IEEE P802.11
Wireless LANs

TGn Sync Proposal Technical Specification

Date: 2005-01-18

Author:

Name	Company	Address	Phone	email
Syed Aon Mujtaba	Agere Systems Inc.	555 Union Blvd., Allentown, PA 18109, U.S.A.	+1 610 712 6616	mujtaba@agere.com

Abstract

This document presents the technical specification for the
MAC and the PHY layer of the TGn Sync proposal to IEEE 802.11 TGn

Notice: This document has been prepared to assist IEEE 802.11. It is offered as a basis for discussion and is not binding on the contributing individual(s) or organization(s). The material in this document is subject to change in form and content after further study. The contributor(s) reserve(s) the right to add, amend or withdraw material contained herein.

Release: The contributor grants a free, irrevocable license to the IEEE to incorporate material contained in this contribution, and any modifications thereof, in the creation of an IEEE Standards publication; to copyright in the IEEE's name any IEEE Standards publication even though it may include portions of this contribution; and at the IEEE's sole discretion to permit others to reproduce in whole or in part the resulting IEEE Standards publication. The contributor also acknowledges and accepts that this contribution may be made public by IEEE 802.11.

Patent Policy and Procedures: The contributor is familiar with the IEEE 802 Patent Policy and Procedures <http://ieee802.org/guides/bylaws/sb-bylaws.pdf>, including the statement "IEEE standards may include the known use of patent(s), including patent applications, provided the IEEE receives assurance from the patent holder or applicant with respect to patents essential for compliance with both mandatory and optional portions of the standard." Early disclosure to the Working Group of patent information that might be relevant to the standard is essential to reduce the possibility for delays in the development process and increase the likelihood that the draft publication will be approved for publication. Please notify the Chair <stuart.kerry@philips.com> as early as possible, in written or electronic form, if patented technology (or technology under patent application) might be incorporated into a draft standard being developed within the IEEE 802.11 Working Group. **If you have questions, contact the IEEE Patent Committee Administrator at <patcom@ieee.org>.**

INTEL-1012
9,973,361

Additional Authors:				
Name	**Company**	**Address**	**Phone**	**Email**
Adrian P Stephens	Intel Corporation			adrian.p.stephens@intel.com
Aleksandar Purkovic	Nortel Networks			apurkovi@nortelnetworks.com
Bo Xia	Intel Corporation			bo.xia@intel.com
Brian Hart	Cisco Systems Inc			brianh@cisco.com
Chiu Ngo	Samsung Electronics Co. Ltd.			cngo@sisa.samsung.com
Daisuke Takeda	Toshiba Corporation			daisuke.takeda@toshiba.co.jp
Darren McNamara	Toshiba Corporation			Darren.McNamara@toshiba-trel.com
Dongjun Lee	Samsung Electronics Co Ltd			djthekid.lee@samsung.com
Eldad Perahia	Cisco Systems Inc			eperahia@cisco.com
Eric Jacobsen	Intel Corporation			eric.a.jacobsen@intel.com
Gyubum Kyung	Samsung Electronics Co Ltd			gyubum.kyung@samsung.com
Huanchun Ye	Atheros Communications Inc			hcye@atheros.com
Huaning Niu	Samsung Electronic Co Ltd			Huaning.Niu@samsung.com
Hui-Ling Lou	Marvell Semiconductors			hlou@marvell.com
Jan Boer	Agere Systems Inc			janboer@agere.com
Jari Jokela	Nokia			jari.jokela@nokia.com
Jeffrey Gilbert	Atheros Communications Inc			gilbertj@atheros.com
Jerry Kim	Samsung Electronic Co Ltd			kimjy@samsung.com
Job Oostveen	Royal Philips Electronics			job.oostveen@philips.com
Jörg Habetha	Royal Philips Electronics			joerg.habetha@philips.com
John Sadowsky	Intel Corporation			john.sadowsky@intel.com
Luke Qian	Cisco Systems Inc			lchia@cisco.com
Masahiro Takagi	Toshiba Corporation			masahiro3.takagi@toshiba.co.jp
Michael Livshitz	Nortel Networks			livshitz@nortelnetworks.com
Monisha Ghosh	Royal Philips Electronics			monisha.ghosh@philips.com
Nina Burns	Nortel Networks			nburns@nortelnetworks.com
Pen C. Li	Royal Philips Electronics			pen.li@philips.com
Pieter-Paul Giesberts	Agere Systems Inc			pgiesberts@agere.com
Richard van Leeuwen	Agere Systems Inc			rleeuwen@agere.com
Ronald Rietman	Royal Philips Electronics			ronald.rietman@philips.com
Seigo Nakao	SANYO Electric Co Ltd			snakao@gf.hm.rd.sanyo.co.jp
Sergey Sukobok	Nortel Networks			ssukobok@nortelnetworks.com
Sheung Li	Atheros Communications Inc.			sheung@atheros.com
Stephen Shellhammer	Intel Corporation			Stephen.j.shellhammer@intel.com
Tomoko Adachi	Toshiba Corporation			tomo.adachi@toshiba.co.jp
Tomoya Yamaura	Sony Corporation			yamaura@wcs.sony.co.jp
Tsuguhide Aoki	Toshiba Corporation			tsuguhide.aoki@toshiba.co.jp
Won-Joon Choi	Atheros Communications Inc.			wjchoi@atheros.com
Xiaowen Wang	Agere Systems Inc			xiaowenw@agere.com
Yasuhiko Tanabe	Toshiba Corporation			yasuhiko.tanabe@toshiba.co.jp
Youngsoo Kim	Samsung Electronics Co., Ltd			KimYoungsoo@samsung.com
Yuichi Morioka	Sony Corporation			morioka@wcs.sony.co.jp

1 Executive Summary

The TGn Sync team respectfully submits its complete proposal. The team is comprised of 802.11 WG participants whose companies represent a broad range of markets, including, PC, Enterprise, Consumer Electronics, Handset, Semiconductor and Public Access. Since these companies are based in the US, Europe and Asia Pacific, the team has had the benefit of a worldwide perspective, both in terms of perceived market demand as well as regulatory concerns (including channel bandwidth limitations, limited channel access, restricted transmission times, etc.).

Our goal has been to align (or "sync" with) 802.11n members to a single set of core positions prior to the TGn down selection phase. By aligning before entrenched positions form, we hope to increase the possibility of a more rapid introduction of an 802.11n standard.

The team has worked as a technical group for many months to develop a TGn proposal, recognizing that whatever decisions were made had to ultimately survive the TGn process and therefore be based solely on technical merit. There were no signed agreements involved in joining the TGn Sync team. There was no attempt to develop joint IP and each company was expected to abide by the same reasonable and non-discriminatory terms that the IEEE bylaws require.

The technical approach balances market expectations with design practicality. The solution is a robust, scalable architecture targeting the least amount of complexity. In fact, the basic configuration delivers a maximum mandatory rate of 243 Mbps using only two antennas. This rate is consistent with the historical trend of 5x with each generation of 802.11 (802.11/2Mbps, 802.11b/11Mbps, and finally 802.11a/ 54Mbps). The proposal also incorporates options for higher rates beyond 600 Mbps. This choice of higher peak data rates was seen as critical for future proofing this next generation of WLAN.

The PHY techniques used to achieve the higher data rates involve a MIMO evolution of 802.11 OFDM PHY with spatial division multiplexing of spatial streams, wider bandwidth options (either 20MHz or 40MHz, but not required in regulatory domains where prohibited) and an optimized interleaver for both the 20 MHz and 40MHz channelizations. Additionally, optional enhancements include advanced FEC coding techniques (LDPC), transmit beamforming with negligible additional cost in the receiving client device and/or additional RF chains. These options provide extra scalability if additional robustness and/or throughput is required.

The TGn Sync proposal also offers seamless interoperability with 802.11 legacy devices. This interoperability is achieved with an enhanced 802.11 preamble design and efficient PHY and MAC level mechanisms, which also provide robustness and cost-effectiveness.

The features added to improve MAC efficiency – for higher throughput and overall system performance – include aggregation, bi-directional data flow, enhanced Block Ack, power management (to reduce power consumption), channel management (including a receiver assisted channel training protocol) and feedback mechanisms that enable rate adaptation. Protection mechanisms are also added to achieve the seamless interoperability and coexistence with legacy

devices mentioned above. The approach supports 802.11e and proposes new features for further enhancement.

1.1 Proposal Structure

The TGn Sync complete proposal meets the IEEE 802.11n PAR, meets all functional requirements, and addresses all mandatory requirements of the comparison criteria.

This document contains a technical specification of the proposed MAC and PHY enhancements.

[5] contains the presentation of the proposal in powerpoint format. (Please note that the summary presentation [4] is no longer up to date).

[6] contains a statement of compliance to the IEEE 802.11 TGn FRCC requirements.

[7] and [8] contain FRCC results (and others) for the PHY and MAC respectively.

[9] and [10] are Excel spreadsheets containing detailed FRCC results from MAC simulations, and [11] describes the methodology used in these simulations.

Table of Contents

1 Executive Summary 3
1.1 Proposal Structure 4
2 References 13
3 Definitions 14
4 Abbreviations and acronyms 15
5 Introduction 19
5.1 Purpose 19
5.2 General description of MAC enhancements 19
5.3 General description of PHY enhancements 20
6 Frame Formats 20
6.1 Control frames 21
6.1.1 Initiator Aggregation Control (IAC) MPDU 21
6.1.2 Responder Aggregation Control (RAC) MPDU 25
6.1.3 Multiple Receiver Aggregate Descriptor (MRAD) MPDU 28
6.1.4 Increase Channel Bandwidth (ICB) MPDU 29
6.1.5 Decrease Channel Bandwidth (DCB) MPDU 30
6.1.6 Block Ack (BA) MPDU 30
6.1.7 Calibration Training MPDU 32
6.2 Data frames 33
6.2.1 MAC Header (MHDR) MPDU 33
6.2.2 Compressed Header Data (CHDATA) MPDU 34
6.3 Management frame formats 35
6.3.1 HT Capabilities 35
6.3.2 Additional HT Information Elements 36
6.3.3 Management Action Frames 38
6.3.4 TSPEC Modification to support Periodic RDR 41
6.3.5 Block Ack Parameter Set field 41
6.4 A-MPDU and A-MSDU aggregation formats 42
6.4.1 Introduction to Aggregation (Informative) 42
6.4.2 Aggregated MPDU format (A-MPDU) 42
6.4.3 Single receiver frame aggregation format 44

6.4.4 Multiple receiver frame aggregation format 45
6.4.5 Aggregated MSDU Format (A-MSDU)..... 47

7 MAC sublayer functional description 50
7.1 Aggregation Exchange Sequences and related rules 50
7.1.1 IAC/RAC Exchange Rules..... 50
7.1.2 Non-aggregate IAC/RAC Rate selection Rules 50
7.1.3 Response Period Scheduling 50
7.1.4 Responder Rules..... 51
7.1.5 Response PPDU Contents 51
7.1.6 Power-saving at the receiver of an aggregate PPDU 51
7.1.7 Protection mechanisms 52
7.1.8 MAC support for closed loop modes..... 56
7.1.9 Aggregated Single Directional frame transfer..... 57
7.1.10 Bi-directional data transfer..... 62
7.1.11 Multiple receiver aggregation (MRA) 65
7.2 MAC Header Compression 69
7.2.1 Introduction (informative)..... 69
7.2.2 Functional description..... 70
7.2.3 Ordering constraints for MHDR and CHDATA MPDUs 71
7.2.4 Multiple receiver aggregates 71
7.3 Multirate support 71
7.4 Quality of Service (QoS) 71
7.4.1 802.11e Requirements 71
7.4.2 Periodic RDR..... 72
7.4.3 Introduction (Informative)..... 72
7.4.4 Meaning of Service Period with Periodic RDR 72
7.5 Enhanced Block acknowledgment 72
7.5.1 Introduction (Informative)..... 72
7.5.2 Implicit BAR mechanism..... 73
7.5.3 Maintaining Acknowledgement State for a BA agreement..... 76
7.5.4 Compressed BA 76

8 MAC sublayer management..... 77
8.1 Coexistence management 77

8.1.1 Operating Modes 77
8.1.2 Receiver capabilities wrt legacy operation in the extension channel 77
8.1.3 Infrastructure BSS Operation 78
8.1.4 Ad-hoc (IBSS) Operation 84
8.2 Channel management 85
8.2.1 Channel Management at the AP 85
8.2.2 Channel Management at the STA 86
8.2.3 Channel Selection Methods 86
8.3 TRMS Power-Saving Management 87
8.3.1 BSS Association, DLP Setup and Mode Changes 87
8.3.2 TRMS Operation 88
8.3.3 Performance Impact of using TRMS mode (Informative) 89
8.3.4 Calibration 90

9 MAC Background (Informative) 92
9.1 Scheduling 92
9.1.1 Introduction 92
9.1.2 Scheduler 1 92
9.1.3 Scheduler 2 92
9.1.4 Retransmissions and Allocation of Time 93
9.1.5 Selecting RDL/RDG values (Informative) 93
9.1.6 Scheduling heuristics for MRA 94
9.2 Method to determine Best channel and channel width 94
9.2.1 First time Channel Selection 95
9.2.2 Subsequent Channel Selection 95
9.3 Background to Protection Mechanisms 95
9.3.1 Introduction to Protection Mechanisms 95
9.3.2 Effect of Spoofing 96
9.3.3 Pairwise spoofing 96
9.3.4 Selecting between Pairwise Spoofing and LongNAV 98
9.3.5 HT PPDU Format with Legacy Protection 98
9.3.6 Protection mechanisms for MRA 100

10 PHY-SAP service specification 103
10.1 PHY Specific Service Parameter List 103

10.1.1 TXVECTOR 103
10.1.2 RXVECTOR 106

11 MIMO-OFDM HT PHY specification 110
11.1 Overview 110
11.1.1 Introduction (Informative) 110
11.1.2 Timing related parameters 119
11.1.3 Mathematical conventions in signal descriptions 119
11.2 PLCP sublayer 120
11.2.1 Basic MIMO Mode 120
11.2.2 Basic Transmit Beamforming (Optional Mode) 147
11.2.3 Advanced Transmit Beamforming (Optional Mode) 154
11.2.4 Low-Density Parity-Check Coding (Optional Mode) 158

List of Figures

Figure 1 – IAC MPDU 22
Figure 2 – RAC MPDU 26
Figure 3 – MRAD MPDU 28
Figure 4 – ICB MPDU 29
Figure 5 – DCB MPDU 30
Figure 6 - Compressed BA MPDU 31
Figure 7 - BA Control field 31
Figure 8 – Calibration Training MPDU 32
Figure 9 – PPDU attributes for Calibration Training MPDUs 33
Figure 10 – MHDR MPDU 34
Figure 11 – CHDATA MPDU 34
Figure 12 – HT Capability element format 36
Figure 13 – HT Capabilities Info field 36
Figure 14 – TRMS element format 38
Figure 15 – Action codes for the HT action category 39
Figure 16 - Block Ack Parameter Set fixed field 41
Figure 17 – A-MPDU format 43
Figure 18 – A-MSDU Structure 48
Figure 19 – Example of TXOP Truncation 53
Figure 20 – Basic Operation of Pairwise Spoofing 54
Figure 21 – Basic Operation of Single Ended Spoofing 55
Figure 22 – LongNAV, untrained unidirectional exchange 58
Figure 23 – LongNAV, untrained unidirectional showing aggregate response 59
Figure 24 – LongNAV, trained unidirectional showing aggregate response 60
Figure 25 – Pairwise Spoofing, Untrained Unidirectional Exchange 61
Figure 26 – Pairwise Spoofing, Trained Unidirectional Exchange 62
Figure 27 – LongNAV Bi-Directional Example 64
Figure 28 – Pairwise Spoofing, Trained BiDirectional 65
Figure 29 – Example of MRA response scheduling with 3 responders 66
Figure 30 – MRMRA Example for Periodic Traffic 68
Figure 31 – MRMRA Example for Aperiodic Traffic 69
Figure 32 - Basic buffer management in Implicit BAR 74
Figure 33 - Buffer management when the first MPDU is incorrect in Implicit BAR (Case 1) 75

Figure 34 - Buffer management when the first MPDU is incorrect in Implicit BAR (Case 2) 75
Figure 35 – 20 MHz-Base Managed Mixed Mode 82
Figure 36 – Data transfer example with timed receive mode switching 88
Figure 37 - Calibration Procedure 91
Figure 38 – Standard Virtual Carrier Sense 97
Figure 39 – Pairwise Spoofing Mechanism 97
Figure 40 – PPDU Format with Legacy Protection 98
Figure 41 – When Received by a Legacy Node 98
Figure 42 – HT Signal Field BPSK Constellation 99
Figure 43 – When Received by an HT Node but MPDU is not Decodable 100
Figure 44 – When Received by an HT Node and the MPDU is Decodable 100
Figure 45: Transmitter Datapath for 2-antenna MIMO in 20MHz 112
Figure 46: Transmitter datapath for 2-antenna MIMO in 40MHz 112
Figure 47: PPDU Format for 2x20 Mandatory Basic MIMO Transmission 112
Figure 48: Upper and Lower sub-channel specification in 40MHz 113
Figure 49: PPDU Format for 2x40 Mandatory Basic MIMO Transmission 113
Figure 50: HT Preamble format 114
Figure 51: Transmitter configuration for 6Mbps in 40MHz 116
Figure 52: PPDU format for 6Mbps duplicate mode 116
Figure 53: Transmitter Datapath with option to perform spatial shaping 117
Figure 54: Timing boundaries for subframes in a HT transmission 120
Figure 55: PPDU Format for N_{TX} antennas 121
Figure 56: PPDU format leveraging CDD transmission 122
Figure 57: HT-LTF pattern for 2 antennas (2 spatial streams) 128
Figure 58: HT-LTF pattern for 3 antennas (3 spatial streams) 129
Figure 59: HT-LTF pattern for 4 antennas (4 spatial streams) 130
Figure 60: Constellation specification for LSIG and HTSIG 131
Figure 61: SIGNAL field bit assignment 131
Figure 62: HT SIGNAL FIELD (HTSIG1 and HTSIG2) bit assignment 134
Figure 63: Generation of CRC bits 137
Figure 64: Puncturing pattern for code rate = 7/8 139
Figure 65: Structure of interleaver 140
Figure 66: Tone Format for 40MHz channelization 143
Figure 67 Inputs and Outputs of IFFT for a 40 MHz channel 143

