 second interval located in the middle of the trial duration, as described below. An identifying tag or signature MUST be placed in each data frame sent to the DUT during the measurement interval, so that it can be correlated with the frames received from the DUT. Note that frames transmitted to the DUT during the measurement interval can continue to be received beyond the end of the measurement interval; these frames MUST be included in the results.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 3 plus the number of security modes.

When testing Access Points with multiple physical or virtual clients, consecutive frames transmitted by the tester to the DUT MUST have different combinations of source and destination addresses, and all possible such combinations of addresses MUST be represented equally within each trial. This distributes the load of transmission and

reception uniformly among the clients. Failure to ensure this can lead to inconsistent results.

#### 5.2.1.4. Analysis and reporting

The instantaneous latency of the DUT is measured (per Section 26.2 of RFC 2544 [RFC2544]) as the difference, in seconds, between the timestamps assigned to a frame transmitted to the DUT and the corresponding frame received from the DUT. The mean of these differences in timestamps over all the data frames received from the DUT in a 1 second interval is computed and reported as the average latency of the DUT.

The latency variation of the DUT is measured and reported as the difference between the maximum and minimum instantaneous latency of all the frames received from the DUT in the same 1 second interval.

The offered load and the observed frame loss rate over this 1 second interval MUST be reported as well. In this context, "frame loss rate" refers to the percentage of frames that were injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as graphs of latency and latency variation versus each of: frame size and signal level. Separate results MUST be reported per configuration and direction of traffic flow.

### 5.2.2. ESS latency and latency variation

#### 5.2.2.1. Objective

To determine the latency and latency variation (a.k.a. jitter) exhibited by the DUT when forwarding unicast data frames between the wired and wireless media (i.e., between the BSS and the DS, as described in 5.2.2 of [802.11]). This test is only applicable to Access Points. The results of this test can be used to estimate the latency and latency variation introduced by an Access Point on delay-sensitive traffic to or from a client.

The general setup for the test comprises one or more virtual or physical clients on the wireless side of the DUT that transfers data to or from one or more virtual clients on the wired side.

#### 5.2.2.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Frame sizes, Frame Spacing, Signal level, Number of STAs, Fragmentation, RTS/CTS usage, Security usage and Uni-directional transfer.

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, and security not used.

#### 5.2.2.3. Procedure

The DUT is initially set up according to the baseline configuration. and data are transmitted to it by the tester at a constant load for the duration of the trial. The latency and latency variation of the DUT are measured over a 1 second interval located in the middle of the trial duration, as described below. An identifying tag or signature MUST be placed in each data frame sent to the DUT during the measurement interval, so that it can be correlated with the frames received from the DUT. Note that frames transmitted to the DUT during the measurement interval can continue to be received beyond the end of the measurement interval; these frames MUST be included in the results.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 2 plus the number of security modes.

When testing Access Points with multiple physical or virtual clients, consecutive frames transmitted by the tester to the DUT MUST have different combinations of source and destination addresses, and all possible such combinations of addresses MUST be represented equally within each trial. This distributes the load of transmission and reception uniformly among the clients. Failure to ensure this can lead to inconsistent results.

#### 5.2.2.4. Analysis and reporting

The instantaneous latency of the DUT is measured (per Section 26.2 of RFC 2544 [RFC2544]) as the difference, in seconds, between the timestamps assigned to a frame transmitted to the DUT and the corresponding frame received from the DUT. The mean of these differences in timestamps over all the data frames received from the DUT in a 1 second interval is computed and reported as the average latency of the DUT.

The latency variation of the DUT is measured and reported as the difference between the maximum and minimum instantaneous latency of all the frames received from the DUT in the same 1 second interval.

The offered load and the observed frame loss rate over this 1 second interval MUST be reported as well. In this context, "frame loss rate" refers to the percentage of frames that were injected into the DUT by the tester, but not forwarded by the DUT to the tester for any reason.

The test results SHOULD be reported as graphs of latency and latency variation versus each of: frame size and signal level. Separate results MUST be reported per configuration and per direction of traffic flow.

