

Ensuring 802.11n and 802.11a/b/g Compatibility

The 802.11n specification defines several mechanisms by which HT (High Throughput) STAs may coexist with non-HT (802.11a/b/g) STAs. This concept is not new, or unique to 802.11n. In fact, many of the issues of coexistence and backwards compatibility are present in 802.11g whereby protection mechanisms are employed allowing ERP (802.11g) STAs to coexist with HR/DSSS (802.11b) STAs. To be sure, the designers of 802.11n certainly understood that in order for the technology to be successfully adopted, it must contend and coexist with legacy WLAN deployments – especially in the enterprise, where significant investment has already been made in (legacy 802.11 a/b/g) WLAN infrastructure. There should be mechanisms for the legacy and 802.11n stations to understand each other and protect themselves from interference created by each other.

This paper takes a detailed look at the impact of coexistence on 802.11n data rate and throughput.

[Table of contents](#)

CSMA/CA and Virtual Carrier Sense	2
802.11 PHYs	2
802.11n Preamble Types	3
802.11n Protection Rules	4
802.11n Protection Frame Exchanges	6
RTS/CTS Frame Exchange	7
CTS-to-self Frame Exchange	7
L-SIG TXOP	8
Conclusion	8

CSMA/CA and Virtual Carrier Sense

CSMA/CA stands for “Carrier Sense Multiple Access with Collision Avoidance”. It is the mechanism by which 802.11 STAs share the transmission medium. An 802.11 STA is not able to transmit and receive at the same time; and this makes it impossible for the STA to directly know that while it is transmitting, that another STA is transmitting at the same time (a collision has occurred). 802.11 STAs use positive acknowledgement (ACK frames) in order to indirectly infer whether or not a collision has occurred. A STA assumes that no collision (or other error) has occurred when an ACK frame is received from the recipient. The ACK frame basically says, “I received your transmission, without error”. When (after a small timeout period) no ACK frame is received from the recipient, an 802.11 STA assumes that a collision (or some other error) has occurred and reschedules the frame for re-transmission.

According to the rules of CSMA/CA and the 802.11 DCF (Distributed Coordination Function), a STA that wishes to transmit, must first “listen” on the channel for a period defined by the DIFS (DCF Interframe Space). If the channel has been sensed as “idle” for this time period (that is, no STA within range is transmitting), the STA may initiate transmission. If before, during or after listening the DIFS period, the channel has been sensed as “busy” (that is, another STA is transmitting), the STA must defer its transmission and “wait” a random back-off interval before trying again. A STA must also perform this random back-off after a transmission as well. When a STA detects a collision (that is, no ACK frame is received), it increases the minimum time it will wait during the random back-off interval. It is in this way that TxOP (transmit opportunity) is distributed to all STAs present in a (DCF) wireless network and individual STAs are prevented from monopolizing the medium.

802.11 supplements CSMA/CA with Virtual Carrier Sense mechanisms. Frames exchanged between STAs carry with them the duration of the frame exchange (including the expected ACK); alerting all other STAs how long a particular transmission will occupy the medium. In this way, STAs know (in advance) how long the current frame exchange will keep the channel busy, and they defer their own transmissions accordingly.

The RTS/CTS (Request-to-send/Clear-to-send) frame exchange addresses the “**hidden node**” problem, where STA A can hear the AP, but not STA B. In this case, there exists the possibility of STA A transmitting to AP at the same time STA B is transmitting to AP (a collision). If STA B sends an RTS frame, and waits for AP to transmit a CTS frame, STA A can hear the CTS frame and defer transmission accordingly, avoiding a collision. RTS/CTS frames carry with them the duration of the complete frame exchange (including the data frame and ACK to follow). STAs update their NAV (Network Allocation Vector – which is essentially a “counter” for how long the medium will be busy) whenever they receive a frame indicating a duration larger than their current NAV value.

All 802.11 frames are preceded by a **radio preamble** and **PLCP header**. The **radio preamble** is used by the receiver to fine tune the manner in which it will receive the transmission. It allows the receiver to intelligently adjust gain settings, and apply frequency and timing correction algorithms, such that the frame will be processed as accurately as possible. The **PLCP (Physical Layer Convergence Procedure) header** following the radio preamble, tells the receiver important information about the transmission. Amongst other things, it indicates the modulation type, coding rate, and length of the transmission. It is the detection and receipt of these radio preambles and PLCP headers that tell the STA whether or not the medium is busy.