Figure 68: Transmit Spectrum Mask for 20MHz channelization 146
Figure 69: Transmit spectrum mask for 40MHz channelization 147
Figure 1: Application of calibration correction vector on Tx 152
Figure 70: Transmitter architecture for TX Beamforming for N_{ss}=2,and N_{TX} = 3 154
Figure 71: Transmit Datapath Architecture for ABF-MIMO mode 157

List of Tables

Table 1 - Interpretation of the ADDBA Parameters 32
Table 2 - Interpretation of the BA Bitmap 32
Table 3 –HT Capabilities 35
Table 4 – Additional HT Information Elements 36
Table 5 - Interpretation of the Buffer Size parameter 41
Table 6 – MPDU delimiter Fields 43
Table 7 Single Receiver Aggregation MPDUs 44
Table 8 - Broadcast/Multicast Receiver Group MPDUs 45
Table 9 - No Response Receiver Group MPDUs 46
Table 10 - MRA Aggregate MPDUs 46
Table 11 – QoS Control Field 48
Table 12 – Address field contents for unicast MPDUs 49
Table 13 - Address field contents for broadcast/multicast MPDUs 49
Table 14 – AP Mode Definitions 78
Table 15 – IBSS Mode Definitions 85
Table 16 – Beacon Elements for HT Channel Management 85
Table 17: Table of parameters for the TX Vector 104
Table 18: Scrambler Initialization Sequence 106
Table 19: Table of parameters for the RX Vector 107
Table 20: Options for Throughput Enhancement 111
Table 21: Feedback Mechanisms 117
Table 22: MIMO Transmission options 118
Table 23: Timing Related Parameters 119
Table 24: List of HT Long Training Fields 126
Table 25: Content of the legacy SIGNAL field 132
Table 26: MCS field content for the basic mode 135
Table 27: Interleaver Parameters 140
Table 28: Amount of Frequency Rotation 141

Table 29: Extended MCS set for ABF-MIMO 156
Table 30: Constellation Mapping for 256-QAM in ABF-MIMO Mode 157

2 References

[1] IEEE Std 802.11®-1999 (Reaff 2003)
Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications

[2] IEEE P802.11e/D8.0, February 2004
(Draft Amendment to IEEE Std 802.11, 1999 Edition (Reaff 2003))
Medium Access Control (MAC) Enhancements for Quality of Service (QoS)

[3] Project Authorization Request for IEEE 802.11n.
11-02-798r7-HT-SG-Draft-PAR.doc

[4] IEEE 802.11-04/887, "TGnSync Proposal Summary"

[5] IEEE 802.11-04/888, "TGnSync Proposal"

[6] IEEE 802.11-04/890, "TGnSync Proposal FRCC Compliance"

[7] IEEE 802.11-04/891, "TGnSync Proposal PHY Results"

[8] IEEE 802.11-04/892, "TGnSync Proposal MAC Results"

[9] IEEE 802.11-04/893, "TGnSync Proposal MAC1 Simulation Results"

[10] IEEE 802.11-04/894, "TGnSync Proposal MAC2 Simulation Results"

[11] IEEE 802.11-04/895, "TGnSync Proposal MAC Simulation Methodology"

3 Definitions

Aggregate: A PSDU transported by the PHY with an aggregate attribute indicating that it contains multiple MPDUs.

Initiator: A STA that holds a TXOP and transmits the first frame in a frame exchange sequence including aggregate PPDUs and/or aggregate control frames (IAC/RAC/MRAD).

LongNAV: A MAC-layer protection mechanism that protects multiple PPDUs exchanged within a TXOP

Mixed Mode: A mode of operation of an HT BSS in which there are co-channel legacy devices present.

OFDM Format: Definition of subcarrier frequencies relative to the channel center frequency, and specification of data bearing and pilot subcarriers.

Originator: A QSTA that sends data using the Block ACK mechanism.

Pairwise Spoofing: A PHY layer protection mechanism described in 7.1.7.2.1.

Preamble: short and long: Portions of the 802.11 HT PPDU used for PHY synchronization and channel estimation.

Pure Mode: A mode of operation of a BSS in which there are no legacy devices.

Receiver: Any STA receiving the current PPDU.

Recipient: A QSTA that receives data using the Block ACK mechanism.

Responder: A STA that responds to an initiator in a frame exchange sequence including aggregate PPDUs and/or aggregate control frames (IAC/RAC/MRAD).

SIGNAL FIELD: Portion of the 802.11 HT PHY PPDU that contains “Rate” and “Duration”, part of the PLCP header.

Symbol: Typically refers to an OFDM symbol, which is a collection of sub-carriers (or tones)

4 Abbreviations and acronyms

Term	Description
ABF-MIMO	Advanced Beamforming MIMO (Transmit Beamforming with Spatial water filling techniques on Multiple Spatial Streams)
ACK	Acknowledgement
AGC	Automatic Gain Control
AP	Access Point
ATOC	Aggregate Table Of Contents
BF-MIMO	Beamforming MIMO (Basic Transmit Beamforming on Multiple Spatial Streams)
CAP	Controlled Access Phase
CDD	Cyclic Delay Diversity
CHDATA	Compressed Header Data MPDU
CRC	Cyclic Redundancy Check
CSI	Channel State Information
CSIT	Channel State Information at the Transmitter
CTS	Clear to Send
DLP	Direct Link Protocol
DCB	Decrease Channel Bandwidth MPDU
ERP	Extended Rate PHY
FCS	Frame Check Sequence
FPD	Following Packet Descriptor
HC	Hybrid Controller
HID	Header Identity
HT	High Throughput
HT-LTF	High Throughput Long Training Field
HT-SIG	High Throughput Signal Field
HT-STF	High Throughput Short Training Field
IAC	Initiator Aggregate Control
ICB	Increase Channel Bandwidth MPDU
IE	Information Element
LLC	Logical Link Control
L-LTF	Legacy Long Training Field

L-SIG	Legacy Signal Field
L-STF	Legacy Short Training Field
MAC	Medium Access Controller
MCS	Modulation Coding Scheme – includes modulation order (BPSK, QPSK, 16-QAM, 64-QAM, 256-QAM) and FEC code rate (1/2, 2/3, 3/4, 7/8)
MFB	MCS Feedback
MHDR	MAC Header MPDU
MIMO	Multiple input, multiple output (Both transmitter and receiver have multiple antennas for signaling over a point-to-point link)
MISO	Multiple input, single output (transmitter has multiple antennas but receiver has a single antenna)
MPDU	MAC protocol data unit
MRA	Multiple Receiver Aggregate – an aggregate that contains MPDUs for multiple receiver addresses.
MRAD	Multiple Receiver Aggregate Descriptor
MRC	Maximum Ratio Combining (a technique to maximize received SNR using antenna signal combining).
MRMRA	Multi-response Aggregate – an MRA with multiple responses from its multiple receivers
MRQ	MCS Request
MSDU	MAC service data unit
NAV	Network Allocation Vector (a MAC-level carrier-sense)
N_{SD}	Number of Data Subcarriers
N_{SS}	Number of Spatial Streams
N_{RX}	Number of Receive Antennas
N_{TX}	Number of Transmit Antennas
OFDM	Orthogonal Frequency Division Multiplexing
PHY	Physical Layer
PLCP	PHY layer convergence protocol
PPDU	PHY protocol data unit
PSDU	PHY service data unit
Pure Mode	A mode of operation of a BSS in which there are no legacy devices
QAP	QoS access point
QoS	Quality of service

QSTA	QoS station
RAC	Royal Automobile Club
RDG	Reverse Direction Grant
RDL	Reverse Direction Limit
RDR	Reverse Direction Request
RDTID	Reverse Direction Traffic Identifier
RPO	Response Period Offset
RR	Rate Recommendation
RSSI	Receive Signal Strength Indicator
RTS	Request to send
RX	Receiver
SAP	Service Access Point
SF	Signal Field
SIFS	Short inter-frame spacing
SIMO	Single input, multiple output (transmitter has a single antenna, receiver has multiple antennas)
SISO	Single input, single output (both transmitter and receiver have a single antenna)
SoP	Start of Packet (i.e. PPDU Packet)
SRA	Single Receiver Aggregate – an aggregate that contains MPDUs for a single receiver address.
STA	Station
SVD	Singular Value Decomposition: For an arbitrary matrix A = USV^H
TBD	To be determined
TC	Traffic Category
TGe	802.11 Task Group e - Quality of Service
TGn	802.11 Task Group n - Enhancements for Higher Throughput
TRMS	Timed Receive Mode Switching
TRQ	Training Request
TS	Traffic Stream
TSID	Traffic Stream Identifier
TSPEC	Traffic specification
TX	Transmitter
TXOP	Transmission opportunity

1
2

5 Introduction

5.1 Purpose

This document defines proposed Enhancements for Higher Throughput within the scope of the Project Authorization Request (PAR) [3] formulated for IEEE 802.11 Task Group n.

5.2 General description of MAC enhancements

The proposed MAC enhancements assume a baseline specification defined by IEEE 802.11 standard [1] and its later amendments, including the draft IEEE 802.11e amendment [2]. The purpose of the enhancements is to significantly improve throughput, providing a maximum throughput of at least 100Mbps, as measured at the MAC data service access point (SAP).

The features supported by this TGn Sync MAC proposal are as follows:

- A Frame Aggregation format that allows aggregation of multiple Data and Control MPDUs in one PPDU. Note, an aggregate can also include a single management frame, in addition to multiple data and control frames per unicast receiver address.

- New Control MPDUs and (aggregate) frame exchange rules that control the exchange of multiple MPDUs per PPDU between a single initiating STA and potentially multiple responding STAs.

- MAC header compression that allows more efficient medium usage when multiple Data MPDUs are aggregated in a single PPDU.

- A receiver assisted channel training protocol.

- A reverse direction data protocol that allows efficient channel access for single and multiple responders within a frame exchange sequence.

- Protection of the new frame exchange sequences may be provided through one of two mechanisms: a MAC-level mechanism based on NAV reservations, and a PHY-level mechanism based on spoofing the legacy PLCP rate/length information.

- QoS enhancements for TSPEC renegotiation and prioritized discarding of QoS MSDUs according to an application-defined loss priority labeling.

- Mechanisms to manage coexistence of legacy and HT devices.

- Channel management and channel selection methods.

- Power saving mechanism to reduce power consumption by only enabling a single receive chain during periods of inactivity.

5.3 General description of PHY enhancements

With the aim to increase the maximum PHY layer data rate, we have introduced two primary techniques that are mandatory in our proposal:

- Transmission of Multiple Spatial Streams via Multiple Transmit Antennas. Our recommendation is to make 2 spatial streams mandatory, with the option to scale to 4 spatial streams.

- Extended Bandwidth signaling. 20MHz channelization shall be mandatory worldwide. Our recommendation is to support 40MHz channelization as mandatory in regulatory domains that will permit wider bandwidth operation

We have also introduced several secondary techniques for increasing throughput. These include:

- A reduced guard interval of 400ns (i.e. half-GI), channel dispersion permitting (e.g. TGn channel model B). Support for half-GI is mandatory in 20MHz.

- Higher channel coding rate of 7/8. Support for 7/8 is mandatory in 20MHz.

- Support for 256-QAM is optional in our proposal.

Due to the susceptibility of a wireless link to multipath fading, especially in a MIMO channel, our proposal introduces two optional techniques for improving the robustness of the link between the transmitter and the receiver, with the aim to keep the client (station) configuration low-cost and low-power:

- Transmitter beamforming, with the understanding that Access Points will be able to support more than 2 antennas, while the client stations would be restricted to 2 antennas.

- Advanced channel coding. Our proposal includes LDPC as an option. LDPC codes exhibit a larger minimum free distance compared to the existing 802.11a convolutional code, and hence are able to better exploit the frequency and spatial diversity in the channel.

Our proposal support full backwards compatibility with existing 802.11a/b/g standards. Our preamble design support PHY layer interoperability with the OFDM preamble from 802.11a/g, while providing a robust mechanism for synchronization of the high throughput data.

6 Frame Formats

This clause specifies the format of new HT frames.

Reserved fields and subfields shall be transmitted as zeroes and shall be ignored by an HT receiver

6.1 Control frames

This section defines new control frame subtypes as listed below:

Control Frame **Subtype**	**Name**
IAC	Initiator aggregate control
RAC	Responder aggregate control
MRAD	Multiple receiver aggregate descriptor
ICB	Increase channel bandwidth
DCB	Decrease channel bandwidth
Calibration Training	Calibration Training frame

6.1.1 Initiator Aggregation Control (IAC) MPDU

6.1.1.1 Introduction

This MPDU provides control of the MCS, size, duration, training and any reverse flow of aggregates in an exchange of aggregates between STAs. It is sent by the initiator STA of an aggregate exchange.

There are a number of groups of fields (IAC elements), which perform specific purposes within the MPDU as described below:

Element Name	**Purpose of item**
RTS – Request to Send	Detects collisions at the receiver Tests state of NAV in the receiver Sets NAV reservation near Transmitter (LongNAV)
TRQ – MIMO training request	Request to train the channel Used for MIMO with implicit Feedback.
MRQ – MCS request	Request for MCS feedback
MFB – MCS Feedback	Closed loop channel adaptation feedback Contains values from which the MCS, which will be used for the next PPDU, may be derived deterministically. Note, when combined with an RDG, the RDG duration is adjusted as appropriate to this MCS value.
FPD – Following packet descriptor	Provides size and default MCS for the next PPDU to be sent by this STA so that its peer can properly set a Pairwise Spoofing value

RDL – Reverse direction limit	Indicates a duration that is a limit on the amount of reverse direction that will be granted Used for Bi-directional frame transfer as defined in section 7.1.10
RDG – Reverse direction grant	Grants a specified duration to be used for a reverse direction PPDU Used for Bi-directional frame transfer as defined in section 7.1.10

Note, although described as "elements", the structure of the IAC MPDU containing these elements is fixed.

6.1.1.2 Format of IAC MPDU

This section defines the format of the IAC MPDU.
The IAC MPDU has a fixed structure. Fields that are not used are present, but reserved.

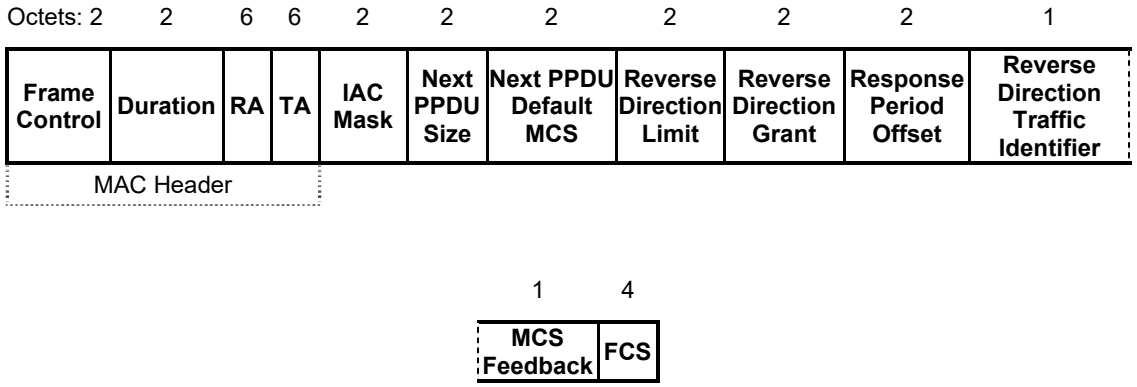


Figure 1 – IAC MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Type is control, subtype is IAC
Duration	2	
Receiver Address	6	
Transmitter Address	6	
IAC Mask	2	A bitmask indicating which IAC elements are present in the packet.
Next PPDU Size	2	Size in bytes of following PPDU that will be sent by this STA (FPD). This is interpreted along with the Next PPDU default MCS to determine the duration of the next

		PPDU. Present when FPD is indicated, otherwise undefined.
Next PPDU Default MCS	2	Default MCS that will be used in the absence of any updated training information to send next PPDU. Present when FPD is indicated, otherwise undefined.
Reverse direction limit	2	Indicates the amount of time in microseconds that is the maximum amount of time that will be granted in a RDG. Present when RDL is indicated the IAC mask field, otherwise undefined.
Reverse direction grant	2	Indicates the amount of time in microseconds that is available for a reverse direction PPDU including any expected response MPDUs and an RAC MPDU. Present when RDG is indicated, otherwise undefined.
Response Period Offset (RPO)	2	Indicates the delay in microseconds between the end of the PPDU containing the IAC MPDU and the start of the response PPDU. This value shall be no less than SIFS. Present when RDG is indicated, otherwise undefined.
Reverse Direction Traffic Identifier (RDTID)	1	Indicates one of: • An AC for which reverse direction grant is valid • A TSID indicating a specific TS for which the reverse direction grant is valid • Unconstrained Present when RDG is indicated, otherwise undefined.
MCS Feedback	1	Contains a recommended MCS value. Present when MFB is indicated, otherwise undefined.
FCS	4	

1

6.1.1.2.1 IAC Mask Field

This field indicates which logical elements are carried in the MPDU.

IAC Mask Field bit	Position	Description
RTS	B0	When set, indicates this is an RTS. The receiver should not generate any response if its NAV is non-zero.
TRQ	B1	When set, indicates a request to train the channel. Used for MIMO with implicit Feedback.
MRQ	B2	When set, indicates a request for MCS feedback.
MFB	B3	When set, indicates an MFB training response is present as defined by the MCS Feedback field.
FPD	B4	When set, indicates that the next PPDU duration may be determined from the next PPDU length and default MCS fields. FPD is set only when using Pairwise Spoofing rules.
RDG	B5	When set, indicates a reverse direction grant is present.
RDL	B6	When set, indicates that a reverse direction limit is present

6.1.1.2.2 MCS Feedback Field

This contains an MCS value containing a recommended MCS value.
MCS values are defined in section 11.2.1.4.2.2.