### 5.2.3. Roaming and reassociation time

#### 5.2.3.1. Objective

The 802.11 protocol enables a client to dynamically disassociate itself from one Access Point and reassociate with another Access Point in the same ESS. This is done to facilitate the mobility (or roaming) of clients within an extended region that constitutes a single logical network covered by more than one Access Points. The rate at which clients can transition from one Access Point to the next hence plays a large part in the quality and reliability of the mobile system; long roaming times can result in lost data and dropped connections. This test seeks to determine the rate at which WLAN clients and Access Points can support roaming functions.

In 802.11 networks, clients are the primary drivers of roaming behavior. A client is responsible for detecting when a roaming operation is required - for instance, because the signal from its Access Point has fallen below some threshold of acceptability - and also for locating some other Access Point and associating with it. Access Points are a contributing factor to the total roaming delay, in terms of the time required for them to accept and complete an association or reassociation with a roaming client. Client and Access Point roaming time contributions are measured separately.

#### 5.2.3.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

> Frame sizes, Number of STAs, RTS/CTS usage, RTS/CTS usage and Test Access Point beacon period.

The baseline DUT configuration for performing this test consists of: RTS/CTS disabled and security not used.

#### 5.2.3.3. Procedure

When performed on a client, this test MUST utilize two separately controllable physical or virtual Access Points belonging to the same service set (i.e., with the same SSID). Both test Access Points MAY be simulated by the same physical device, but both MUST be of similar signal strength as measured at the DUT location.

During the test, both Access Points are started simultaneously, and the DUT is allowed to authenticate with one or both of them, and then associate with one or the other of them. The association with the specific test Access Point MUST be verified by causing the DUT to transfer one or more frames of test data to it. The test Access Point with which the DUT is associated is disabled or otherwise prevented from transmitting beacons to the DUT and responding to DUT frames. The DUT should then discover that it is unable to communicate with the first test Access Point and associate with the second test Access Point. The association with the second test Access Point MUST be verified by causing the DUT to transfer one or more frames of test data to it.

When performed on an Access Point, the tester authenticates a single virtual or physical client with the DUT, and then measures the time required to perform an association procedure of the test client with the DUT (as per subclause 11.3 of IEEE 802.11 [802.11]). It then measures the time required to perform a reassociation procedure of the test client with the DUT. An authentication failure of the test client MUST keep the tester from attempting an association for the client.

In both cases, the DUT is initially set up and tested according to the baseline configuration. After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 1 plus the number of security modes.

In the case of Access Points, the test MAY be repeated with one or more additional clients associated with the DUT, in order to measure the ability of the DUT to handle associations and reassociations at various database capacities.

#### 5.2.3.4. Analysis and reporting

In the case of a client, the roaming time of the DUT is measured as the time, in seconds, from the Target Beacon Transmission Time (see 11.1.2.1 of IEEE 802.11 [802.11]) of the first missing beacon from the first test Access Point after it has been disabled, to the last bit of the Association or Reassociation Request frame transmitted to the second test Access Point by the DUT. The beacon period used or

observed for both test Access Points MUST be reported as well.

In the case of an Access Point, the association response time of the DUT is measured as the time, in seconds, from the last bit of the Association Request frame transmitted by the tester to the last bit of the Association Response frame transmitted to the tester by the DUT. The reassociation response time of the DUT is measured as the time, in seconds, from the last bit of the Reassociation Request frame transmitted by the tester to the last bit of the Reassociation Response frame transmitted to the tester by the DUT.

Separate results MUST be reported per DUT configuration. If additional clients are used for Access Point testing, the number of additional clients used in each trial MUST be reported.

### 5.2.4. Rate adaptation time

#### 5.2.4.1. Objective

Rate adaptation is used by 802.11 devices to maintain connectivity and continue to transfer data successfully (at lower rates) in spite of a decreasing signal-to-noise ratio, as may happen when mobile stations roam from one location to another. The slower rates employ more robust and error tolerant modulation formats that require less signal-to-noise ratios. This test determines the signal levels at which an 802.11 device switch from one rate to another, and also measures the time taken for the device to detect and perform the rate change.