802.11 PHYs

For the purposes of this study, it is useful to think of STAs in terms of the PHY (physical) layer that is in use by them, as opposed to their 802.11 amendment designator(s). 802.11 has defined seven PHYs as of 802.11n. Ignoring the IR (Infrared) and FHSS (Frequency Hopping Spread Spectrum) PHYs, they are:

- Clause 15 (DSSS PHY for 2.4 GHz, defined in the original 802.11 specification)
- Clause 17 (OFDM PHY for 5 GHz, defined in the 802.11a amendment)
- Clause 18 (HR/DSSS PHY for 2.4 GHz, defined in the 802.11b amendment)
- Clause 19 (ERP PHY for 2.4 GHz, defined in the 802.11g amendment)
- Clause 20 (HT PHY for 2.4 and 5 GHz, defined in the 802.11n amendment)

Each of these PHYs define the operating band(s), radio preamble / PLCP format(s), modulation types and operating rules that a STA may use. It is important to understand PHY compatibilities, particularly the radio preambles, since they are used to sense the medium's condition as busy or not. The point here is that devices with incompatible PHYs can not sense each other as occupying the medium – which would clearly lead to a break-down in the CSMA/CA system. 802.11 copes with this by mandating that later-amendment PHYs (operating in the same band as previously defined PHYs) must also support those previously defined PHYs:

- STAs implementing the Clause 18 PHY, must also implement the Clause 15 PHY.
- STAs implementing the Clause 19 PHY, must also implement the Clause 18 and 15 PHY.
- STAs implementing the Clause 20 PHY (in 2.4 GHz), must also implement the Clause 19, 18 and 15 PHY.
- STAs implementing the Clause 20 PHY (in 5 GHz), must also implement the Clause 17 PHY.

This way backwards-compatibility (and thus, a properly functioning CSMA/CA system) is maintained across amendments to the standard. A STA with only the Clause 19 PHY implemented, can not interpret transmissions from (or sense) a STA using the Clause 20 PHY. That is, unless it first uses the Clause 19 PHY to “protect” the Clause 20 PHY transmission. The STA need not care that the Clause 19 STA can not interpret the Clause 20 transmission; it only cares that the Clause 19 STA has enough information to know how long to defer itself from transmitting. This is what is meant by protection mechanisms. Basically, STAs need to transmit “protection frames” such as RTS/CTS, or CTS-to-self, at rates defined by a Clause which all STAs present can interpret.

802.11n Preamble Types

The Clause 20 HT PHY addresses coexistence and protection right from the start. Three preamble types are defined:

- Non-HT Preamble (identical to the OFDM preamble used in 802.11 a/g)
- HT Mixed Format Preamble (an HT preamble preceded by a non-HT preamble)
- HT Greenfield Format Preamble (a purely HT preamble)

The HT (High Throughput) Mixed Format preamble is mandatory; and is used whenever there are non-HT stations present. The Non-HT Preamble is also mandatory, providing legacy 802.11 a/b/g operation. The Greenfield Preamble is optional, and may be used when all STAs present support (HT) Greenfield operation.