6.1.1.3 Use of IAC elements according to context

This section defines what IAC elements are valid in the following three contexts:

- Within an SRA
- Within an MRA, when a response is permitted for that receiver
- Within an MRA, when no response is permitted for that receiver

Element Name	SRA Context	MRA Context (response)	MRA Context (non-response)
RTS – Request	Valid	Valid (see note	Not Valid

to Send		below)	
TRQ – MIMO training request	Valid	Valid	Not Valid
MRQ – MCS request	Valid	Valid	Not Valid
MFB – Use this MCS selection	Valid	Valid	Valid
FPD – Following packet descriptor	Valid – implies Pairwise Spoofing rules	Not Valid	Not Valid
RDL – Reverse direction limit	Valid	Valid	Not Valid
RDG – Reverse direction grant	Valid	Valid	Not Valid

Note – the RTS element may be included in an MRA. For this to be useful, the initiator should aggregate together an IAC (RTS) for each responder and send these as an aggregate. These IACs should contain a short RDG to collect RAC (CTS) responses. The responder will interpret this as a request to check its NAV is idle during its scheduled response period. If the NAV is idle, a response is sent. Using this method, the initiator can test the NAV state of multiple STAs.

6.1.2 Responder Aggregation Control (RAC) MPDU

6.1.2.1 Introduction

This MPDU provides control of the MCS, size, duration, training and any reverse flow of aggregates in an exchange of aggregates between STAs. It is sent by the responder STA of an aggregate exchange.
There are a number of groups of fields (RAC elements), which perform specific purposes within the MPDU as described below:

RAC Element Name	Purpose of item
CTS – Clear-to-Send	Response to RTS indicates that the medium is clear
TRQ – MIMO training request	Request to train the channel Used for MIMO with implicit Feedback.
MRQ – MCS request	Indicates a request for MCS feedback
MFB –MCS Feedback	Closed loop channel adaptation feedback Contains values from which the MCS, which will be used for the next PPDU, may be derived deterministically. Note, when combined with an RDG, the RDG duration is adjusted appropriate to this MCS value.
RDR – Reverse direction request	Request for a reverse direction data flow Contains requested total data MPDU length (bytes) and default MCS

6.1.2.2 Format of RAC MPDU

This section defines the format of the RAC MPDU.
The RAC MPDU has a fixed structure. Fields that are not used are present, but reserved.

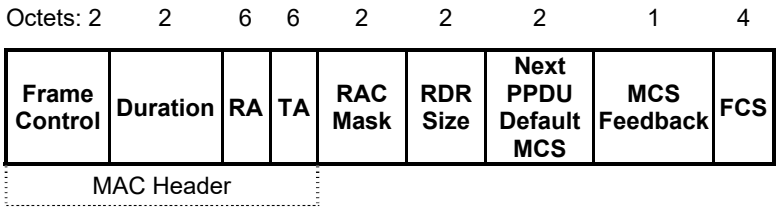


Figure 2 – RAC MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Type is control, subtype is RAC
Duration	2	
Receiver Address	6	
Transmitter Address	6	
RAC Mask	2	A bitmask indicating which RAC elements are present in the packet.
RDR Size	2	Size in bytes of requested reverse direction flow. This is interpreted along with the Next PPDU default MCS to determine the duration of the next PPDU. Present when RDR is indicated, otherwise undefined.
Next PPDU Default MCS	2	Default MCS that will be used in the absence of any updated training information to send next PPDU Present when RDR is indicated, otherwise undefined.
MCS Feedback	1	Contains a recommended RCS value.. Present when MFB is indicated, otherwise undefined.
FCS	4	

6.1.2.2.1 RAC Mask Field

This field indicates which logical RAC elements are carried in the MPDU.

RAC Mask Field bit	Position	Description
CTS	B0	When set, indicates this is a CTS.
TRQ	B1	When set, indicates a request to train the channel. Used for MIMO with implicit Feedback.
MRQ	B2	When set, indicates a request for MCS feedback.
MFB	B3	When set, indicates an MFB training response is present as defined by the MCS Feedback field.
RDR	B4	When set, indicates a request for reverse direction data-flow is present as described by the RDR Size

		and Next PPDU Default MCS fields.

6.1.2.2.2 MCS Feedback Field

As defined in IAC MPDU.

6.1.3 Multiple Receiver Aggregate Descriptor (MRAD) MPDU

6.1.3.1 Introduction

The MRAD MPDU is the first frame inside a MRA aggregate. It defines the contents of the aggregate. The MRAD may be used by a receiving STA to save power. Based on the contents of the MRAD, the receiving STA may disable its receiver after it has received any broadcast frames or frames addressed to its MAC address.

6.1.3.2 Format of MRAD MPDU

This section defines the format of the MRAD MPDU.

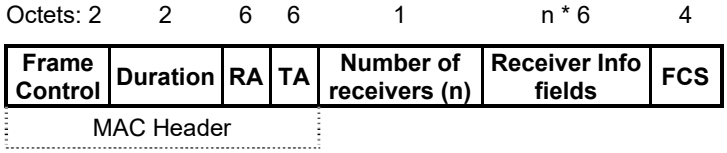


Figure 3 – MRAD MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Frame type is Control, subtype MRAD
Duration	2	
Receiver Address	6	The RA field is the broadcast group address.
Transmitter Address	6	The TA field is the address of the STA transmitting the MRA aggregate.
Number of receivers	1	The number of receivers (n) for which MPDUs are included inside the MRA aggregate.
Receiver Info fields	n * 6	n Receiver Info fields, one for each Receiver Address in the MRA aggregate.
FCS	4	

The order of the Receiver Info fields is equal to the order in which MPDUs for the corresponding receivers appear in the MRA aggregate.

6.1.3.2.1 Receiver Info field

The format of the Receiver Info field is as described below.

Field	Size (bytes)	Description
Receiver Address	6	Contains the receiver's MAC address

6.1.4 Increase Channel Bandwidth (ICB) MPDU

6.1.4.1 Introduction

The ICB frame is sent by an HT AP to start a 40 MHz period in 20 MHz-base managed mixed mode. When 40 MHz HT STAs operating in 20 MHz receive this frame, they shift to 40 MHz analogue mode.

6.1.4.2 Format of ICB MPDU

The frame format for the ICB frame is defined as is in Figure 4.

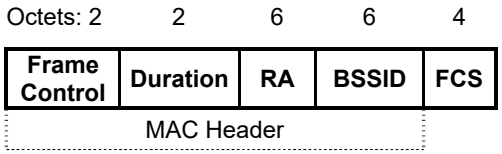


Figure 4 – ICB MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Frame type is Control, subtype ICB
Duration	2	Duration of the 40MHz phase plus transition periods
Receiver Address	6	The RA field is the broadcast group address.
BSSID	6	
FCS	4	

This frame is used in the 20 MHz-base managed mixed mode and sent by the HT AP. When switching to 40 MHz channel, the Duration field will be set to cover the 40 MHz phase plus the transition periods between 20 and 40 MHz operation. This is to have HT STAs switch back to 20 MHz channel even if they do not receive the DCB frame at the end of the 40 MHz period.

6.1.5 Decrease Channel Bandwidth (DCB) MPDU

6.1.5.1 Introduction

The DCB frame is sent by an HT AP to end a 40 MHz period in 20 MHz-base managed mixed mode. When 40 MHz HT STAs operating in 40 MHz receive this frame, they shift to 20 MHz analogue mode.

6.1.5.2 Format of DCB MPDU

The frame format for the DCB frame is defined as is in Figure 4.

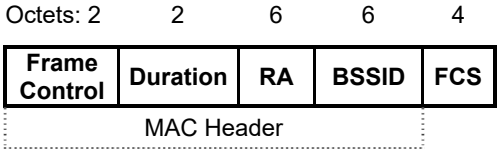


Figure 5 – DCB MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Frame type is Control, subtype DCB
Duration	2	Duration of the 20MHz period, in TU.
Receiver Address	6	The RA field is the broadcast group address.
BSSID	6	The BSSID field is the address of the STA contained in the AP.
FCS	4	

This frame is used in the 20 MHz-base managed mixed mode and sent by the HT AP.
When switching back to the 20 MHz channel, the Duration field will at least cover up to the start of the next 40 MHz period. The Duration field in the DCB frame contains this duration in TUs.

6.1.6 Block Ack (BA) MPDU

This section defines modifications to the BA frame introduced in 802.11e.

6.1.6.1 Introduction

This section defines a compressed variant of the 802.11e BA MPDU. The variant is distinguished by setting the "Compressed" field of the BA Control field to 1.
Two compression mechanisms exist:

- Support for non-fragmented BA. This reduces the bitmap size to 1 bit per MSDU.
- Truncation of the bitmap to reduce the number of MSDUs acknowledged in the bitmap.

The use of these mechanisms is negotiated when the BA agreement is created through the exchange of ADDBA frames, and is constant for a BA agreement throughout its lifetime. This means that a BA for a particular BA agreement is always the same size.

6.1.6.2 Format of BA MPDU

The frame format for the BA frame is defined as in Figure 6.

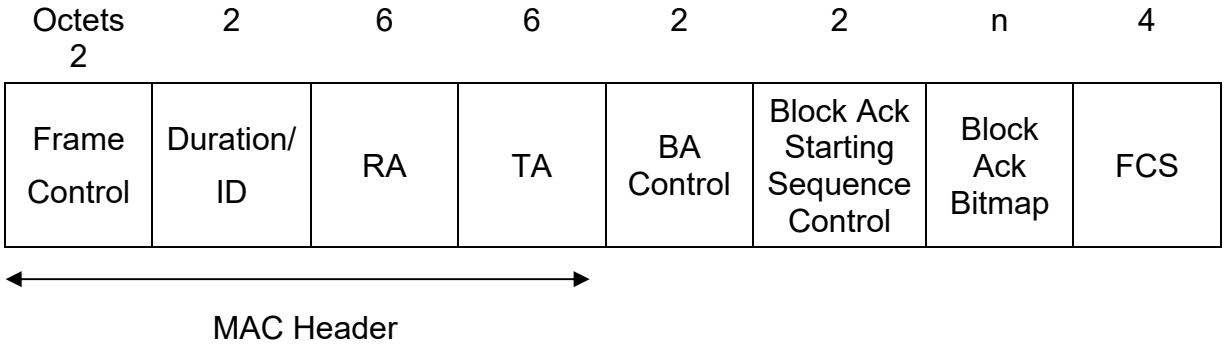


Figure 6 - Compressed BA MPDU

Only the BA control and Block Ack Bitmap fields are modified.

The BA Control Field is shown in Figure 7.

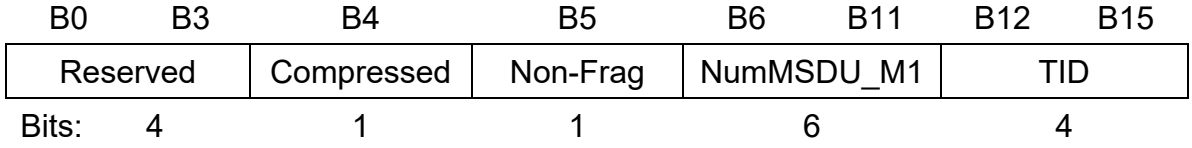


Figure 7 - BA Control field

The Compressed, Non-Frag and NumMSDU_M1 are new fields.

The Compressed field selects between legacy 802.11e format and compressed format. When Compressed is set to 0, the BA is interpreted according to legacy 802.11e. The Compressed field is always set to 1 when the BA is set up between HT STAs.

The Non-Frag bit indicates whether acknowledgement is per MSDU or per MPDU. It affects the interpretation of the bitmap field as defined below.

The NumMSDU_M1 field defines the number of MSDUs that are acknowledged by the present bitmap, less one. The NumMSDU_M1 field set as 0 indicates 1 MSDU is acknowledged by the bitmap and when it is 63, 64 MSDUs are acknowledged.

The maximum value of NumMSDU_M1 and the length of the Block Ack Bitmap is defined by the BA agreement, and is fixed for the lifetime of the BA agreement. They are derived from the ADDBA parameters, Non-Frag and Buffer Size, as defined below

Table 1 - Interpretation of the ADDBA Parameters

Non-Frag subfield value	Maximum value of NumMSDU_M1	Bitmap length n (octets)
1	Buffer Size – 1	ceil (Buffer Size/8)
0	Min (Buffer Size, 64) – 1	Buffer Size x 2

The bitmap is interpreted according to the setting of the Compressed, Non-Frag and NumMSDU_M1 fields as defined in the following table.

Table 2 - Interpretation of the BA Bitmap

Compressed	Non-Frag	NumMSDU_M1	Bitmap interpretation
0	0	0	Legacy 802.11e interpretation
1	1	0 to 63	Each bit set acknowledges the successful reception of a single MSDU in order of sequence number
1	0	0 to 63	Each bit set acknowledges reception of a single MPDU. The bitmap is interpreted as in legacy 802.11e except that only the first (NumMSDU_M1+1) x 16 bits are significant.

Other combinations of these fields are reserved.

6.1.7 Calibration Training MPDU

The frame format of the Calibration Training MPDU is defined in Figure 8.

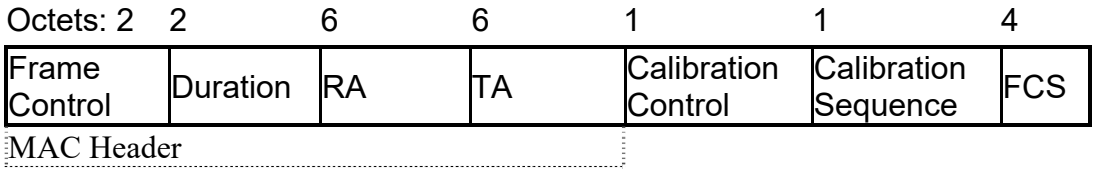


Figure 8 – Calibration Training MPDU

The Calibration Control field contains:

- Number of antennas at the initiating STA *Ni* (2 bits)
- Number of antennas at the responding STA *Nr* (2 bits)
- Position (2 bits). This field is set to '0' by the STA transmitting the first frame that initiates the calibration training frame exchange sequence. The field is incremented in

subsequent Calibration Training frame transmissions in the calibration training frame exchange sequence.

- *Reserved* (2 bits)

The Calibration Sequence field is incremented at the transmission of the Calibration Training frame that initiates the calibration procedure. The remaining frames in the same calibration procedure use the same Calibration Sequence value.

The Calibration Traning frame is carried in a PPDU of attributes defined in according to the value of the position field.

Position Field	PPDU	TXVECTOR NLTF
0	Non-Aggregate	
1	Non-Aggregate, Sounding	*Ni*
2	Non-Aggregate, Sounding	*Nr*

Figure 9 – PPDU attributes for Calibration Training MPDUs

The Calibration Training frame shall be transmitted in a non-aggregate PPDU as part of the calibration training exchange defined in 8.3.4.sounding PPDU that includes an HT-LTF of appropriate length.

6.2 Data frames

To allow for MAC header compression two new MPDUs are defined: MAC Header (MHDR) MPDU and Compressed Header Data (CHDATA) MPDU.
These MPDUs are present only in aggregates. Because the space of Data Subtypes is fully used, the unused MPDU type with value "11" is defined to be the Extended Data (EDATA) type. The MAC header (MHDR) and Compressed Header Data (CHDATA) MPDUs defined here are subtypes of EDATA.

6.2.1 MAC Header (MHDR) MPDU

6.2.1.1 Introduction

This MPDU provides the MAC header for subsequent CHDATA MPDUs.

6.2.1.2 Format of MHDR MPDU

This section defines the format of the MHDR MPDU.

Octets: 2	2	6	6	6	1	1	2	4
Frame Control	Duration	Address 1	Address 2	Address 3	HID	Reserved	QoS Control	FCS



Figure 10 – MHDR MPDU

Note, the MHDR MPDU contains no data.

Field	Size (bytes)	Purpose
Frame Control	2	Frame type is Data.
Duration	2	
Address 1	6	
Address 2	6	
Address 3	6	
HID	1	The Header ID field that links the MHDR MPDU with corresponding CHDATA MPDUs
Reserved	1	
QoS Control	2	
FCS	4	

6.2.2 Compressed Header Data (CHDATA) MPDU

6.2.2.1 Introduction

This MPDU contains the actual data and a (compressed) header that refers to a MHDR MPDU for full MAC header information.

6.2.2.2 Format of CHDATA MPDU

This section defines the format of the CHDATA MPDU.

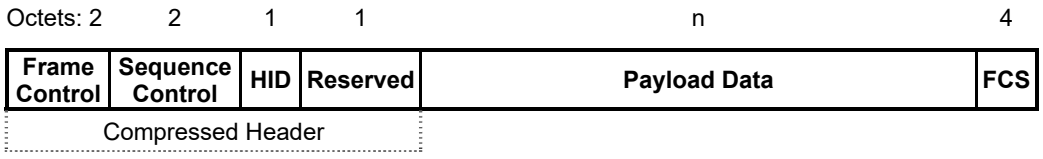


Figure 11 – CHDATA MPDU

Field	Size (bytes)	Purpose
Frame Control	2	Frame type is EDATA, subtype is CHDATA

Sequence Control	2	
HID	1	The Header ID field that links the CHDATA MPDU with corresponding MHDR MPDU
Reserved	1	
FCS	4	

6.3 Management frame formats

6.3.1 HT Capabilities

Support for any optional features or behavior will be declared through a STA's capabilities or management elements.

Table 3 –HT Capabilities

Capability	Contents
HT Supported	Indicates that the device is an HT device, supporting all the mandatory features of this document. This capability is indicated by the presence of the HT capability element.
Number of Rx Antennas	Indicates the number of Rx Antennas
Max number of Rx spatial channels	Max number of spatial channels the device can receive: in the range 1 to 4.
Supported MCS set	Defines the MCS values supported by a STA. An MCS set contains a bitmap of size 128 bits. Bit 0 maps to MCS 0. A bit is set to indicate support for that MCS. See section **11.2.1.4.2.2** for a list of MCSs.
Advanced Coding Capability	Advanced Coding Capability Indicates support for advanced coding. This field has the same format as the advanced coding field defined in section **10.1.1.8**, where a bit set as one indicates that that coding type is supported.
Supported Channel Width Set	Contains: {20} for a device only capable of 20MHz operation {20, 40} for a device capable of both 20MHz and 40MHz operation.