This test can be performed on both Access Points and clients.

#### 5.2.4.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

> Frame size, Offered load, Number of STAs (in the case of Access Points), RTS/CTS usage, Fragmentation, and Security.

The baseline DUT configuration for performing this test consists of: RTS/CTS disabled, fragmentation and security not used

#### 5.2.4.3. Procedure

When performed on an Access Point, this test MUST utilize a one or more controllable physical or virtual clients to generate test traffic to and from the DUT. When performed on a client, the test MUST utilize a controllable physical or virtual Access Point with a

traffic generator capable of causing the DUT to send and receive traffic. The signal strength at the DUT location MUST be controllable in steps of 3 dB or better.

The test is performed in two parts. The first part establishes the signal levels at which the DUT switches from one PHY data rate to another. The second part uses these results to further determine the time taken for the DUT to perform these data rate steps.

At the start of the first part of the test, the necessary authentication and association procedures are used to establish a connection with the DUT (one per physical or virtual client, in the case of tests on Access Points). Data are then exchanged with the DUT at the maximum data rate corresponding to the DUT PHY type, and at the highest signal level that will not overload the DUT. Data transfer MUST be verified as taking place both from and to the DUT. The tester progressively reduces its transmitted signal level by steps of 3 dB or less while transferring data to and from the DUT, and records the PHY data rate of the frames received from the DUT for each signal level. At least 1 second of data transfer (at the configured offered load) MUST be performed at each signal level before stepping to the next. The process is continued until the DUT is found to be operating at the minimum data rate for the DUT PHY type, or until the tester has reached the lower limit of its transmit power range. The range of signal levels used MUST be reported along with the test results. The signal levels reported MUST be referenced to the DUT receiver.

During the second part of the test, the same authentication and association process is used to establish a connection with the DUT. Data are then exchanged with the DUT at the maximum data rate corresponding to the PHY type of the DUT, and at the highest signal level that will not overload the DUT. The signal level of the traffic output by the tester is then reduced to a value that is known (from the preceding part) to cause the DUT to reduce its PHY data rate to a lower value, and the time required for the DUT to detect and respond is measured. The tester SHOULD continue this process for each of the rates supported by the DUT. At least 1 second of data transfer MUST be performed at each signal level before stepping to the next. The range of signal levels used MUST be referenced to the DUT receiver and MUST be reported along with the test results.

The DUT is initially set up and tested according to the baseline configuration. After this configuration has been tested, the tester MAY repeat the test process with a new configuration, until the desired number of configurations has been exercised. The connection with the DUT SHOULD be broken (i.e., the tester disassociates and deauthenticates) and then re-established prior to starting the next

trial. The maximum number of additional test configurations is 2 plus the number of security modes.

#### 5.2.4.4. Analysis and reporting

The rate adaptation levels of the DUT are reported as the average signal levels, in dBm referenced to the DUT receiver input, that cause the DUT to change its PHY data rate from a given value to the next lower value. For example, a DUT having an 802.11b PHY (Clause 18 of IEEE 802.11 [802.11]) will have three values reported, namely, the received signal levels that cause a transition from 11 Mb/s to 5.5 Mb/s, from 5.5 Mb/s to 2 Mb/s, and from 2 Mb/s to 1 Mb/s, respectively.

The rate adaptation times of the DUT MUST be measured from the first packet transmitted by the tester at the lower signal level to the first packet received from the DUT at the lower rate. Taking the same example, these times would be measured at the 11 Mb/s - 5.5 Mb/s, 5.5 Mb/s - 2 Mb/s and 2 Mb/s - 1 Mb/s transition points, respectively.

As the process of rate adaptation is heavily influenced by RF and analog effects, this test SHOULD be performed multiple times and the results averaged, the standard deviation of the results SHOULD also be reported.