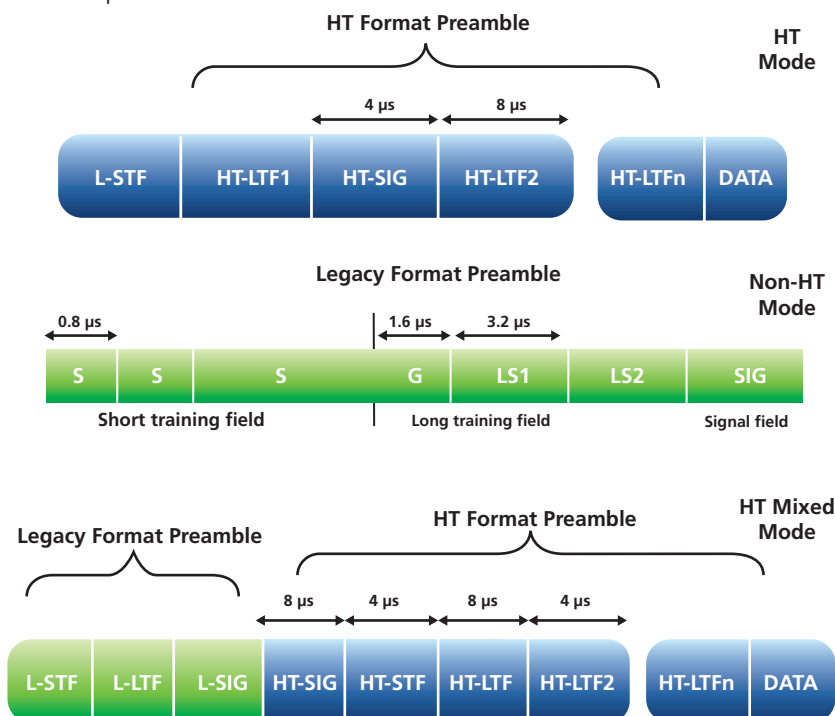


Figure 1. HT, legacy and mixed-mode preambles

802.11n Protection Rules

A Clause 20 STA (hereinafter referred to as an HT STA) advertises information about what types of STAs are observed to be present, using the Beacon and Probe Response frames. These frames carry the HT Information Element, which include the following fields:

- Operating Mode
- Non-Greenfield STAs Present
- OBSS Non-HT STAs Present

The **Operating Mode field** indicates the operating mode of the BSS, and may take one of four possible values:

- 0: All STAs in the BSS are HT STAs (no non-HT STAs are present in the BSS)
- 1: Non-HT STAs are present in the primary and/or the secondary channel
- 2: All STAs in the BSS are HT STAs, however at least one STA supports only 20 MHz operation
- 3: One or more non-HT STAs are present in the BSS

The **non-Greenfield STAs present field** indicates whether or not all HT STAs that are associated are Greenfield capable. The **OBSS Non-HT STAs present field** indicates whether or not an overlapping BSS has non-HT STAs associated to it.

In addition to the fields present in the HT Information Element, STAs use the ERP Information Element to determine at what rates protection frames may be sent. When the Use_Protection bit field is set to 1, it indicates that Clause 15 or Clause 18 STAs are present (and thus, protection frames may not be sent at ERP (Clause 19) data rates). It is the combination of both the HT Information Element and the ERP Information Element which determine what type is STA is the “least capable” in the BSS – which allows STAs to know which PHY(s) may be used for protection frames.

For completeness, STAs are directed to “propagate” whatever knowledge they have about what other STAs are present, by setting the field described appropriately, even if the AP has not yet done so. In fact, an AP is directed to set its own operating mode based upon what it is able to observe, as well as what other STAs have indicated they have observed. In this way, information about the presence of non-HT, or non- ERP STAs “ripples” throughout the BSS (and any overlapping BSS).

802.11n Protection Frame Exchanges

Depending on the type of HT transmission and the values contained in the HT and ERP Information Elements, several protection frame exchanges are allowed:

- RTS/CTS (Request-to-send/Clear-to-send)
- CTS-to-self
- Using a legacy or mixed mode format preamble, transmit a frame which requires a response frame
- L-SIG TXOP Protection

Now let us examine the impact on data rate (throughput) by RTS/CTS, CTS-to-self and L-SIG TXOP protection frame exchanges.

RTS/CTS Frame Exchange

In the following examples, our transmitting and receiving HT STAs are using 40 MHz wide channels, with short guard interval and two spatial streams (full 2x2); yielding PHY data rates of 300 Mbps. When using RTS/CTS to protect the 40 MHz HT transmissions, non-HT duplicate format frames are used and this essentially equates to using legacy 6 Mbps frames on each of the 20 MHz halves of the 40 MHz channel.

Using a payload size of 1,500 bytes, we may calculate a total frame exchange time of just over 300 μ sec¹ (X-axis in figure below); which gives us a maximum link layer throughput of 36.84 Mbps. The following figure illustrates the frame exchange.