6.3.1.1 HT Capability element

The HT Capability element contains a number of subfields that are used to advertise optional HT capabilities at an HT device. The HT Capability element is present in beacons, association and re-association request frames and in probe response frames. The presence of the HT Capability element indicates that the device is an HT device. The HT Capability element is defined in Figure 12.

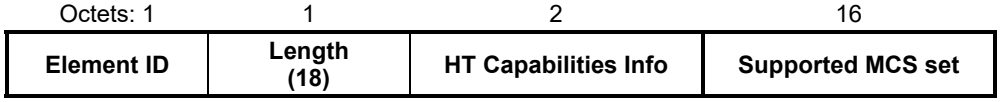


Figure 12 – HT Capability element format

The HT Capabilities Info field is 2 octets in length, and contains capability information bits as defined in Figure 13.

B0 B1	B2 B3	B4	B5
Number of rx Antennas	**Max number of Rx spatial channels**	**Advanced coding capability**	**Supported channel width set**

Figure 13 – HT Capabilities Info field

The Number of rx antennas field contains the number of rx antennas supported, minus 1.
The max number of rx spatial channels filed contains the maximum number of rx spatial channels supported, minus 1.
The advanced coding capability field indicates support for advanced coding if its value is 1.
The supported channel width set indicates {20} if its value is 0, and {20, 40} if its value is 1.
Note: All APs and STAs support TRMS operation in other STAs with a zero hold-on time.

6.3.2 Additional HT Information Elements

The AP will signal information in new information elements to manage the BSS.
This information is defined in the table below.

Table 4 – Additional HT Information Elements

Name	Description	Use
Control Channel	Channel number of the control 20MHz channel	In a mixed mode, all broadcast and non-aggregated control frames are sent so that they can be received by a legacy device operating on the control channel.
Extension Channel Offset	Values: –1, 1 indicate offset of the extension channel from the control channel. Value of 0 indicates no extension channel is present.	To locate the 40MHz channel in combination with the control channel
Permitted Width	Defines the channel widths	A STA can use any channel

Set	that may be used in the BSS. Values: {20} or {20,40}	width in transmitting a unicast non-control frame supported by the destination STA, and present in the permitted width set.
Operating Mode	Either: Mixed: (there are legacy devices co-channel) Pure: (there are no legacy devices co-channel or the AP is operating internally in managed mixed mode) Base-20 Managed mixed: (there may be legacy devices in both the control and the extension channel, see section 8.1.3.4) Mixed-Capable IBSS	Used to control the use of protection mechanisms (see section7.1.7). Also selects between legacy beacon and HT beacon generation in an IBSS network (see section 8.1.3.2).
Channel Extension Indication	Values: 0,1 1 means that 40 MHz capable HT STAs in the BSS should extend their channel bandwidth to 40 MHz according to the information from the Extension Channel Offset information element.	Present in Beacon frames in CFP when 20 MHz-base managed mixed mode is used. It indicates that transition to a 40 MHz period will start in the current CFP after the reception of the Beacon.
Extension Channel Access Timeout	Integer >=1 Specifies a time limit (in TU) after which the AP gives up the transition to 40 MHz and the 40 MHz capable HT STA switches back to 20 MHz if it has not received CTS-to-self or Beacon on the extension channel.	Present in Beacon/Probe Response frames when 20 MHz-base managed mixed mode is used. Note: The AP sets the local extension channel access timer to the Extension Channel Access Timeout minus the duration of CTS-to-self or Beacon frame.
Basic MCS set	An MCS set contains a bitmap of size 128 bits. Bit 0 maps to MCS 0. A bit is set to indicate support for that MCS.	Present in Beacon/Probe Response frames to indicate what MCS values shall be supported by all devices in the BSS.
Duplicate Legacy	Indicates that this beacon is a duplicate transmitted as	Present in legacy beacons transmitted by a HT AP during

Beacon	part of managed mixed mode operation. HT STAs should ignore duplicate legacy beacons regardless of BSSID.	managed mixed mode operation. (See section 8.1.3.3.)

6.3.2.1 TRMS Information Element

The TRMS feature is managed by the TRMS information element defined below.
The TRMS Information Element (IE) may be included in association and reassociation request frames, DLP request and response frames and TRMS Update action frames. This IE includes the following fields:

Field	Description	Use
TRMS active	Indicates whether or not TRMS operation is active	If TRMS is active, a STA transmitting to this STA shall operate the procedures defined in section 8.3.
Hold-on time	The amount of time in 100s of milliseconds that the STA operates with full receive capability after transmitting a frame. A value of zero indicates that it retains full receive capability for 1 slot time (+ SIFS) after transmitting a frame for which a response is expected.	A hold-on time of zero indicates that the STA wishes to reduce receive capability on completing a sequence exchange. All APs and STAs support TRMS operation in other STAs with a hold-on time of zero. It is an option of a STA to support TRMS tracking for non-zero values of the hold-in time.

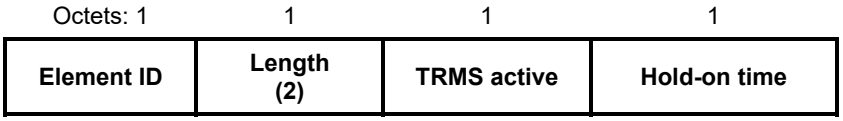


Figure 14 – TRMS element format

The TRMS active field indicates that TRMS is active when it takes the value 1, and TRMS is inactive when it takes the value 0.

6.3.3 Management Action Frames

These frame are defined to be action frames using the 802.11e Action frame format.
They share a common category code "HT".

The frames using the "HT" category code are defined in Figure 15.

Action Code	Meaning
1	TRMS Update
2	Calibration Measurement
3	Calibration Correction

Figure 15 – Action codes for the HT action category

6.3.3.1 TRMS Update Action Frame

The TRMS Update Action frame provides a mechanism for the STA to notify the AP and its peers of changes in its operating mode. This frame includes the TRMS IE defined in section 6.3.2.1.

The frame body of the TRMS update contains the information shown in the table below:

Order	Information
1	Category
2	Action
3	Dialog Token
4	Status Code
5	TRMS IE

All fields except the TRMS IE are defined in 802.11e **[2]**.

6.3.3.2 Calibration Measurement Action Frame

Calibration Measurement frame is a management action frame of category Calibration. It contains the following fields.

Order	Information	Size
1	Calibration Control	1
2	Calibration Sequence	1
3	MCMR Segment Sequence	1
4	MIMO Channel Measurement Report	3 × *Ns* × *Ni* × *Nr*

Where:

- *Ns* = Number of subcarriers in the Calibration Training symbols
- *Ni* = Number of antennas at the STA that transmitted the Calibration Request
- *Nr* = Number of antennas at the STA that is transmitting the Calibration Response

All messages in the same calibration procedure use the same Calibration Sequence value.
The Calibration Control field (1 octet) consists of:

- Number of antennas at the initiating STA *Ni* (2 bits)
- Number of antennas at the responding STA *Nr* (2 bits)
- Calibration Complete (1 bit). The Calibration Complete bit is set to '1' if the STA transmitting the Calibration Measurement frame declares the calibration procedure complete without requiring the other STA to transmit a Calibration Correction Vector frame.
- Reserved (3 bits)

The MCMR Segment Sequence (1 octet) contains the segment sequence number when the MIMO Channel Measurement Report has been segmented. The MCMR Segment Sequence of the first segment is 0. The complete MIMO Channel Measurement Report contains $Ns \times Ni \times Nr$ complex coefficients.

Each complex coefficient (3 octets) is represented by its 12-bit real and 12-bit imaginary parts. Hence the total length of the complete MIMO Channel Measurement Report is $3 \times Ns \times Ni \times Nr$ octets.

The complete MIMO Channel Measurement Report may need to be transmitted using more than one MAC frame. In the mandatory configuration with $Ns = 114$, $Ni = 2$, $Nr = 2$, there will be 1368 octets (456 coefficients) in a single MIMO Channel Measurement Report element. The maximum length of the MIMO Channel Measurement Report field is 5472 octets (1824 coefficients) for the case with $Ns = 114$, $Ni = 4$, and $Nr = 4$. If the MIMO Channel Measurement Report is longer than 1368 octets, the first 1368 octets are included in the first frame (MCMR Segment Sequence = 0) and the remaining octets are included in subsequent frames (MCMR Segment Sequence = 1, 2, etc.).

6.3.3.3 Calibration Correction Action Frame

Calibration Correction frame is a management action frame of category Calibration Correction. It contains the following fields.

Order	Information	Size
1	Calibration Control	1
2	Calibration Sequence	1
3	MIMO Correction Vector	$3 \times Ns \times Nr$
Where: • *Ns* = Number of subcarriers in the Calibration Training symbols • *Nr* = Number of antennas at the STA that transmitted the Calibration Response		

All frames in the same calibration procedure use the same Calibration Sequence value.

The Calibration Control field (1 octet) consists of:

- Number of antennas at the initiating STA *Ni* (2 bits)
- Number of antennas at the responding STA *Nr* (2 bits)
- Reserved (4 bits)

The MIMO Correction Vector contains *Ns* × *Nr* complex coefficients.
Each complex coefficient (3 octets) is represented by its 12-bit real and 12-bit imaginary parts.
Hence the total length of the MIMO Correction Vector field is 3 × *Ns* × *Nr* octets. The maximum length of the Correction Vector field corresponds to *Ns* = 114, *Nr* = 4 is 1368 octets.

6.3.4 TSPEC Modification to support Periodic RDR

This section defines modifications to the 802.11e TSPEC structure in support of Periodic reverse direction data.
A new 1-bit field is added to the TSPEC Info Field called "Periodic RDR".
This field takes values 0 and 1.
When "Periodic RDR" is 1, the TSID is associated with a Periodic RDR. This is valid only with TSPECs indicating periodic uplink. This is interpreted by the HC to mean that, instead of transmitting a QoS CF-Poll to the STA to satisfy its TSPEC, the HC periodically transmits an IAC containing an RDG for the specified TSID to satisfy the TSPEC.
See also section 7.4.2.

6.3.5 Block Ack Parameter Set field

This section defines changes to the 802.11e Block Ack Parameter set field of the ADDBA frame.

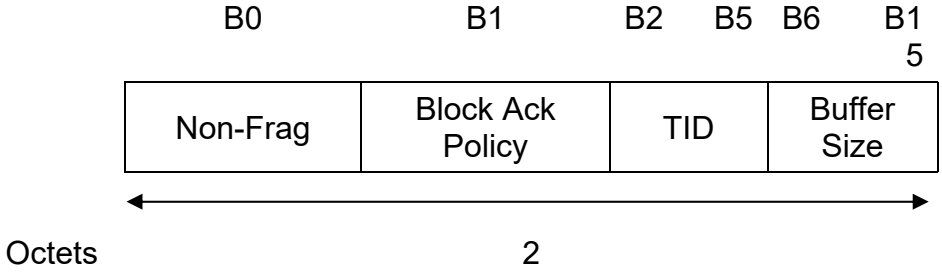


Figure 16 - Block Ack Parameter Set fixed field

The Non-Frag subfield is set to 1 if fragmentation is not used.
The compressed BA format is used whenever the Block Ack negotiation is set up between the HT STAs.
The Buffer Size subfield contains the number of buffers at the receiver. Its interpretation depends on the setting of the Non-Frag subfield. The maximum value of the NumMSDU_M1 parameter used in the compressed BA is also derived from the Buffer Size and Non-Frag fields as specified in the following table.

Table 5 - Interpretation of the Buffer Size parameter

Non-Frag subfield	Buffer Size interpretation	NumMSDU_M1

value		
1	Buffer size contains the maximum number of MSDUs that can be sent within the BA window. Maximum value is 64	NumMSDU_M1 = Buffer Size – 1
0	Buffer size contains the maximum number of MPDUs that can be sent within the BA window. Maximum value is 1024 (i.e., 64 x 16)	NumMSDU_M1 = min (Buffer Size, 64) - 1

In an ADDBA request frame, the Buffer Size subfield is advisory and the recipient selects a value equal to or smaller than the initial value.

6.4 A-MPDU and A-MSDU aggregation formats

6.4.1 Introduction to Aggregation (Informative)

There are two types of aggregation defined in this specification: A-MPDU aggregation and A-MSDU aggregation.

A-MPDU aggregation joins multiple MPDUs together and transports them in a single PSDU. A-MPDU aggregation is defined in this section. A-MPDU aggregation can be considered to be a layer that operates at the bottom of the MAC.

A-MSDU aggregation joins together multiple MSDUs and creates a single larger MSDU that is transported in MPDUs. A-MSDU aggregation is defined in 6.4.5. A-MSDU aggregation can be considered to be a layer that operates at the top of the MAC.

Note, the term "aggregate" used elsewhere in this document refers to A-MPDU aggregation, unless specifically referring to A-MSDU aggregation.

The MAC interprets the PSDU as a single MPDU, or as an A-MPDU according to the *Aggregate* TXVECTOR/RXVECTOR parameter, that is robustly transported in the PLCP header.
The MAC separates the MPDUs of an A-MPDU using a robust MPDU delimiter. It is robust in the sense that loss of a single delimiter does not necessarily cause loss of multiple MPDUs in an aggregate.
Note, MPDUs are separately protected by a CRC. Loss of an individual MPDU does not imply loss of all MPDUs in a PPDU.
The Duration fields in the MAC Headers of all MPDUs in an aggregate carry the same value.

6.4.2 Aggregated MPDU format (A-MPDU)

An A-MPDU consists of a number of MPDU delimiters each followed by an MPDU.

Except when it is the last MPDU, padding octets are appended (if needed) to make each section a multiple of 4 octets in length.

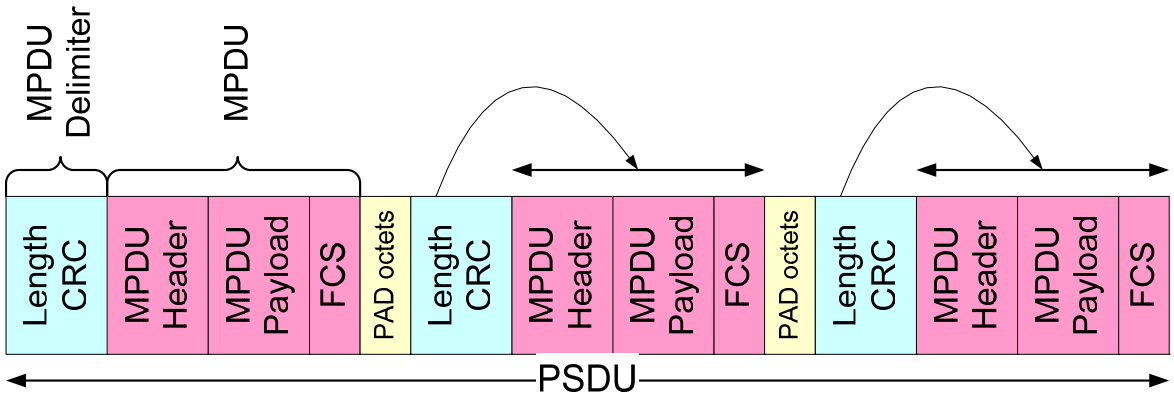


Figure 17 – A-MPDU format

The MPDU delimiter is 4 octets in length and contains the fields defined below:

Table 6 – MPDU delimiter Fields

MPDU delimiter Field	Size (bits)	Description
Reserved	4	-
MPDU length	12	Length of the MPDU in octets
CRC	8	8-bit CRC of the preceding 16-bits. As defined in IEEE 802.11 section 14.3.2.2.3
Unique pattern	8	Unique pattern that may be used to detect an MPDU delimiter when scanning for a delimiter. The unique pattern is set to the ASCII value for character 'N'.

The purpose of the MPDU delimiter is to robustly delimit the MPDUs within the aggregate. Robust in this case means that the structure of the aggregate can usually be recovered when one or more MPDU delimiters are received with errors. Individual delimiters have the same BER as the surrounding MPDUs, and so can be lost.

Note, a delimiter with MPDU length zero is valid. This can be used to introduce padding between MPDUs, or at the end of the A-MPDU.

6.4.2.1 Deaggregation (Informative)

The receiver checks the MPDU delimiter for validity based on the CRC. It may also check that the length indicated is within the PSDU HT-LENGTH indicated in RXVECTOR.
If the MPDU delimiter is valid, the MPDU is extracted from the aggregate. The next MPDU delimiter is expected at the first multiple of 4 octets immediately after the current MPDU. This process is continued until the end of the PPDU is reached.
If the MPDU delimiter is not valid, the deaggregation skips forward 4 octets and checks to see if the new location contains a valid MPDU delimiter. It continues searching until either a valid

delimiter is found, or the end of the PSDU is reached based on the RXVECTOR PSDU HT length field. [1]
This algorithm is expressed in the following pseudo-code (note, this is not optimized for efficiency):

```
ParseA-MPDU(length)
{
      Int offset = 0;   /* Byte offset from start of PPDU */
      While (offset+4 < length)
      {
            If valid_MPDU_delimiter(offset) &&
                  get_MPDU_length(offset) <= 2346 &&
                  get_MPDU_length(offset) <= (length – (offset+4))
            {
                  /* Valid delimiter */
                  Receive_MPDU(offset+4, get_MPDU_length(offset));
                  /* advance multiple of 4 bytes */
                  offset += 4 + 4*((get_MPDU_length(offset)+3)/4);
            }
            Else
            {
                  /* delimiter invalid, advance 4 bytes and try again */
                  offset += 4;
            }
      }
}
```

Note: the unique pattern is used to reduce computation required while scanning for a valid delimiter. In this case the receiver tests each possible delimiter for a matching unique pattern. Only when a match is discovered does it then check the CRC.

6.4.3 Single receiver frame aggregation format

An aggregate is a sequence of MPDUs carried in a single PPDU carrying the aggregate attribute. In a single-receiver aggregate, all MPDUs are addressed to the same receiver address. The aggregate contains MPDUs in the order defined below.

Table 7 Single Receiver Aggregation MPDUs

MPDU	Multiplicity	Description
IAC or RAC	At most one	Initiator or responder aggregate control. Used when it is necessary to include a logical element as described in sections 6.1.1 and 6.1.2
ACK	At most one	Contains an acknowledgement of at most one MPDU in the previous aggregate that required immediate acknowledgement.