### 5.2.5. Beacon interval and timing

#### 5.2.5.1. Objective

To determine the rate at which beacons are transmitted, as well as the accuracy with which the actual transmission time of beacons matches the advertised transmission time as indicated by the DUT. A number of 802.11 network functions are affected by the rate and accuracy of beacon generation. For instance, clients in power-save mode are expected to wake up just prior to the expected transmission of a beacon from an Access Point, and remain awake until the beacon has been received and indicates that no data is pending for them. If beacons are irregular or absent, this would cause clients to either miss beacons (increasing response time) or remain awake for excessive periods (wasting power).

This test may be performed on both Access Points and clients. In the case of clients, this test is performed only if the client supports the IBSS (ad-hoc) mode of operation, as clients operating in infrastructure BSS mode do not transmit beacons.

#### 5.2.5.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Nominal beacon interval.

#### 5.2.5.3. Procedure

The DUT is configured to generate beacons at the nominal beacon interval. The tester then measures the beacon inter-arrival time and the variation in inter-arrival time over the duration of the trial, as and also captures the Beacon Interval field from the received beacon frames for comparison with the measured times.

Care MUST be taken to ensure that extraneous traffic does not occur at or near the nominal Target Beacon Transmission Times (i.e., the expected arrival time of beacons), as beacon frames are required to defer to ongoing traffic.

In the case of Access Points, the test MAY be repeated with one or more virtual or physical clients associated with the DUT, in order to measure the ability of the DUT to properly generate beacons at various database capacities.

#### 5.2.5.4. Analysis and reporting

The average beacon interval of the DUT is measured and reported as the average time, in seconds, between the first bit of consecutive beacon frames received by the tester. It MUST be computed over the duration of the trial.

The beacon interval variation is measured and reported as the difference, in seconds, between the minimum and maximum time between the first bit of consecutive beacon frames received by the tester. It MUST be computed over the duration of the trial.

The beacon interval accuracy of the DUT is measured and reported as the difference between the measured average beacon interval and the value of the Beacon Interval field of the beacon frames from the DUT, expressed as a percentage of the measured average beacon interval. The Beacon Interval field MUST be converted to seconds prior to performing the computation. If the value of the Beacon Interval field is modified by the DUT during the trial, this MUST be reported with the test results.

### 5.2.6. Reset recovery time

#### 5.2.6.1. Objective

To determine the speed with which a DUT recovers from a device or software reset. This test is only applicable to Access Points.

The rapidity with which an Access Point recovers from a reset condition affects the perceived availability and stability of a wireless network. For example, an excessive time required to recover from a reset can force clients to roam to other Access Points, cause higher-layer connections to be dropped, and so on.

#### 5.2.6.2. Test parameters

The following parameter MUST be configured prior to each trial:

Reset duration - The test SHOULD be carried out with a reset duration of at least 10 seconds.

#### 5.2.6.3. Procedure

The DUT is set up according to its baseline configuration. The tester first sets up a single virtual or physical client to serve as a test client and probe the DUT. The test client is caused to send a continuous stream of broadcast Probe Request frames to the DUT over the duration of the trial period, with a nominal interval between Probe Request frames of 5 milliseconds. The Probe Response frames from the DUT are identified and timestamped.

During the middle of the trial period, the DUT is reset. The time stamps associated with the last received Probe Response frame just prior to the reset, and the first received Probe Response frame just following the reset, are recorded. The tester MUST verify that the Probe Response frames are generated by the DUT (e.g., by comparing the BSSID of the Probe Response frames with the MAC address of the DUT).

A power-interruption reset test MUST be performed. If the DUT is capable of a software reset and/or a hardware reset, then the test SHOULD be repeated with the software and hardware resets. The results MUST be reported separately.

#### 5.2.6.4. Analysis and reporting

The reset recovery time MUST be measured and reported as the time, in seconds, between the last received Probe Response frame just prior to the reset and the first received Probe Response frame just following the reset.

## 5.3. Capacity tests

The tests described in this section measure frame storage and database related characteristics of the DUT. These tests are significant in that they quantify the ability of the DUT to handle the bursty traffic loads and large number of clients that are expected to be found in enterprise LANs. These tests are applicable only to Access Points.