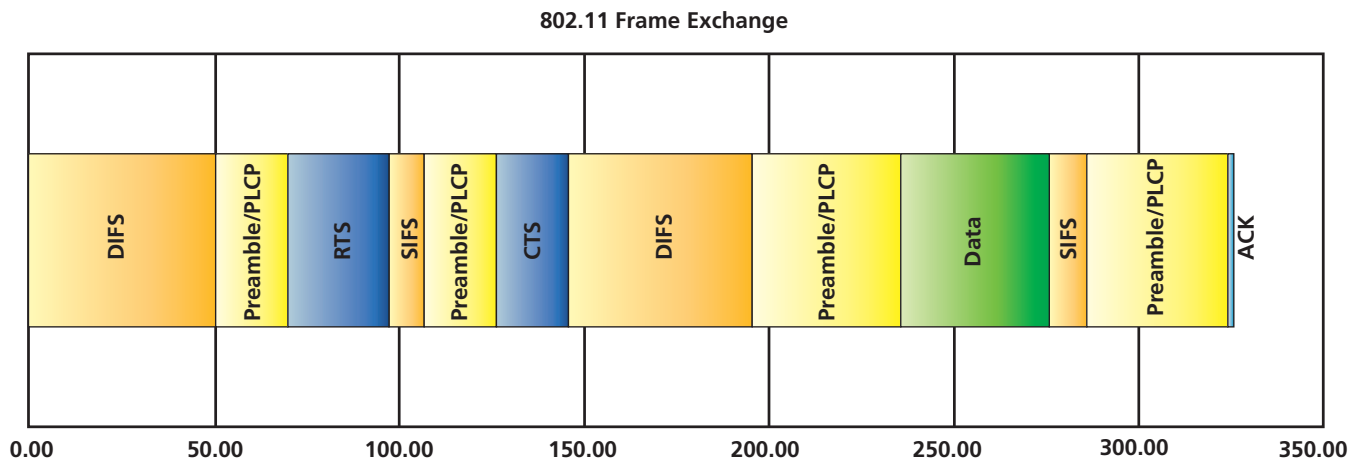


Figure 2: 300 Mbps PHY Data Rate RTS/CTS Frame Exchange (No Frame Aggregation)

We may further calculate an overhead of 88% by way of Interframe Spacing (depicted in orange), Radio Preamble/ PLCP (depicted in yellow) and control frames (in blue).

For completeness, let us examine the same frame exchange with full frame aggregation (65535 bytes); pictured in the following figure. As can be expected, our maximum link layer throughput is increased – to nearly 258 Mbps, however, we still incur an overhead of 14%. This illustrates that even in the best-case scenario in terms of aggregation efficiency (which is not usually the case), the protection overhead is still significant.

¹ Frame exchange time calculated based on values defined by the 802.11 specification