[1] This procedure will occasionally wrongly interpret a random bit-pattern as a valid delimiter. When this happens, the MAC will attempt to interpret a random MPDU. The MAC will discard it based on a bad MAC CRC check. The overall effect may be to lose one or more MPDUs from the aggregate very infrequently.

BA	At most one per Block ACK agreement in the appropriate direction	Block ACK MPDUs
QoS Data, MHDR or CHDATA	No limit, except as constrained by Block ACK protocol and TXOP duration When MAC Header compression is applied the aggregate shall include at least one MHDR MPDU per TC and receiver.	Data MPDUs sent under Block-ACK policy. These MPDUs can come from multiple TCs, provided that they comply with the TC policy of the current TXOP. Apart from QoS Data MPDUs the aggregate may also contain MHDR- and CHDATA MPDUs (both of Type 'EDATA').
BAR	At most one per Block ACK agreement (immediate Block ACK policy).	Immediate Block ACK Requests
Any MPDU requiring an immediate response	At most one	At most one MPDU that requires an Ack MPDU response may be transmitted. This includes BAR sent under delayed Block ACK policy, Data MPDU (including non-QoS), Management MPDU.

At most one single Block ACK Request using the delayed Block ACK policy is permitted. This is acknowledged by an acknowledgement MPDU in the following aggregate.

6.4.4 Multiple receiver frame aggregation format

A multiple-receiver aggregate contains MPDUs that are addressed to multiple receivers. The multiple receiver aggregate (MRA) is distinguished by the presence of the MRAD MPDU that lists the receiver addresses referenced in the aggregate. The MRA contains a MRAD MPDU, followed by one or more groups of MPDUs organized by receiver address.
Each group is either a response group or a non-response group. The rules for contents of the two groups are similar, except that the response group contains an IAC that defines a response period.
A Broadcast/Multicast Data Group contains MPDUs in the order defined below.

Table 8 - Broadcast/Multicast Receiver Group MPDUs

MPDU	Multiplicity	Description
Data or management MPDU	No limit, except as constrained TXOP duration	Data or management MPDUs with the broadcast or the same multicast destination address.

A Scheduled Response Receiver Group shall contain the same MPDUs as a Single Receiver Aggregate (see section 6.4.3), addressed to the same receiver, in the order defined in Table 7.

A No Response Receiver Group shall contain the following MPDUs, addressed to the same receiver, in the order defined below.

Table 9 - No Response Receiver Group MPDUs

MPDU	Multiplicity	Description
IAC	At most one	Initiator aggregate control MPDU Note this cannot contain an RDG.
ACK	At most one	Contains an acknowledgement of at most one MPDU in the previous aggregate that required immediate acknowledgement.
BA	At most one per Block ACK agreement in the appropriate direction	Block ACK MPDUs
QoS Data, MHDR or CHDATA	No limit, except as constrained by Block ACK protocol and TXOP duration When MAC Header compression is applied the aggregate shall include at least one MHDR MPDU per TC and receiver.	Data MPDUs sent under Block-ACK policy. These MPDUs can come from multiple TCs, provided that they comply with the TC policy of the current TXOP. Apart from QoS Data MPDUs the aggregate may also contain MHDR- and CHDATA MPDUs (both of Type 'EDATA').

The MRA aggregate shall contain the following (groups of) MPDUs in the order defined below.

Table 10 - MRA Aggregate MPDUs

MPDU	Multiplicity
MRAD	One
Broadcast/Multicast Data Group	One per broadcast or multicast destination address. Only permitted when the BSS is operating in pure mode.
Scheduled Response Receiver Group or No Response Receiver Group	One per receiver address.

Scheduled Response Receiver Groups and No Response Receiver Groups may be mixed in any order.
Data MPDUs can come from multiple TCs, provided that they comply with the TC policy of the current TXOP.
Data MPDUs with the same receiver address are included in consecutive locations in the MRA aggregate. The order of (groups of) Data MPDUs for different receiver addresses is equal to the order of the Receiver Info fields in the MRAD MPDU.

6.4.5 Aggregated MSDU Format (A-MSDU)

This section describes the format of Aggregated MSDU (A-MSDU). The purpose of the A-MSDU is to allow multiple MSDUs being sent to the same receiver to be aggregated and sent in a single MPDU, which improves the efficiency of the MAC layer, particularly when there are many small MSDUs such as TCP acks. Support for A-MSDU is mandatory at the receiver. However, the transmitter is free to use A-MSDU or not based on information such as traffic characteristics and link conditions. The maximum length of an A-MSDU is limited so that the length of an MPDU containing an (unfragmented) A-MSDU shall not exceed 4095 bytes.

Figure 18 shows the structure of a Data MPDU containing an A-MSDU, which is a sequence of *n* MSDU sub-frames. Each sub-frame consists of a sub-frame header followed by an MSDU and 0-3 bytes of padding. The MSDU shall not exceed 2304 byes in length. A sub-frame (except the last) is padded so that its length is a multiple of 4 bytes. The last sub-frame has no padding.

The sub-frame header contains three fields: Destination Address (DA), Source Address (SA), and Length. [2] The DA and SA fields are interpreted as described in Table 12 and Table 13 below. The length field contains the length in bytes of the MSDU.

[2] Note, these fields are arranged in the same order as the 802.3 frame format.

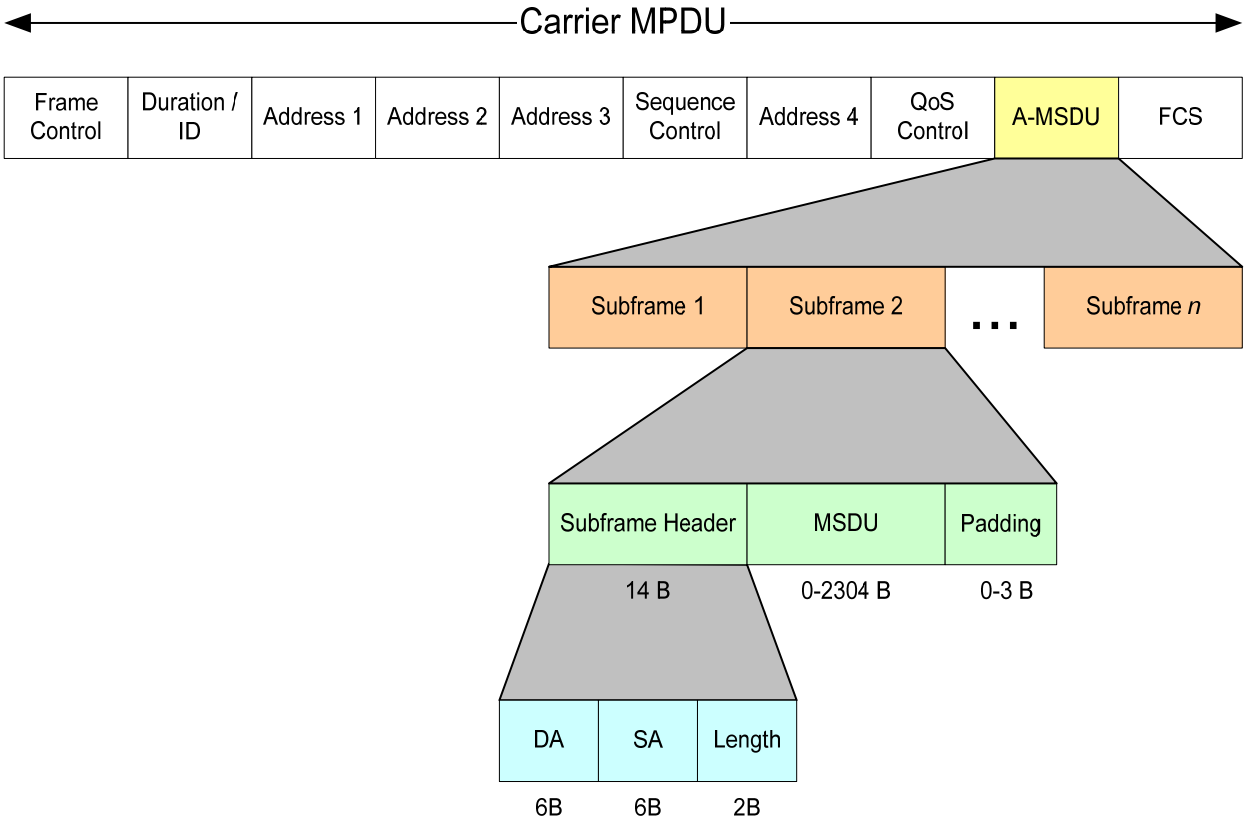


Figure 18 – A-MSDU Structure

The MPDU containing the A-MSDU is one of the QoS data subtypes. The presence of the A-MSDU is signaled by a bit in the QoS control field, whose contents are described in the table Table 11.

Table 11 – QoS Control Field

Applicable Frame (sub) Types	Bits 0-3	Bit 4	Bit 5-6	Bit 7	Bits 8-15
QoS CF-Poll and QoS CF-Ack+CF-Poll frames sent by HC	TID	EOSP	Ack Policy	Reserved	TXOP limit
QoS Data+CF-Poll and QoS Data+CF-Ack+CF-Poll frames sent by HC	TID	EOSP	Ack Policy	A-MSDU Present	TXOP limit
QoS Data and QoS Data+CF-Ack frames sent by HC	TID	EOSP	Ack Policy	A-MSDU Present	QAP PS Buffer State
QoS Null and QoS CF-Ack frames sent by HC	TID	EOSP	Ack Policy	Reserved	QAP PS Buffer State
QoS Data and QoS Data+CF-Ack frames sent by non-AP QSTAs	TID	0	Ack Policy	A-MSDU Present	TXOP duration requested
	TID	1	Ack Policy	A-MSDU Present	Queue size

QoS Null and QoS CF-Ack frames sent by non-AP QSTAs	TID	0	Ack Policy	Reserved	TXOP duration requested
	TID	1	Ack Policy	Reserved	Queue size

The DA and SA are passed between the MAC and its service client in the MAC service primitives. The mapping of DA, SA to the addresses in the MPDU header is described by the following tables. It is possible to have different DA and SA in the same A-MSDU as long as they all map to the same receiver address.

Table 12 – Address field contents for unicast MPDUs

To DS	From DS	Address 1	Address 2	Address 3	Address 4	DA	SA
0	0	DA	SA	BSSID	X	-	-
0	1	DA	BSSID	-	X	-	SA
1	0	BSSID	SA	-	X	DA	-
1	1	RA	TA	-	-	DA	SA

Table 13 - Address field contents for broadcast/multicast MPDUs

To DS	From DS	Address 1	Address 2	Address 3	Address 4	DA	SA
0	0	Bcast Addr.	SA	BSSID	X	DA	-
0	1	Bcast Addr.	BSSID	-	X	DA	SA
1	0	BSSID	SA	-	X	DA	-
1	1	RA	TA	-	-	DA	SA

Notes:
1. Fields marked with X are not present
2. Fields marked with "–" are present but their contents are reserved
3. Broadcast/multicast MPDU transfer follows the procedure described in Section 9.2.7 of IEEE 802.11-2003. Thus broadcast/multicast frames sent by non-AP stations are directed at the AP and are ACKed by the AP. In this case, a STA may aggregate both broadcast/multicast frames and unicast frames. The same is not true for frames sent by AP, since broadcast/multicast and unicast frames are mapped to different receiver address (Address 1).

After A-MSDU aggregation, the A-MSDU is treated like any other MSDU. It be be transported using either QoS Data or CHDATA MPDUs. It may be fragmented.

7 MAC sublayer functional description

7.1 Aggregation Exchange Sequences and related rules

7.1.1 IAC/RAC Exchange Rules

This section describes how IAC/RAC MPDUs are used in a bi-directional exchange.
An exchange protected using LongNAV rules shall start with the exchange of non-aggregate IAC and RAC MPDUs at a basic rate. An RTS/CTS exchange may be substituted for an IAC/RAC exchange if none of the features of the IAC/RAC are required by the initiator. Control frames intended to modify the NAV (IAC, RAC, CF-End and SCB) are sent at a basic rate.
A STA shall respond to a PPDU carrying a non-aggregate IAC MPDU with a non-aggregate RAC response MPDU. Note, in this case the IAC MPDU does not include an RDG value. The response RAC is generated regardless of the absence of the RDG.
Following the initial exchange, for all PPDUs for which a response is permitted, the initiator may include an IAC in the aggregate containing the permitted response grant.
When the initiator schedules a response, the initiator shall allow enough RDG time for a response RAC and any requested BAs. The responder does not need to include an RAC if no elements of the RAC need to be transmitted. If the response is a single MPDU (i.e. a Block Ack), the responder may send this as an aggregate or non-aggregate.
Note: under LongNAV rules, the initiator can re-use any time not used by the responder in its response. Under pairwise spoofing rules, this is not the case.

7.1.2 Non-aggregate IAC/RAC Rate selection Rules

An initiator that needs to set the NAV in its locality (i.e. LongNAV rules for initial exchange) shall use a basic rate. In a mixed BSS, it shall use a basic legacy rate and PPDU format.
A responder responding to an IAC with an RAC MPDU assumes that it is required to set the NAV in its locality. It shall use a basic rate. In a mixed BSS, it shall use a basic legacy rate and legacy compatible PPDU format.

7.1.3 Response Period Scheduling

The initiator may grant reverse direction to:

- Respond to an RDR contained in the previous PPDU received from this STA
- Respond to a periodic RDR TSPEC at the start of its service period
- Respond to a transmission failure in the previous PPDU received from this STA
- Allow for an RAC MPDU
- Allow for a BA response to any BAR it is transmitting to the STA in its current PPDU.
- Allow for an Ack response to any MPDU it is transmitting that requires an ACK response.

The grant is in units of time and should use the latest MCS value for that STA to convert the known amount of data to grant value. This MCS value is the latest value either sent to the STA (as an MFB element, including one in the current PPDU) or received from the STA (as a default MCS).

7.1.3.1 Retransmission Rules

When transmitting an MRA containing periodic RDR, or when transmitting a SRA under the Pairwise Spoofing rules, the initiator shall allocate in its next RDG sent to a responder during the current frame exchange sequence, a grant for MPDUs transmitted by the responder that were not received successfully at the initiator.

7.1.4 Responder Rules

An HT STA that receives an IAC containing an RDG shall, in response, transmit a single non-aggregate MPDU or a single receiver aggregate PPDU addressed to the initiating HT STA. This response shall begin no earlier than the scheduled response period and end no later than the end of the scheduled response period. If the IAC contains an RTS element, the responder shall generate no response if its NAV is non-zero at the start of the response period, otherwise the responder shall respond with an RAC containing a CTS element.
A responder should respond to the MRA initiator at the starting time of its granted response period, regardless of whether it is receiving frames or is physically sensing the medium busy at the starting time of its RDG.

7.1.5 Response PPDU Contents

When sending a response PPDU, the IAC's Reverse Direction Traffic Identifier (RDTID) identifies what Data MPDUs are eligible to be transmitted, as follows.

When the RDTID contains a TSID, only MPDUs of this traffic stream may be transmitted. This is typically used with a periodic TSPEC to gather uplink MPDUs for that TSPEC.

When the RDTID contains an AC, only MPDUs of the specified access category may be transmitted. This is typically used by an initiator that gained access to the channel using EDCA to ensure that only MPDUs of that AC are sent by a responder.

When the RDTID contains the value "unconstrained", MPDUs of any traffic stream or user priority may be transmitted. This may be used by an initiator during a TXOP gained through HCCA.

7.1.6 Power-saving at the receiver of an aggregate PPDU

An HT STA may sleep through part of the aggregate PPDU under the following conditions:

- It correctly received an initial MPDU that was not a MRAD, and was addressed to another receiver
- It correctly received an initial MPDU that was a MRAD that did not contain its receiver address or broadcast address

If the HT STA correctly receives a MRAD that caused it to remain awake because it contained the broadcast address, the HT STA may sleep as soon as it decodes a unicast receiver address in the aggregate.
If the HT STA correctly receives a MRAD that caused it to remain awake because it contained its address, the HT STA may sleep once it has decoded an MPDU containing its receiver address followed by one that contains a different receiver address.

7.1.7 Protection mechanisms

7.1.7.1 MAC layer protection

The following techniques may be used to achieve protection using a MAC layer technique:

- A contention-free-period of a beacon containing a CF parameter set.
- The HC may deliver a polled TXOP to an HT STA. The whole of the TXOP may be protected by the duration field in the QoS CF-poll.
- The HC uses its privileged channel access (after a PIFS) and starts a CAP with transmission of a CTS-to-self MPDU.
- A TXOP may be started with an RTS/CTS or IAC/RAC exchange, providing NAV protection of the duration specified in these control MPDUs.

7.1.7.1.1 Long NAV

When a STA has a TXOP, it may set a long NAV to protect multiple PPDUs using a single protection MAC layer protection exchange, e.g., RTS/CTS.
The STA may be able to accurately predict the duration of these PPDUs, in which case it can set the duration values in the preceding protection exchange accurately.
However, it may not be able to predict the duration accurately. Setting a longer NAV allows it to respond to the following events:

- Retries of failed transmissions in the current exchange
- Adaptation of transmit parameters by training feedback during the current exchange
- Transmission of MSDUs arriving at the MAC Data SAP during the current exchange
- Reverse direction grants

LongNAV protection is defined as selecting a duration value, limited by the remaining duration in the current TXOP and setting the NAV to this value using one of the MAC layer protection techniques.

7.1.7.1.2 Truncation of TXOP

In the case that a STA gains access to the channel using EDCA and uses LongNAV to protect a duration value, and then runs out of MPDUs to transmit, the STA may transmit a CF-End provided that the remaining duration is long enough to transmit this MPDU.
This shall be interpreted by HT STAs and is interpreted by legacy STAs as a NAV reset – i.e. they reset their NAV timer to zero at the end of the PPDU containing this MPDU. [3].