### 5.3.1. Burst capacity

#### 5.3.1.1. Objective

To determine the ability of the DUT to forward bursts of back-to-back data frames typically seen in heavily loaded networks, especially when multiple clients are sending data to a single Access Point. The results are indicative of the efficiency, performance and capacity of the frame buffering functions implemented within the DUT under high load conditions.

#### 5.3.1.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

> Frame sizes, Frame spacing, Burst size, Inter-burst gap, Signal level, Number of STAs, Fragmentation, RTS/CTS usage and security settings.

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, security not used.

#### 5.3.1.3. Procedure

The DUT is initially set up according to the baseline configuration, using a starting combination of frame size, frame spacing, burst size and inter-burst gap (IBG). For a given offered load and burst size, the required IBG is computed per Appendix A. In order to assure a reasonable distribution of traffic amongst multiple sources, the number of virtual or physical test clients used in this test SHOULD be at least 8. The test SHOULD be performed with three traffic configurations: traffic flow from the wireless to the wired interface of the DUT, traffic flow from the wired to the wireless interface, and traffic flow from the wireless to the wireless interface (intra-BSS).

The tester then transmits traffic to the DUT with the configured burst characteristics, which MUST be held constant over the duration of the trial, and measures the frame loss. If the frame loss is zero, the burst size is increased, keeping the offered load constant, and

the trial is repeated. If frame loss is found, the burst size is decreased (again keeping the offered load constant) and the trial is repeated. The process continues until the maximum burst length is found that the DUT can forward without loss at a given offered load and frame size.

The tester SHOULD then repeat the test with different combinations of offered load and frame size, but with the same baseline configuration. After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional configurations is 2 plus the number of security modes. In addition, the tester MAY repeat the test with different signal levels.

The traffic generated by the tester MUST be such that consecutive frames are generated with different combinations of source and destination addresses, and all possible such combinations of addresses MUST be represented equally within each trial. Further, the tester SHOULD ensure that consecutive frames within a burst originate from different physical or virtual clients. This distributes the load uniformly among the physical or virtual clients.

#### 5.3.1.4. Analysis and reporting

The burst capacity of the DUT at a given offered load is computed and reported as the maximum number of back-to-back frames that the DUT will handle without the loss of any frames. (See Section 26.4 of RFC 2544 [RFC2544].) The results SHOULD be reported as graphs of burst capacity versus each of: offered load, frame size, and signal level. Separate results MUST be reported per configuration.

### 5.3.2. Authentication and association database capacity

#### 5.3.2.1. Objective

To determine the number of clients that a DUT can successfully support at one time. This test is only applicable to Access Points.

IEEE 802.11 WLANs implement a connection-oriented protocol with authentication and connection setup (association) being performed at the link layer. The number of clients that can be supported within the coverage area of an Access Point is thus ultimately limited by the capacity of the association database within the Access Point, even if the bandwidth requirements of the clients are well within the capacity of the device. This test therefore measures the ability of a DUT to support high concentrations of clients in a small region, such as within a conference room.

#### 5.3.2.2. Test parameters

The following parameters MUST be configured prior to each trial as specified in Section 3.5.14:

Security settings and Number of STAs.

In addition, the following test parameters MUST be configured to be the same for all trials:

Association Timeout - The tester MUST wait a predetermined amount of time for the DUT to respond to an association request with an association response. If the DUT fails to respond within this time, that association attempt MUST be considered to have failed, and a timeout error SHOULD be reported for that association attempt. The minimum value of the association timeout MUST be at least 10 milliseconds, and MUST NOT exceed 100 milliseconds.

Association Retry Limit - The tester SHOULD retry a failed association attempt (i.e., where the DUT accepts the association request but fails to respond with an association response within the specified timeout). The number of retries performed by any physical or virtual client MUST NOT exceed the pre-set association retry limit. If the number of retries reaches this limit, a retry limit error SHOULD be reported for that client.

Both the association timeout and the association retry limit MUST be reported with the test results.

The baseline DUT configuration for performing this test consists of: security not used.