802.11 Frame Exchange

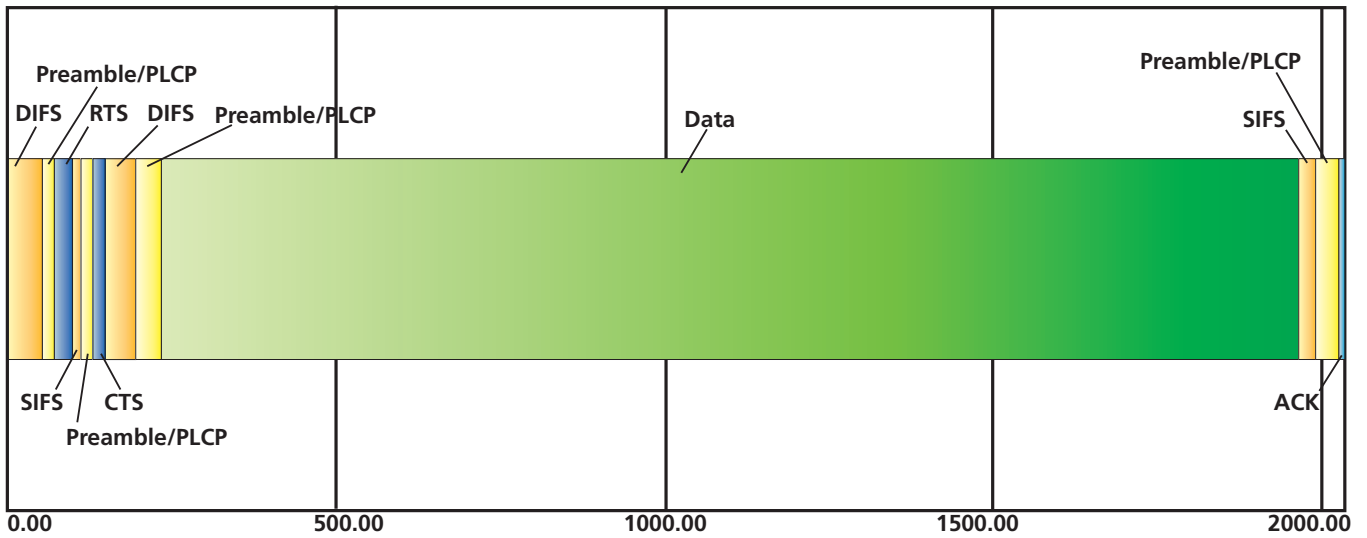


Figure 3: 300 Mbps PHY Data Rate RTS/CTS Frame Exchange (With Frame Aggregation)

Issues of efficiency aside, the RTS/CTS exchange does provide a more robust non-HT signaling mechanism as compared to CTS-to-self, or L-SIG TXOP. RTS/CTS provides at least two NAV-based (MAC layer) and three PHY layer signals to non-HT STAs; as well as addressing the “hidden node” problem.

In the next examples, our transmitting and receiving HT STAs are using 20 MHz wide channels, with short guard interval and two spatial streams (full 2x2); yielding PHY data rates of 144.44 Mbps. When using RTS/CTS to protect the 20 MHz HT transmissions, legacy (clause 17 or 19) format frames are used. In the case of clause 19 protection frames at 6 Mbps, the overhead is essentially the same as in the non-HT duplicate format case. Thus, our calculations are modified only for the rate at which the data and ACK frames are being set.

Using a payload size of 1,500 bytes, we may calculate a total frame exchange time of 369 μsec (X-axis in figure below); which gives us a maximum link layer throughput of 32.50 Mbps. The following figure illustrates the frame exchange.

802.11 Frame Exchange

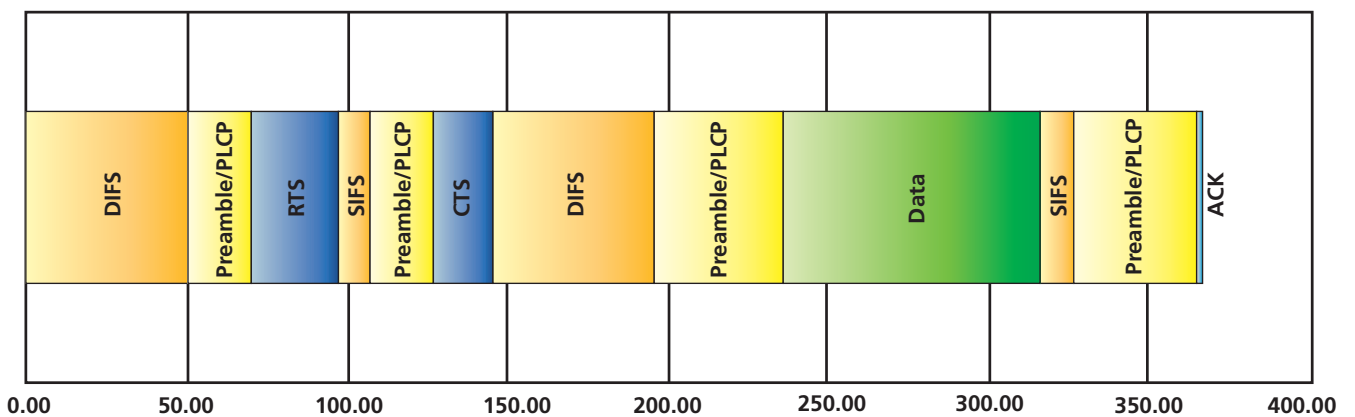


Figure 4: 144 Mbps PHY Data Rate RTS/CTS Frame Exchange (No Frame Aggregation)

We incur an overhead tax of 77%. Again, let us compare the same frame exchange at full frame aggregation. Pushing 65k bytes in our data frame gets us to a maximum link layer throughput of 133.90 Mbps, an overhead of about 7%. The following figure illustrates the frame exchange.