[3] Note, the transmission of a single CF-End MPDU by the initiator resets the NAV of devices hearing the initiator. There may be devices that could hear the responder that had set their NAV that do not hear this NAV reset. Those devices will be prevented from contending for the medium until the original NAV reservation expires

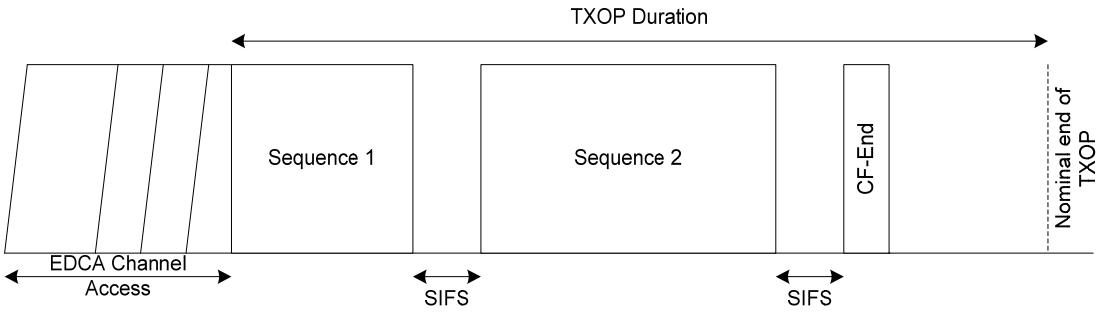


Figure 19 – Example of TXOP Truncation

In the example above the device accesses the medium using EDCA channel access and then engages in two sequences of PPDUs in the initiator role (these could be to the same receiver address, or different receiver addresses). Each sequence may include multiple PPDUs sent and received.
At the end of the second sequence, the initiator has no more data that it can send that fits within the TXOP, so it truncates the TXOP by transmitting a CF-End MPDU.
HT STAs and legacy STAs that receive this MPDU reset their NAV and can start contending for the medium immediately after this MPDU.

7.1.7.2 PHY layer protection – Spoofing

The Spoofed NAV Duration is virtually set in the PHY Header by using the Length and Rate Field of the legacy SIGNAL Field.
Previously, the Rate field declares the Rate that the packet is coded in after the PHY Header, and the Length field declares the Length of the packet (after the PHY header) in Bytes. When a legacy node receives this SIGNAL Field, it starts to decode the rest of the packet in the specified rate and will continue to do so until the end of Length / Rate time.
The Spoofed NAV uses this characteristic of the Length and Rate Fields, so that the (Length / Rate) – (EIFS – DIFS) is equivalent to the intended NAV Duration. A legacy node that is spoofed by these two fields will continue reception for that (Length / Rate) – (EIFS – DIFS) time, preventing it to start transmission during that period. This way transmission protection can be achieved without requiring that legacy nodes can receive the MPDU contents.

7.1.7.2.1 Pairwise Spoofing

Pairwise spoofing is a method of providing protection for an SRA aggregate exchange using the spoofing mechanism at both initiator and responder.

Pairwise Spoofing works as follows.

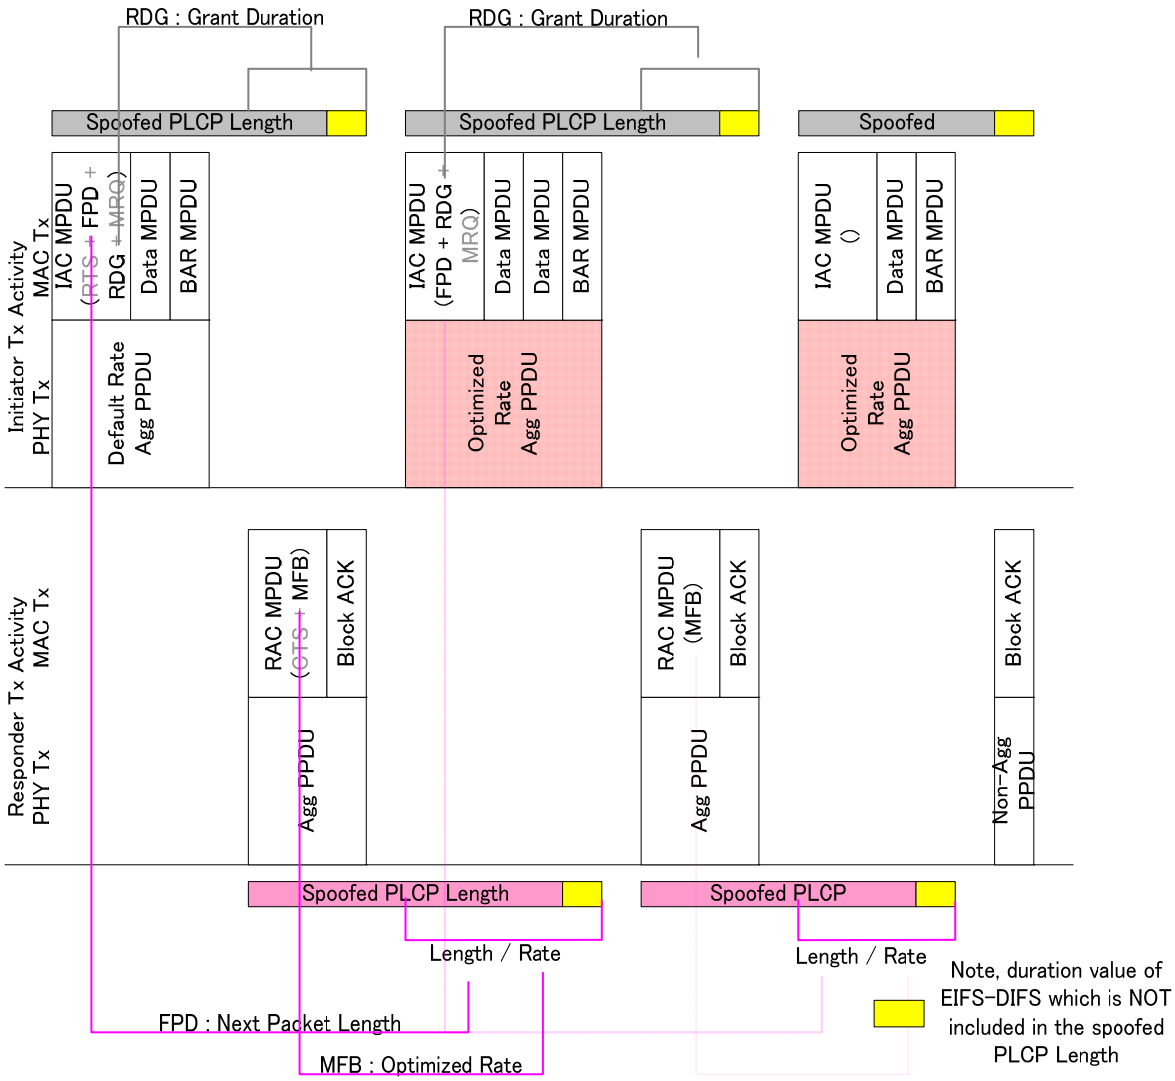


Figure 20 – Basic Operation of Pairwise Spoofing

The transmitter of a Data packet informs the recipient of the length of the packet it intends to transmit by indicating the FPD (Following Packet Descriptor) or RDR (Reverse Direction Request) field. When a STA receives an FPD or an RDR, it calculates the length, in time, it will take to receive that packet by dividing the length by the rate that the packet will be sent at. The MCS of the packet may be specified by the recipient using an MFB (MCS Feed Back) field for optimization purposes, otherwise it uses the default MCS notified in the FPD / RDR Field. The STA shall add to the length of the packet indicated in the FPD or RDR the length of any response MPDUs (e.g. RAC, BA, Ack) implied by its own transmission in determining the duration to spoof.

7.1.7.2.2 Single Ended Spoofing

Single ended spoofing provides protection of MRA aggregate exchanges using the spoofing mechanism at the initiator.

Single Ended Spoofing works as follows.

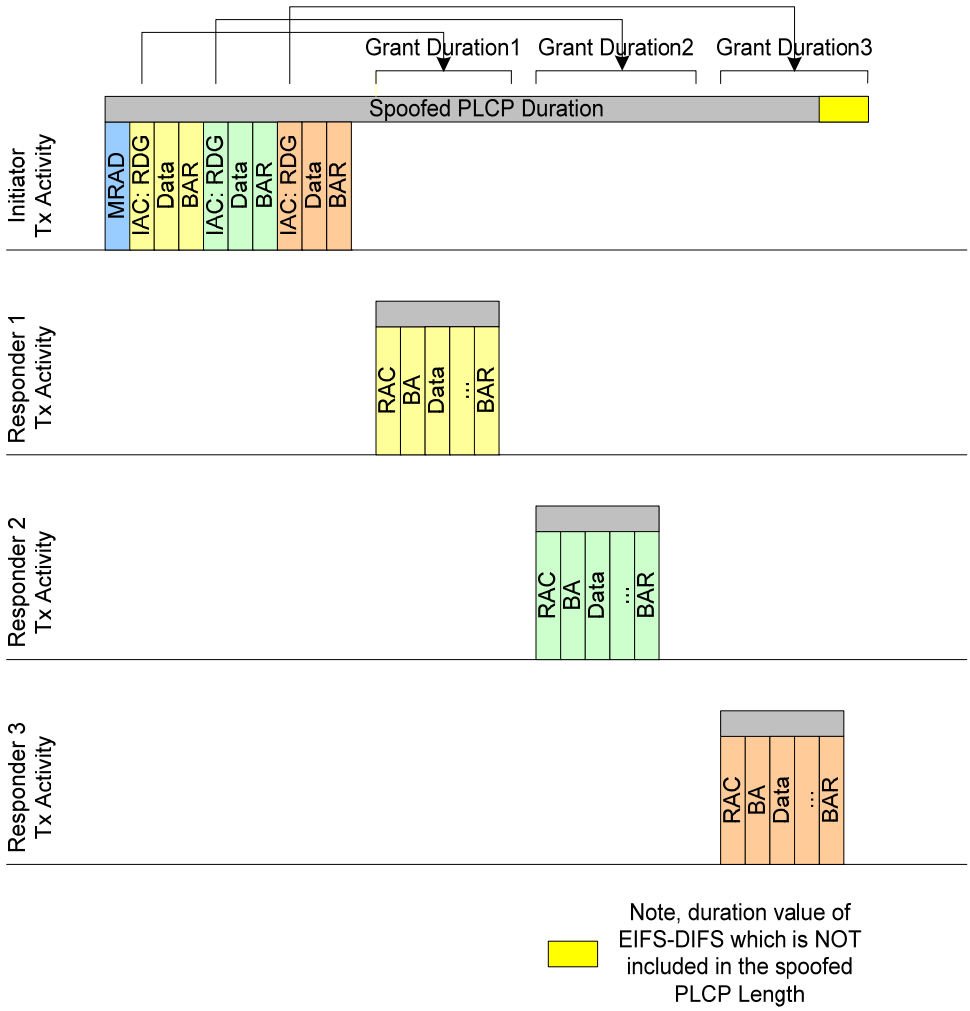


Figure 21 – Basic Operation of Single Ended Spoofing

In the case where single ended spoofing is used, an initiator should set its spoofed duration according to the duration it grants to reponses using the RDG field. The spoofed duration should be set so that it covers from the start of the initiator's frame to the end of the last grant duration minus (EIFS – DIFS).
In the figure above, single ended spoofing is used to protect responses to an MRMRA packet. The use of single ended spoofing is not limited to just MRMRA but may also be used for protection of an SRA sequence.

7.1.7.3 HT Device NAV and EIFS Rules

An HT device that receives a PPDU examines the legacy and HT rate and length fields from which it calculates legacy and HT durations.
It determines that spoofing is in operation if the legacy duration is not equal to the HT duration.
If spoofing is in operation, it shall set its NAV to the value of legacy duration + (EIFS – DIFS) – HT duration at the end of the received PPDU, unless the NAV has been set by reception of a duration field value within the PPDU. (See section 9.3.5 for a explanation of why this is necessary).
Note: EIFS is of less value in an HT network than a legacy network. It is difficult to define exactly what duration should be protected. Also use of spoofing and LongNAV provides NAV protection throughout aggregate exchanges containing MPDUs that a 3rd party HT STA is

unlikely to be able to hear. 3rd party HT STAs should not use EIFS at the end of an aggregate exchange because that would unfairly benefit the two peer STAs involved in the exchange.
The following rule defines EIFS selection for an HT STA:
An HT STA that receives an MPDU with a bad CRC shall only use EIFS if its NAV is zero at the start of the PPDU containing the MPDU. The condition is never true during an aggregate exchange, but is true for isolated Data/ACK.

7.1.8 MAC support for closed loop modes

7.1.8.1 Closed-loop MCS training

Support for closed-loop MCS adaptation is provided by the exchange of MRQ (MCS Request) and MFB (MCS Feedback) elements carried in IAC/RAC control frames.
An MCS Request can be included in any IAC or RAC. Its purpose is to request the peer STA to measure the characteristics of the link and return information in an MFB that allows the link to be used more effectively.
There are three types of MCS training exchange: normal, unsolicited and unanswered.
There is no negotiation or indication about which type of training exchange will take place. The peer initiating the exchange shall be prepared for a normal or missing response. It shall also respect unsolicited MFB.
A STA receiving an MFB should use the information contained in it to adapt its transmission parameters. A STA using the Pairwise Spoofing mechanism shall use the information received in the MFB to define the transmission parameters of the next PPDU to be transmitted in the current exchange sequence to that peer.

7.1.8.1.1 Normal Adaptation

In a normal training exchange, an MFB element is present in the next aggregate sent by the receiver of the MRQ element, and the aggregates are separated by a SIFS interval.

7.1.8.1.2 Unsolicited Adaptation

In an unsolicited exchange, a STA may send an MFB element without any matching MRQ element at any stage of an exchange.

7.1.8.1.3 Unanswered Training Exchange

In an unanswered training exchange, a STA ignores an MRQ element and refuses to provide an MFB element. In the case of pairwise spoofing, the STA that has sent the ignored MRQ element, shall use the default MCS defined in the FPD / RDR element.

7.1.8.2 Closed-loop MIMO training

Support for closed-loop MIMO adaptation is provided by the exchange of a TRQ (MIMO Training Request) element carried in IAC/RAC control frames and an HT sounding PPDU.
A TRQ can be included in any IAC or RAC. Its purpose is to request the peer STA to return an HT sounding PPDU that allows the requesting STA to use better parameters for subsequent transmissions to the peer STA.
Use of Closed-loop MIMO training is an option of the STA sending the TRQ. A STA that does not support closed-loop MIMO will never send an IAC or RAC with the TRQ flag set. All STAs support the creation of an HT sounding PPDU.

A STA that receives a PPDU carrying a TRQ element should ensure that its response PPDU is sent with the TXVECTOR SOUNDING parameter set to 1.

7.1.9 Aggregated Single Directional frame transfer

This section contains examples of aggregate exchanges for a single receiver address and for data transfer from the initiator to the responder.
In order to use aggregation, protection or legacy protection is required as defined by the operating mode.
In the LongNAV case, protection is achieved through the IAC/RAC exchanged transmitted at a rate such that all devices can receive these frames. An RTS/CTS exchange may be used if none of the IAC elements (apart from RTS) is required.
This is followed by a sequence of PPDUs. The PPDUs transmitted by the initiator shall be aggregate PPDUs. The PPDUs transmitted by the responder may be aggregate or non-aggregate PPDUs and must carry only response MPDUs and training feedback. These response PPDUs must not include any Data or Block ACK Request MPDUs.
The initiator prevents reverse direction data flow by defining an RDG to be large enough only for the expected response MPDUs: an RAC, one or more BAs and up to one ACK.

7.1.9.1 LongNAV examples

The first example shows use of LongNAV, unidirectional.
A pair of IAC/RAC MPDUs establishes NAV protection of the TXOP.
There then follows a sequence of transmissions of aggregates and reverse direction Block ACK responses from the responder. No RAC MPDUs are present in the response PPDUs because none of the logical elements are required to manage the exchange.

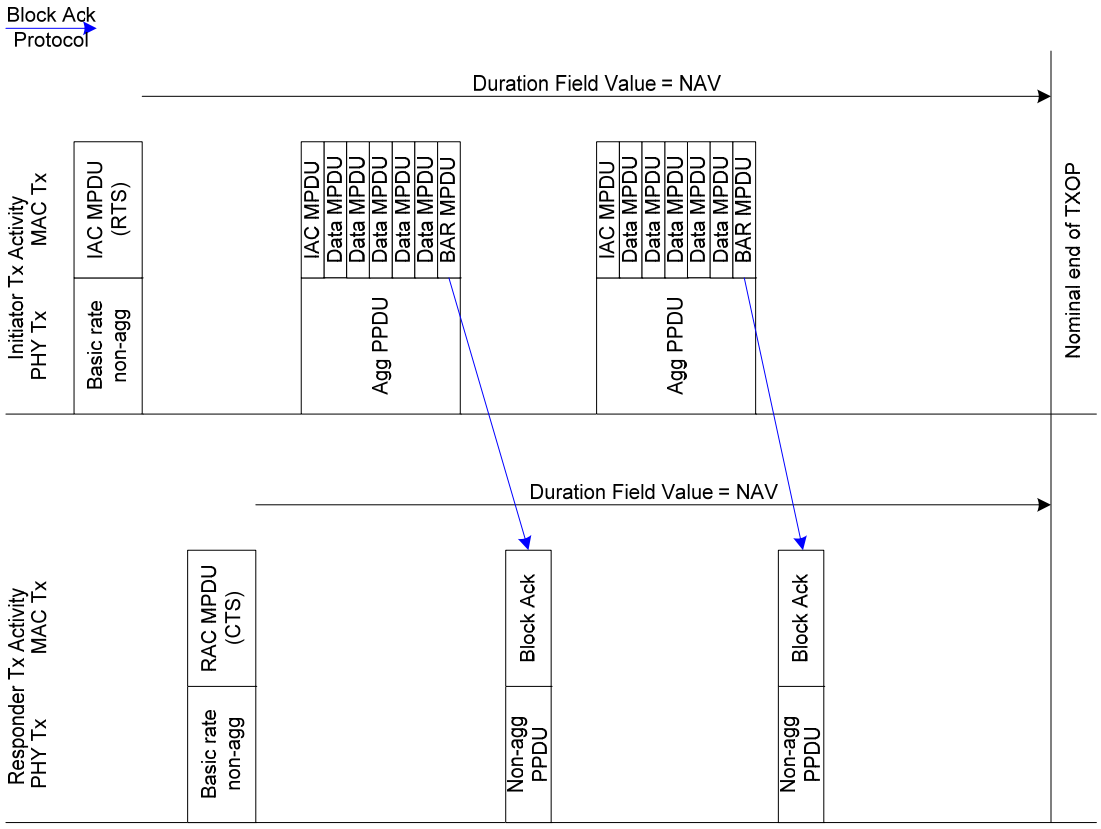


Figure 22 – LongNAV, untrained unidirectional exchange

The responder may want to request some reverse direction data transfer, or to send a training response in addition to sending a BA. In this case, it uses an aggregate to contain this information in its single response PPDU.