#### 5.3.2.3. Procedure

The DUT is first set up according to the baseline configuration. The tester then causes the specified number of virtual or physical test clients to authenticate and associate themselves with the DUT, one client at a time, and measures the number of clients that the DUT can successfully associate. Each client is authenticated and associated in turn; the tester MUST NOT present a new client to the DUT until the authentication and association for the previous client has been completed. An authentication or association failure for a given client SHOULD not cause the tester to stop the trial. Note that the number of clients that can associate with the DUT need not equal the number of clients that can authenticate with it.

After the authentication and association of test clients with the DUT, the tester MUST verify the associations by causing the

successfully associated test clients to transmit data frames to one another, and ensure that these data frames are properly forwarded by the DUT. The rate at which verification data frames are transmitted to the DUT MUST be well below the intra-BSS throughput supported by the DUT. The tester MUST ensure that at least one data frame originating from each client is forwarded.

If the DUT deauthenticates one or more clients during the data transfer phase, these MUST be counted as authentication failures. If the DUT disassociates one or more clients during this phase, these MUST be counted as association failures. If none of the test data frames transmitted by a physical or virtual client are forwarded successfully, this MUST be treated as a verification failure. Clients for which authentication, association or verification failures have been detected MUST NOT be included in the count of successfully associated clients.

The tester SHOULD track the association identifiers (AIDs) returned by the DUT and SHOULD detect the situation where the same AID is issued to two different clients that are associated with the DUT.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is equal to the number of security modes.

The tester MUST remove the test client authentications and associations from the DUT database after the completion of each trial by performing the 802.11 deauthentication procedure for each associated client.

#### 5.3.2.4. Analysis and reporting

The authentication database capacity of the DUT is computed and reported as the maximum number of clients that can be simultaneously authenticated with it. A client that initially authenticates, but is subsequently deauthenticated by the DUT prior to the end of the trial, MUST NOT be counted towards the authentication database capacity.

The association database capacity of the DUT is computed and reported as the maximum number of clients that can simultaneously associate with it. (The association database capacity is always less than or equal to the authentication database capacity.) A client for which authentication, association or verification failures are detected during the trial MUST not be counted towards the association database capacity.

Verification failures MUST be reported along with the test results. Duplicate AIDs, if found, SHOULD be reported along with the results.

### 5.3.3. Power-save buffer capacity

#### 5.3.3.1. Objective

To measure the buffer capacity of the DUT when supporting clients in power management mode. This test is only applicable to Access Points.

Access Points that support clients in power management mode (i.e., sleeping clients) are required to accept and buffer frames on behalf of these clients, and forward them when the clients wake up. This test measures the power management mode storage capacity of the DUT, and hence its ability to support a large number of associated but sleeping clients.

#### 5.3.3.2. Test parameters

The following parameters are relevant to this test, and MUST be configured as specified in Section 3.5.14:

Frame sizes, Number of STAs, Fragmentation, RTS/CTS usage and Security usage.

In addition, the following test parameter MUST be configured prior to each trial:

Power-Save Poll Delay - this is the delay interposed between the announcement by the DUT that data are available for a given client and the retrieval of data by that client by means of a PS-Poll frame (see subclause 7.2.1.4 of IEEE 802.11 [802.11]) or other means. This delay should not exceed the frame aging time of the DUT. It SHOULD be kept the same for all trials.

The baseline DUT configuration for performing this test consists of: fragmentation off, RTS/CTS disabled, and security not used.

#### 5.3.3.3. Procedure

The DUT is initially set up according to the baseline configuration. The tester associates the required number of virtual or physical clients with the DUT, and causes these clients to enter power-save (sleep) mode. The tester MUST verify that the clients have entered power-save mode successfully and that the DUT is no longer transmitting data to the sleeping clients. The tester then transmits a predetermined number of test data frames to the DUT for forwarding to each of the sleeping clients. After all of the test data frames

have been sent to the DUT, the tester MUST wait for the power-save poll delay and then MUST cause each of the sleeping clients to wake up and retrieve their data. The number of frames received by each client from the DUT is counted.