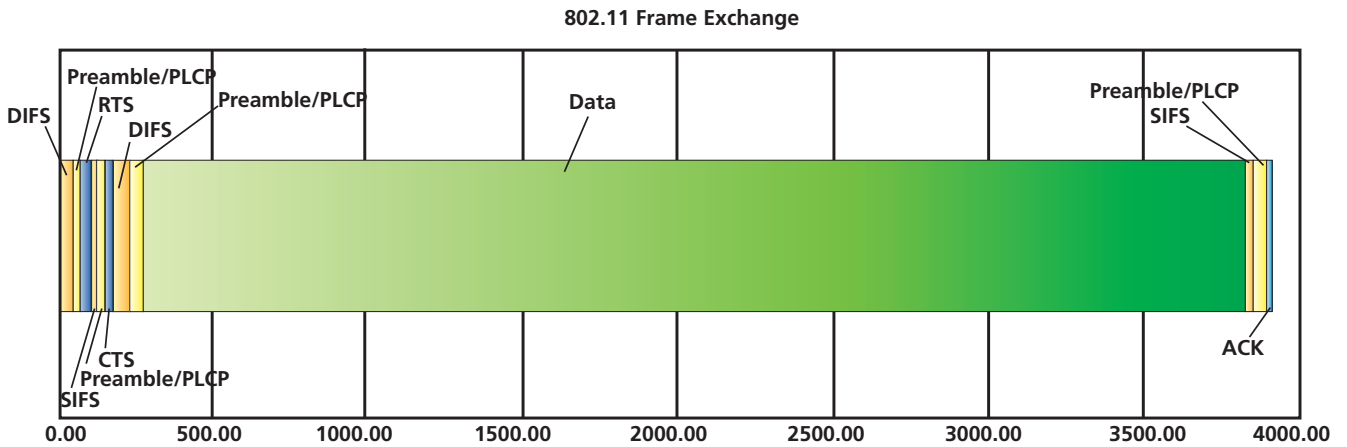


Figure 5: 144 Mbps PHY Data Rate RTS/CTS Frame Exchange (With Frame Aggregation)

CTS-to-self Frame Exchange

The CTS-to-self frame exchange, while more efficient in terms of link layer throughput, does not provide as many NAV (MAC layer) or PHY based non-HT signaling mechanisms; nor does it address the “hidden node” problem. This protection method allows 802.11n devices to transmit a CTS frame to itself and ensure that the neighboring legacy devices will use the timing information to protect following 802.11n frames. The CTS frame must be transmitted using legacy data rates. The following figure illustrates the CTS-to-self frame exchange, which is essentially the same as the RTS/CTS exchange, sans the RTS (and its associated IFS and preambles).




Figure 6: 300 Mbps PHY Data Rate CTS-to-Self Frame Exchange (No Frame Aggregation)

In the above figure, our transmitting and receiving HT STAs are using 40 MHz wide channels, with short guard interval and two spatial streams (full 2x2); yielding PHY data rates of 300 Mbps. The 802.11 PHY (physical) data rate is the speed at which data is transferred over the wireless communications medium; usually expressed in terms of Mbps (mega bits per second). Using a payload size of 1,500 bytes, we may calculate a total frame exchange time of 269 μsec (X-axis in figure below); which gives us a maximum link layer throughput of 44.60 Mbps, or 85% overhead. This gives us an additional 7.76 Mbps of throughput as compared to the same STAs using RTS/CTS, or nearly 3% more of our PHY Data Rate.

For completeness, let us examine CTS-to-self when used with the highest PHY Data Rate offered by 802.11n. At 600 Mbps (40 MHz, 4 spatial streams, SGI) we incur an overhead tax of 92%. As can be seen in the following figure the total transaction takes 248.86 μ sec (X-axis in figure below), of which only 40 μ sec is spent on data bits. We get a maximum link layer throughput of 48.22 Mbps. CTS frames are transmitted using legacy data rates.