Note: The initiator does not know which of these the responder will elect to do, and so provides an RDG that supports the transmission of the longest response.

In the example below, the responder uses an aggregate to return a MFB (MCS Feedback) plus a Block ACK, in response to a MRQ (MCS Request) in the IAC and a BAR (Block ACK Request).

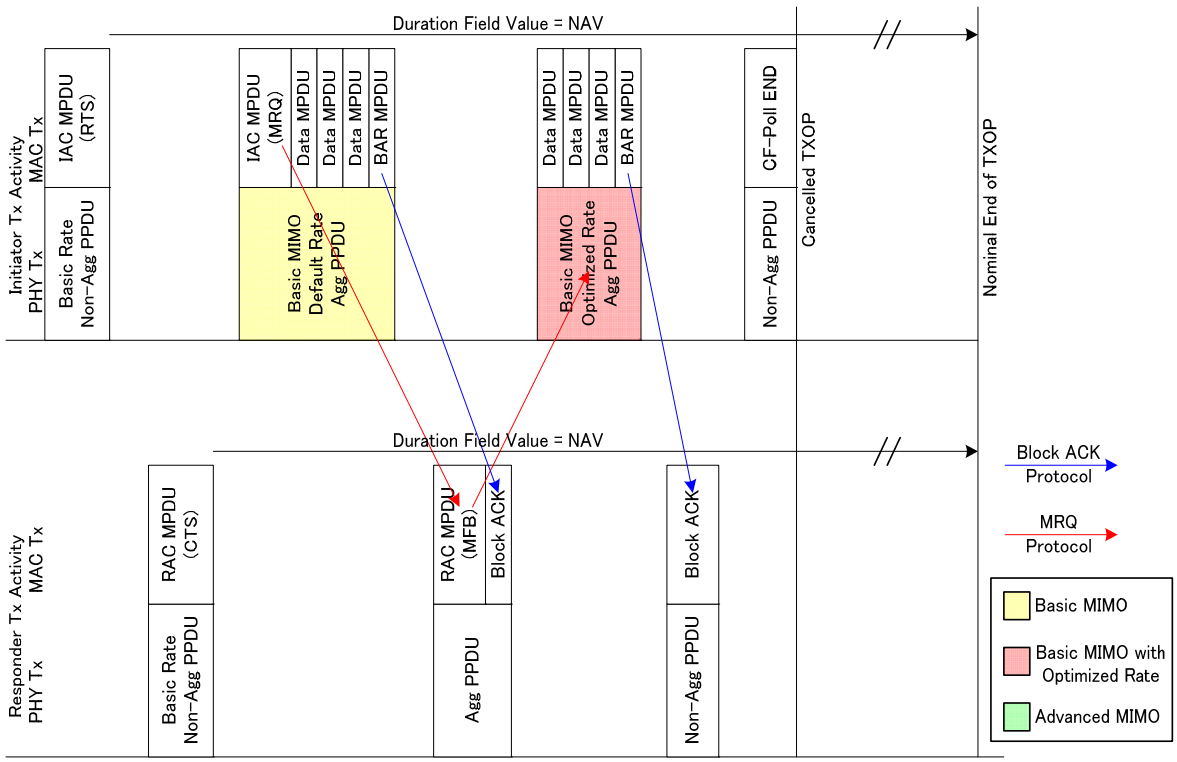


Figure 23 – LongNAV, untrained unidirectional showing aggregate response

The next diagram shows an example when training is involved.
First, a pair of IAC/RAC packets are sent at a basic rate (SISO mode), so that any third party nodes, including legacy nodes can set its NAV correctly. In the next IAC packet, the TRQ (Training Request) element is indicated to show that the following packet will be sent at an "Advanced MIMO Mode (MIMO with implicit Feedback)". In response to the TRQ, the responder will send its response PPDU with the TXVECTOR SOUNDING parameter set to 1. Also, in the second IAC, the MRQ (MCS Request) element is indicated to request a MCS Feedback, so that the modulation scheme may be selected by the responder.
When the initiator chooses to end the TXOP before the nominal end of the intended TXOP, it sends a CF-Poll End packet at a Basic Rate to cancel the TXOP.

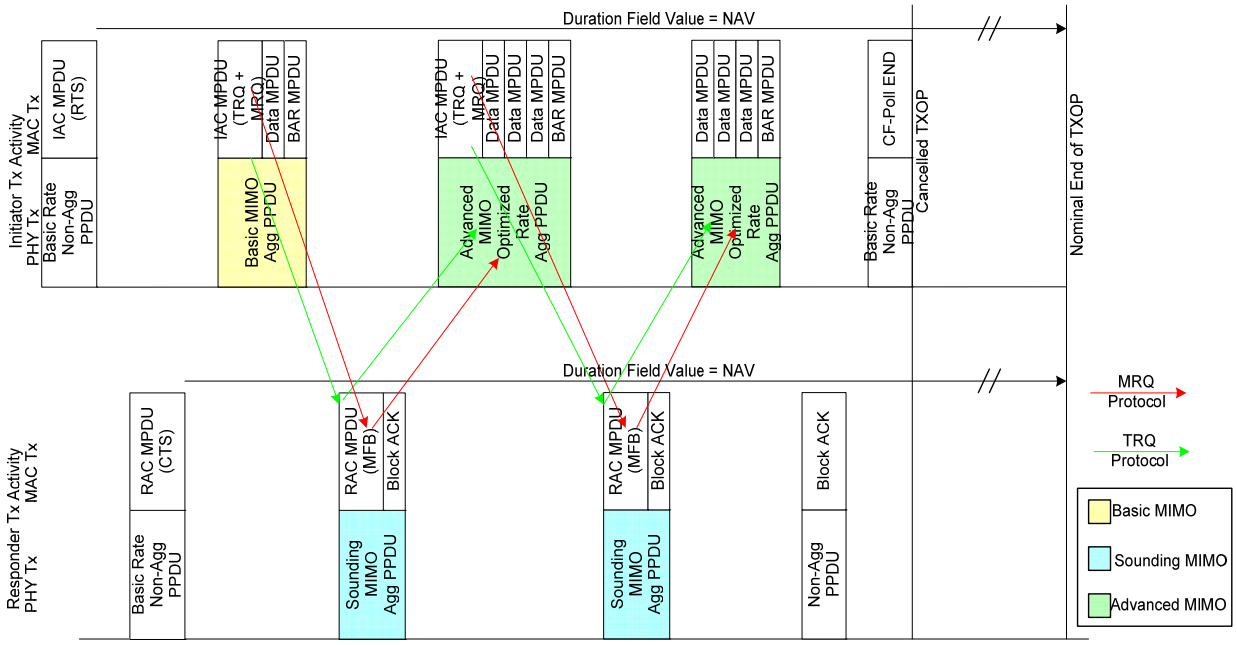


Figure 24 – LongNAV, trained unidirectional showing aggregate response

7.1.9.2 Pairwise Spoofing example

The next set of examples shows uses of Pairwise Spoofing, unidirectional.

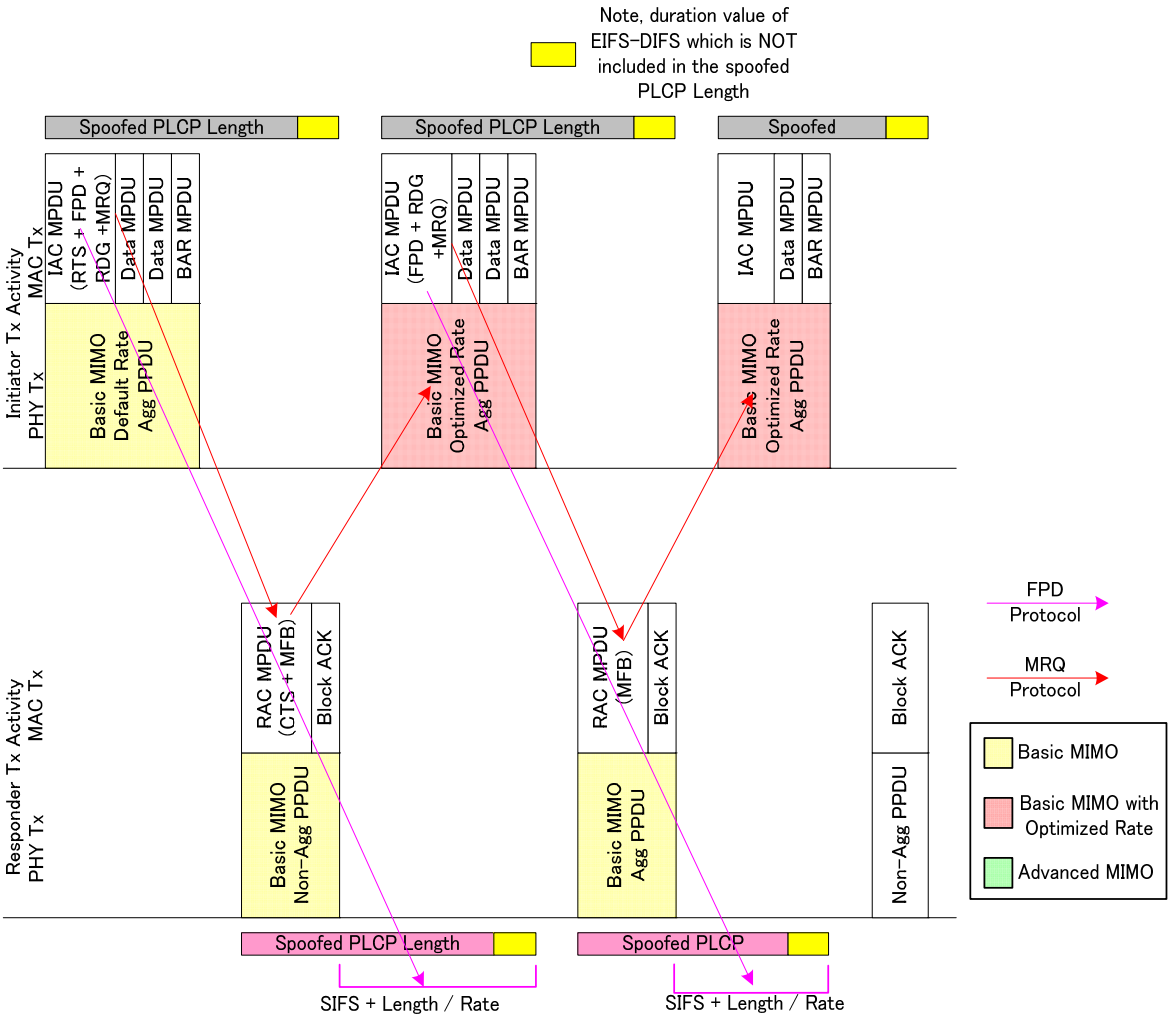


Figure 25 – Pairwise Spoofing, Untrained Unidirectional Exchange

An initiator starts with an IAC MPDU containing an FPD (Following Packet Descriptor) and an RDG big enough to hold the RAC response PPDU. Also, an MRQ is present to request MCS selection by the responder. The PPDU carrying the IAC MPDU will have a spoofed duration, set by the Rate and Length Field of the SIGNAL, which protects the RAC MPDU and Block ACK MPDU of known size.

The Responder answers with a RAC MPDU containing an MFB (MCS Feed Back) element. The MFB species the MCS at which it expects to receive the next PPDU from the initiator. When this MFB Field is not present, the Initiator will send its next PPDU at the default MCS specified in the FPD Field. The PPDU containing the RAC MPDU will have a spoofed duration of Length (in the FPD Field) divided by the Rate (either the default MCS or the MCS selected in the MFB) minus (EIFS – DIFS). For details on Spoofed NAV, see section 8.1.7.1.

After receiving the RAC MPDU, the initiator will send its aggregated MPDU at the MCS specified by the MFB Field. This aggregated MPDU will have a spoofed duration, which protects the fixed sized RAC MPDU and the expected response from the BAR (Block Acknowledgment Request) that it has sent.

This process will continue until the expiry of Initiator's TXOP or when the Initiator has no more packets to send.

The following example considers what happens when there is training involved.

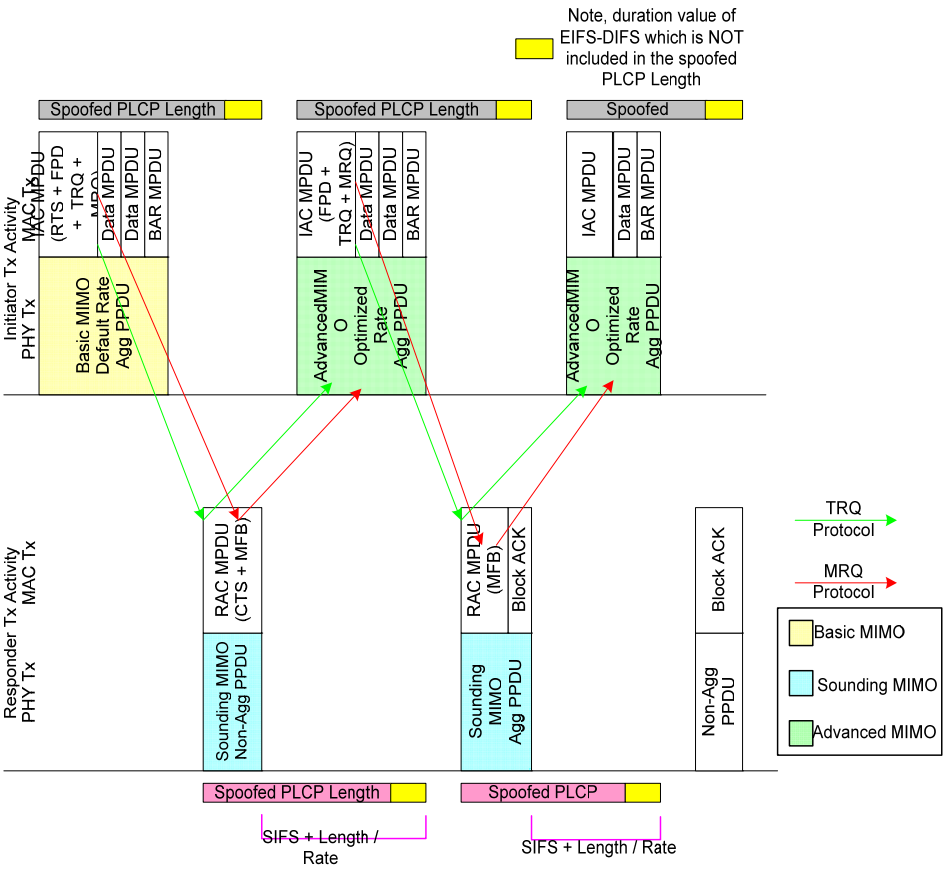


Figure 26 – Pairwise Spoofing, Trained Unidirectional Exchange

In Figure 26, the Initiator's MRQ element is used to optimize MCS selection at the responder. This optimized MCS is carried in the RAC MPDU.
Also, the TRQ (Training Request) element is present to show that the following packet will be sent at an "Advanced MIMO Mode (MIMO with implicit Feedback)". In response to the TRQ, the responder will send its response PPDU with the TXVECTOR SOUNDING parameter set to 1.

7.1.10 Bi-directional data transfer

This section defines the rules for setting of the RDL, RDR and RDG elements in order to achieve reverse direction data transfer, and then illustrates the use of these rules with example sequences.

7.1.10.1 RDL/RDR/RDG Rules

A STA that is an initiator may advertise willingness to provide reverse direction data flow by including a RDL element in its IAC.

A STA that is a responder may request a reverse direction allocation from the initiator using an RDR element. The RDR element contains the length of data in the PPDU it would like to send, plus a default MCS. An MFB element may accompany the RDR, and the initiator may determine an updated MCS based on their contents.

In determining an RDR value, a STA shall only consider those MPDUs that meet any RDTID constraint contained in the previous IAC.

The initiator receiving the RDR information (and possibly augmented by the MFB) calculates a grant duration. The RDG element in the next IAC contains the allocation for reverse direction data, plus any additional grants as defined in section 7.1.3.

The initiator is free to grant any RDG it wishes (compliant with the TXOP rules) including: no duration, the exact duration implied by the RDR, or any other duration determined by the implementation. Under pairwise spoofing rules an initiator should not grant an RDG that is longer than the RDR (plus additional duration as described in section 7.1.3), otherwise it may cause the channel to be used inefficiently.

7.1.10.2 TXOP and RDG

An initiator that grants one or more response periods using RDG shall ensure that that its current PPDU and any scheduled response periods all fit entirely within the current TXOP.
It is up to the initiator to decide whether to also allow for any response it may have to make to the transmissions made by the responder during the allocated RDG. For example, the PPDU sent next by the responder may include multiple BAR MPDUs.
The initiator may make an allowance for BA responses to these BARs in making its RDG allocation, in which case it will have time to transmit a response before the end of the TXOP, or it may not make allowance for its own responses. In the latter case, the responder may send a BAR, but there may be no time to send the BA from the initiator in the current TXOP. The responder will assume that the BAR was not received and will retry it subsequently.