The same number of frames MUST be transmitted to each sleeping client during a particular trial. The tester MAY associate another virtual or physical client with the DUT to serve as a means of injecting test data frames, or MAY forward test data frames via the wired interface of the DUT. The tester SHOULD verify that the DUT signals (by means of its beacons) the presence of data for a given test client before causing the client to wake up and obtain the buffered data. The tester MUST ensure that the test clients continue to poll for and retrieve buffered data from the DUT until the DUT signals (again via its beacons) that no more data are present for the clients.

The test SHOULD be repeated with different frame sizes and numbers of clients.

After the baseline configuration has been tested, the tester MAY repeat the process with a new configuration, until the desired number of different configurations have been exercised. The maximum number of such additional test configurations is 2 plus the number of security modes.

#### 5.3.3.4. Analysis and reporting

The power save buffer capacity of the DUT, for a given frame size and number of clients, is computed and reported as the total number of frames received by all of the virtual or physical clients from the DUT after they have woken up and requested their data. The results SHOULD be reported as graphs of power save buffer capacity versus each of: frame size, and number of clients. The test results MAY report the results for individual clients as well as the total for all of the clients. Separate results MUST be reported per configuration.

## 4. Security Considerations

**Documents of this type do not directly affect the security of the** Internet or of corporate networks as long as benchmarking is not performed on devices or systems connected to operating networks.

Note that performance tests SHOULD be done on with adequate isolation between the DUT and the remainder of the network, or with security systems enabled, to avoid the possibility of compromising the performance of operating networks in some manner.

## 5. IANA Considerations

There are no IANA actions requested in this memo. (Note to RFC Editor: This section may be removed upon publication as a RFC.)

## 6. References

### 6.1. Normative References

[RFC2119] Bradner, S. "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997

[RFC2544] Bradner, S. and McQuaid, J., "Benchmarking Methodology for Network Interconnect Devices", RFC 2544, March 1999.

[RFC2889] Mandeville, R. and Perser, J., "Benchmarking Methodology for LAN Switching Devices", RFC 2889, August 2000.

[RFC1242] Bradner, S., Editor, "Benchmarking Terminology for Network Interconnection Devices", RFC 1242, July 1991.

[RFC2285] Mandeville, R., "Benchmarking Terminology for LAN Switching Devices", RFC 2285, June 1998.

### 6.2. Informative References

[802.11] ANSI/IEEE Std 802.11, "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," ISO/IEC 8802-11:1999(E), ISBN 0-7381-1658-0, 1999.

[RFC1042] Postel, J. and Reynolds, J., "A Standard for the Transmission of IP Datagrams over IEEE 802 Networks," RFC 1042, February 1988.

## 7. Author's Addresses

Tom Alexander
VeriWave, Inc.
9600 Oak Street
Portland, OR, 97223
email: tom@veriwave.com
phone: +1 503 473 8358

Scott Bradner
Harvard University
29 Oxford St.
Cambridge, MA, 02138
email: sob@harvard.edu
phone: +1 617 495 3864

## Appendix A. Intended load computations

Calculating intended load for 802.11 media access is complicated by the number of different parameters that need to be accounted for as well as the random effect of backoff and management overhead. This appendix provides formulas for the theoretical maximum capacity of the media, actual intended load, and inter-burst gap.

Note that the instantaneous capacity of the 802.11 medium changes from transmission to transmission due to the effects of random backoff after each transmission. The formulas presented here are therefore expected to be applied over a large volume of traffic, rather than individual frames or bursts of frames. In addition, the parameters used in the formulas change for different 802.11 physical layers and also different data rates used within a particular physical layer.

### A.1 Calculating theoretical maximum media capacity

The theoretical maximum media capacity is calculated assuming constant-size data frames, transmitted with the minimum frame spacing according to the 802.11 protocol, with no collisions or retries occurring.

The following input parameters are defined:

LENGTH - MAC Data frame size in bytes, including FCS. For fragmented transfers, this is the size of each fragment.

SPEED - PHY data rate for the MAC portion of a DATA frame, in bits/second.

PLCPTIME - Time required to transmit the PLCP header for the given 802.11 PHY type and data rate, in seconds.