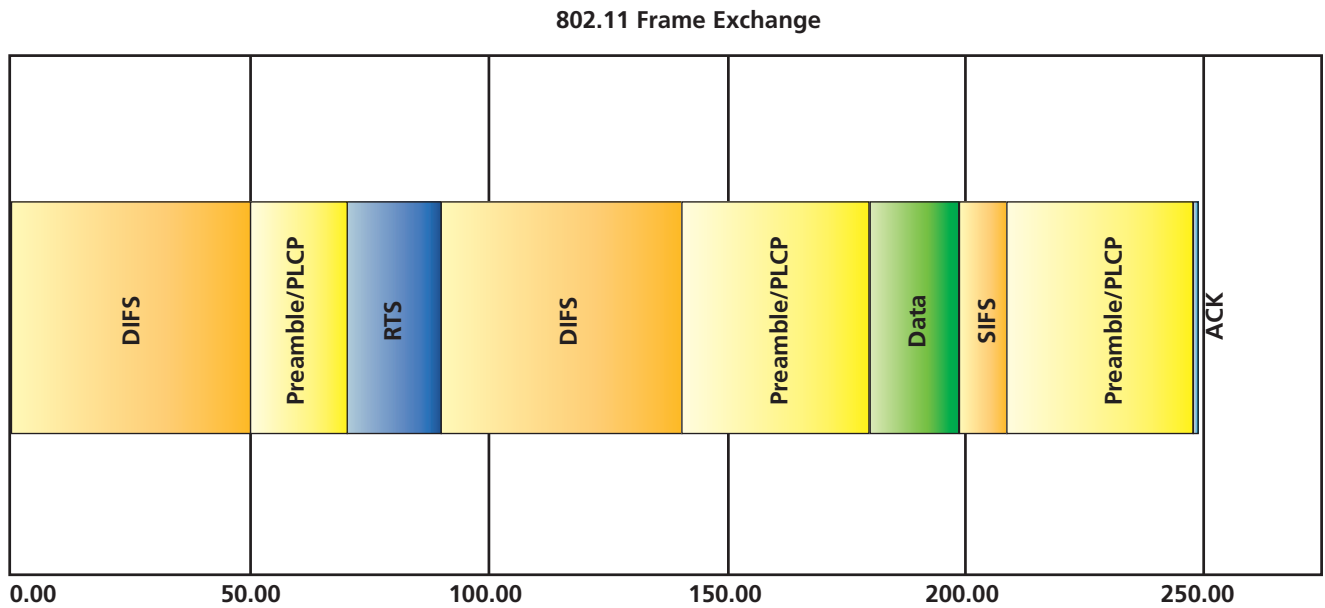


Figure 7: 600 Mbps PHY Data Rate CTS-to-Self Frame Exchange (No Frame Aggregation)

Even if we factor in full frame aggregation, we still incur a 21% overhead tax. As can be seen in the following figure, at 65k bytes per data frame, the frame transaction takes 1.1 msec (X-axis in figure below), yielding a maximum link layer throughput of 475.47 Mbps. The point here being, that even if the more efficient CTS-to-self protection mechanism is used, the higher the PHY Data Rate, the higher the penalty is in terms of overhead.