SLOTTIME - The slot time for the given 802.11 PHY type and data rate, in seconds.

DIFS - The Distributed Interframe Space (see subclause 9.2.10 of IEEE 802.11 [802.11]), in seconds.

SIFS - The Short Interframe Space (see subclause 9.2.10 of IEEE 802.11 [802.11]), in seconds.

CWmin - The minimum contention window duration (see subclause 9.2.4 of IEEE 802.11 [802.11]), in slot times.

The following intermediate values are calculated first:

TXTIME - Time required to transmit a single Data frame or fragment. For transfers that do not involve an RTS/CTS exchange, this is the time taken to transmit the Data frame plus an immediately following ACK frame (see 9.2.8 of IEEE 802.11 [802.11]). For transfers involving an RTS/CTS exchange, this is the time taken to transmit an RTS, CTS, Data and ACK frame.

For RTS/CTS based transfers:

TXTIME = (PLCPTIME * 4) + (SIFS * 3) + (((LENGTH + 48) * 8) / SPEED)

For transfers not involving RTS/CTS:

TXTIME = (PLCPTIME * 2) + SIFS + (((LENGTH + 14) * 8) / SPEED)

AMFI - Average Minimum Frame Interval. This is the minimum legal interval between the start of a Data frame and the start of the immediately following Data frame, averaged over a large number of Data frames.

AMFI = TXTIME + DIFS + ((CWmin * SLOTTIME) / 2)

The theoretical maximum capacity of the medium (abbreviated as CAP), in bits/second, is then given by:

CAP = (LENGTH * 8) / AMFI

The above formula does not take into account overhead due to management frames such as beacons and probe requests/responses. The tester SHOULD separately account for management frame overhead during a trial and subtract this overhead from the calculated theoretical capacity in order compensate for the capacity loss due to these frames.

## A.2 Calculating constant intended load

The calculations in this section deal with a constant (steady) load generated by the tester (i.e., a constant frame pattern). Burst loads are covered in the next section.

If the DUT is not to be overloaded, the intended unidirectional traffic load can range from 0 to 100% of the theoretical maximum media capacity previously calculated (0 to 50% in the case of bidirectional traffic streams). See Section 3.5.1 of RFC 2285 [RFC2285] for a full definition of Iload. For the purposes of this document, the intended load is expressed as a percentage of the

theoretical maximum media capacity, and calculated as Iload using the following formula:

Iload = (LOAD / CAP) * 100

where LOAD is the load in bits/second, and CAP is calculated as in Section A.1.

In order to actually generate traffic at Iload values less than 100%, the tester must insert extra spacing between frames to reduce the traffic load. This extra spacing is referred to here as EFG (Excess Frame Gap), and is calculated as follows:

EFG = AMFI * ((100 / Iload) - 1)

The actual frame interval therefore becomes (AMFI + EFG). The traffic pattern generated by the tester hence consists of a Data frame, the corresponding ACK frame (from the DUT), a gap equal to the DIFS plus the average minimum backoff time, and a further gap equal to EFG.

Generating Iload values greater than 100% requires that the tester violate the backoff rules of the 802.11 protocol. The tests in this document do not require Iload values greater than 100%.

### A.3 Calculating burst intended load

This section deals with the computation of intended load when the traffic pattern is bursty. A bursty pattern comprises a series of back-to-back Data/ACK exchanges separated by a DIFS, followed by a gap, followed by another series of back-to-back exchanges, and so on. The gap between bursts (referred to as the IBG) is selected based on the intended load. In addition, the IBG is calculated such that the Iload for bursty and constant traffic are directly comparable. (See Section 3.4.3 of RFC 2285 [RFC2285] for a discussion of IBG.)

The following input parameters are defined, in addition to those defined above:

BURST - Length of burst in frames.

For a given Iload, the IBG is calculated as:

IBG = DIFS + (AMFI * BURST * ((100 / Iload) - 1))

Note that the IBG is measured from the last bit of the ACK frame of the last data frame in a burst to the first bit of the preamble of the first data frame in the next burst.

## Full copyright statement



## Intellectual Property