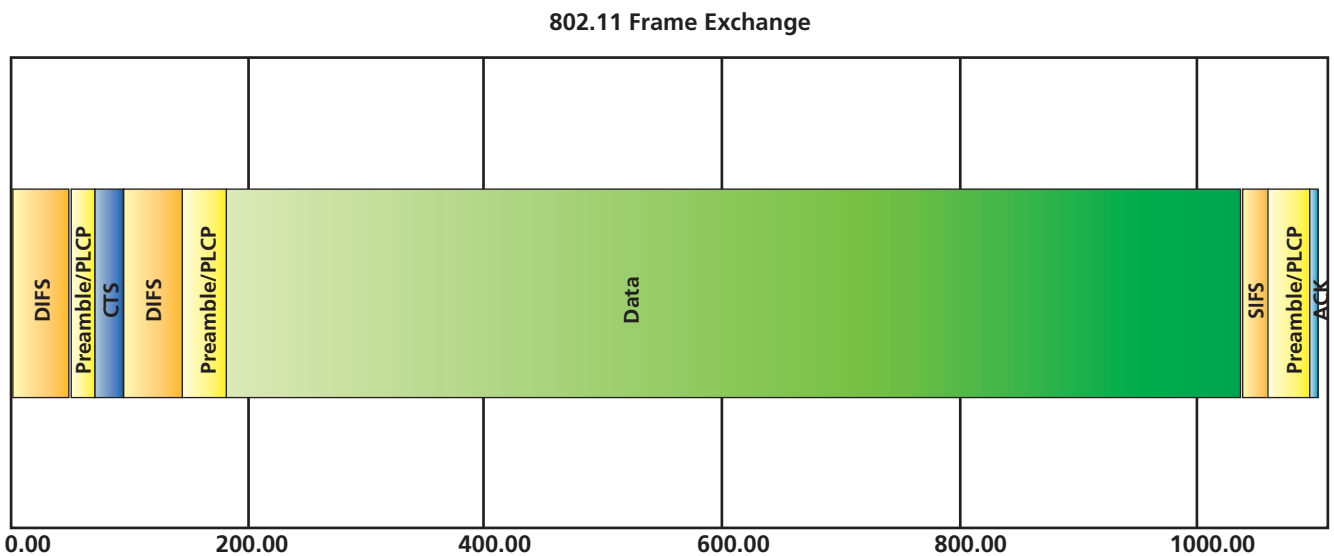


Figure 8: 600 Mbps PHY Data Rate CTS-to-Self Frame Exchange (With Frame Aggregation)

L-SIG TXOP

Unlike RTS/CTS or CTS-to-self, L-SIG TXOP protection does not incur additional protection frame overhead at the MAC layer. L-SIG TXOP protection is achieved by “spoofing” the legacy format portion of the PLCP, such that its duration covers the HT exchange. When compared to Greenfield, there is a 16 μsec penalty paid (for the legacy portion of the mixed-mode preamble). On first look, this clearly leads to a much more efficient protection mechanism – however, it should be noted that L-SIG TXOP provides no NAV (MAC layer) based input to the legacy STAs that are present. This leads to a greater probability for packet collisions. It is probably for this reason that 802.11n specifically states that “implementers are advised to include a NAV based fallback mechanism”.




Figure 9: Basic LSIG-TXOP Operation

Conclusion

802.11n defines several mechanisms by which HT and non-HT STAs can coexist with each other. At this time, the penalties paid for coexistence are probably well-justified expenses for the prize of maintaining backwards-compatibility with previous generations of hardware. However, once application layer throughput demands the higher bandwidth offered by 802.11n, and/or HT technology matures enough to achieve the highest data rates defined by the standard, and/or its deployment becomes ubiquitous, the cost of coexistence may be more fully factored into deployment strategy. With the exception of greenfield deployments, protection overhead is likely to be a part of the HT enabled WLAN for the immediate future. However, by understanding these protection mechanisms and their cost, WLAN managers may be able to minimize the amount of protection overhead present in their networks.

Protection Mechanism	Channel Bandwidth	Aggregated Frame Size	PHY Data Rate	Link Layer Throughput	Protection Overhead
RTS-CTS	40 MHz	1500 Bytes	300 Mbps	36.84 Mbps	88%
RTS-CTS	40 MHz	65535 Bytes	300 Mbps	258 Mbps	14%
RTS-CTS	20 MHz	1500 Bytes	144 Mbps	32.5 Mbps	77%
RTS-CTS	20 MHz	65535 Bytes	144 Mbps	133.9 Mbps	7%
CTS-to-Self	40 MHz	1500 Bytes	300 Mbps	44.60 Mbps	85%
CTS-to-Self	40 MHz	1500 Bytes	600 Mbps	48.22 Mbps	92%
CTS-to-Self	40 MHz	65535 Bytes	600 Mbps	475.47 Mbps	21%
L-SIG TXOP					<0.1%

Figure 10: 802.11n Protection Mechanism summary

About Fluke Networks' Wireless Portfolio

Fluke Networks delivers a complete portfolio of wireless solutions to ensure the security, performance and compliance of wireless networks. Innovative solutions support the entire wireless LAN lifecycle, from planning and deployment to ongoing troubleshooting and dedicated security monitoring, combating all wireless issues that threatens business uptime and integrity.

Contact Fluke Networks: Phone **408-753-1500** or **Email: info@flukenetworks.com**.

Fluke Networks
P.O. Box 777, Everett, WA USA 98206-0777

Fluke Networks operates in more than 50 countries worldwide. To find your local office contact details, go to www.flukenetworks.com/contact.

©2011 Fluke Corporation. All rights reserved.
Printed in U.S.A. 10/2011 4111466A D-ENG-N