7.1.10.3 LongNAV Bi-Directional Example (Informative)

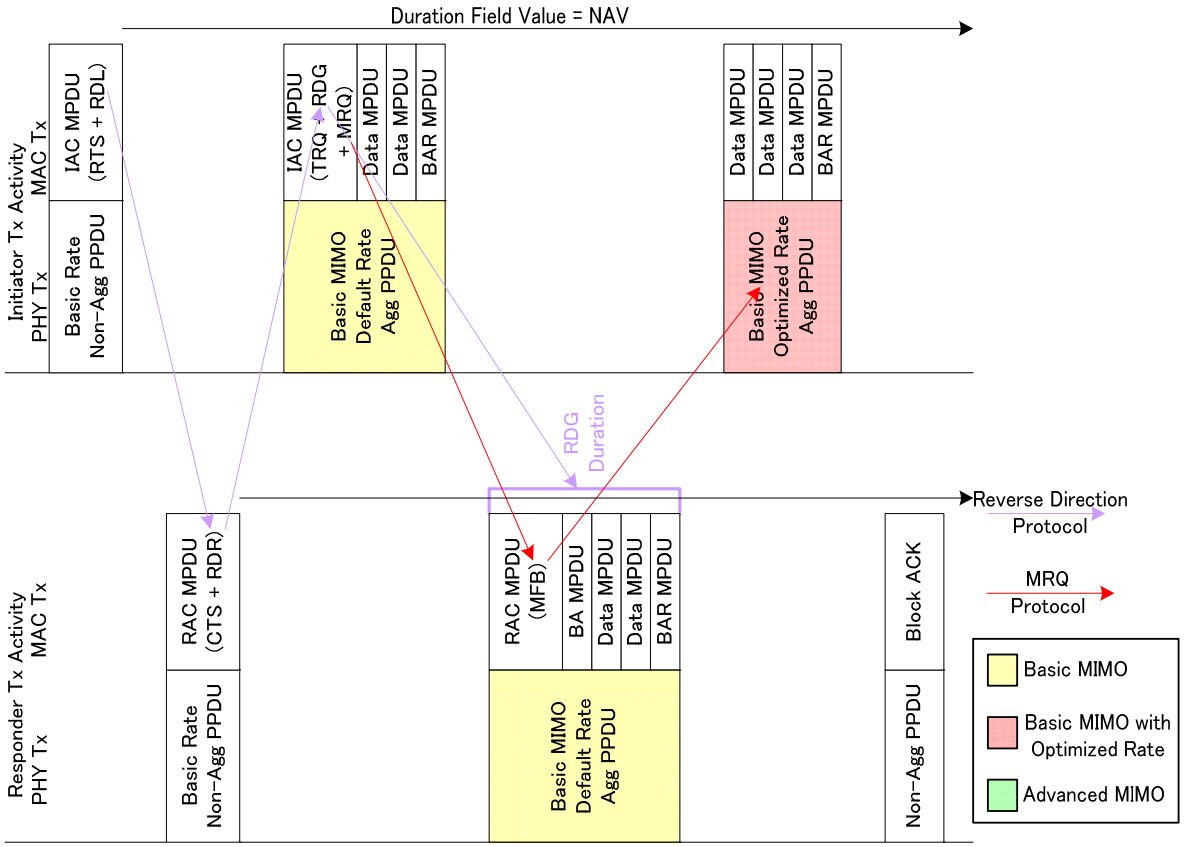


Figure 27 – LongNAV Bi-Directional Example

Initiator sets the granted RDG duration long enough to include the duration of reverse data plus any expected responses (RAC + BAs). The duration granted by the RDG for reverse data may be less than the requested duration, in which case the responder must reduce the amount of data it sends. The responder can also itself reduce the amount of reverse data it sends. Any response PPDU duration no longer than the RDG duration is valid.

7.1.10.4 Pairwise Spoofing Example (Informative)

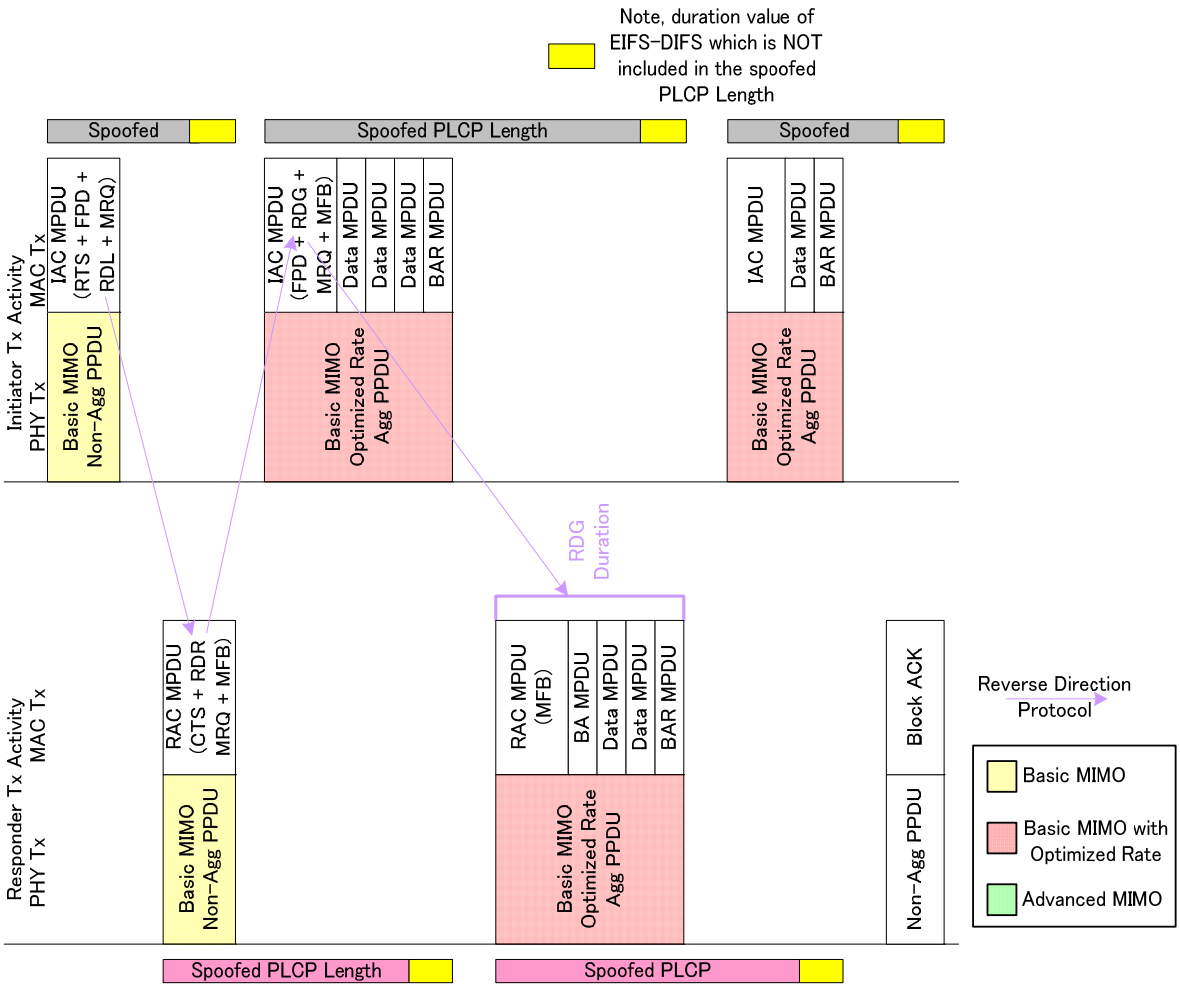


Figure 28 – Pairwise Spoofing, Trained BiDirectional

In Pairwise Spoofing Bi-Directional Transfer, the RDG Duration will include the duration of the reverse direction data plus any expected responses (RAC + BAs). The Spoofed NAV will include the duration of the PPDU transmission from the initiator plus the duration of reverse direction data and any expected responses (RAC + BAs).

7.1.11 Multiple receiver aggregation (MRA)

This section defines rules relating to the generation of multiple receiver aggregates by an initiator and relating to generation of responses to an MRA. It includes example sequences.

7.1.11.1 Rules and Scheduling of MRA

Generation of MRA PPDUs is optional for an HT STA. The ability to receive an MRA and generate a response to an MRA PPDU is mandatory for all HT STAs.
The transmitter of a MRA containing a reference to a periodic RDR TS is always an HT AP. A non-AP HT STA can also generate an MRA with or without responses (see example at end).
An HT STA may send MRA aggregates addressed to multiple HT STAs and request scheduled responses from none, some or all of those HT STAs.

The initiator schedules the response periods according to local policy, except as specified here. The first response period shall start no less than a SIFS after the end of the MRA PPDU. Typically the initiator will arrange that response periods are separated by SIFS, but there is no normative requirement for this to be the case.
An example of the scheduling scheme is shown in Figure 29, where the MRA addresses 3 HT clients.

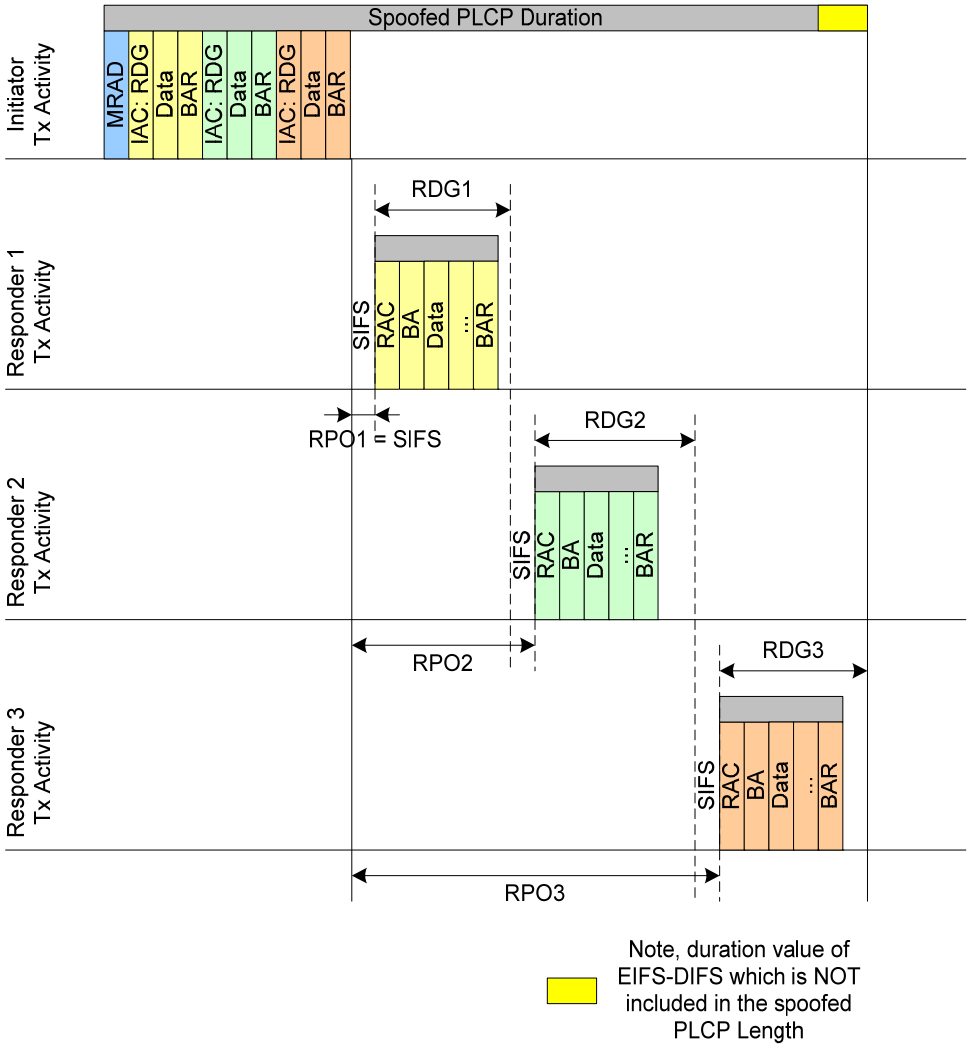


Figure 29 – Example of MRA response scheduling with 3 responders

7.1.11.1.1 Interspersing of periodic and aperiodic traffic

The HT AP may intersperse periodic and aperiodic traffic within the same MRA. It can intersperse periodic and aperiodic transmissions but it cannot intersperse periodic and aperiodic RDG to a single HT STA in the same response period since only one RDTID value can be included in an IAC.
Within the same sequence of aggregates, the HT AP may change the RDTID type to a particular receiver. So, an HT AP could service a periodic RDR TS and transition to servicing best effort traffic to the same HT STA in a later PPDU.

7.1.11.2 Protection Rules for MRA

An MRA exchange may be protected using either LongNAV or single-ended spoofing.
In the LongNAV case the NAV of all devices that can hear the initiator has already been set by transmission of a previous MPDU. In this case, the MRA PPDU should not carry a spoofed duration because it limits the ability of the initiator to recover unused RDG at the end of the response period. The single-ended spoofing case is applicable when there is no apriori NAV setting. In this case, the initiator sets the spoofed duration of the MRA PPDU to include from the start of this PPDU up to the end of the last scheduled response period as indicated by the IAC MPDUs within the aggregate.
In this spoofing case, the initiator does not generate FPD elements in its IAC MPDUs, and so the responders do not generate a spoofed duration for their response PPDUs.
In the case of LongNAV protection, the initiator may continue with a transmission a SIFS after the actual end of the last response PPDU. In the case of spoofing protection, the initiator may continue its transmission a SIFS after the end of the last response period, regardless of the actual length of the response.

7.1.11.3 Rate Selection

When a STA transmits a MRA to multiple receivers, it shall ensure that the MRA is sent using transmit parameters that should make the PPDU receivable by all receivers in the MRAD.
It has to respect their receiver capabilities (e.g. supported MCS values, number of rx antennas, optional coding capabilities) and select a “lowest common” set of parameters that should be consistent with reception at all these STAs.
If it contains multicast/broadcast Data, the rate is further constrained so a rate in the basic MCS set.

7.1.11.4 Interaction of MRA and TRQ+MFB (Informative)

Consider the two possible uses for MRA:

1. Periodic data
2. Best Effort data

In the periodic case, if the lifetime of the training information is longer than the interval between service periods, the initial MRA can contain training requests for multiple STAs. The STAs can generate forward direction rate feedback and a reverse direction training request in their responses. The initiator can then generate a final MRA containing the reverse direction rate feedback. So at the end of each cycle, rate feedback has been exchanged in both directions for use in the next service period.
In the best effort case, the initial MRA contains the training requests (and RDL values). The responses contain the rate feedback (and RDR). After the first MRA and responses, the forward direction is now trained, so an MRA containing data can be transmitted and also containing rate feedback and RDG for the reverse direction.

7.1.11.5 Examples of MRMRA frame transfer (Informative)

7.1.11.5.1 Periodic Traffic

To support periodic RDR TSs, the AP periodically sends an IAC containing a response period specific to that TS. These IACs may be aggregated into an MRA.
The PPDU is then followed by one or more response periods during which the responder sends any RAC, any BA, its data for the specified TS and a BAR.
At the end of the responses, the initiator then sends an aggregate containing Block ACKs. It may also contain IACs scheduling response grants for any transmissions from responders that were not correctly received as well as retransmissions of data that were not acknowledged.
The cycle of MRA followed by responses terminates when either there is no more time in the service period, or when all data to and from the initiator has been successfully received.

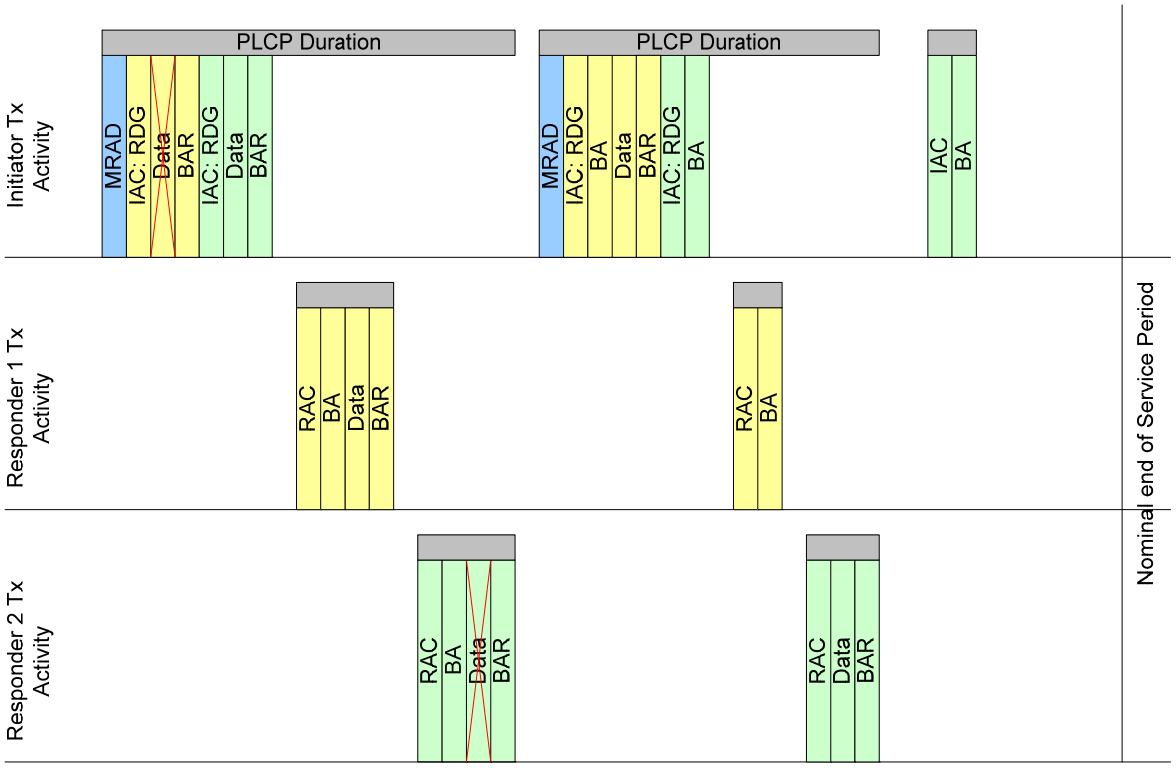


Figure 30 – MRMRA Example for Periodic Traffic

In the example shown above, an AP has two STAs with established periodic RDR in the uplink direction (the downlink is also periodic, and part of the same composite schedule).
The first PPDU is an MRA containing a MRAD that lists the two responders. For each responder there is an IAC containing a grant, some data and a BAR. The spoofed duration of the first PPDU includes the MRA and the two response periods.
The first responder generates a response PPDU containing an RAC (not strictly necessary), BA, Data, BAR. The BA indicates that the data from the initiator was not received correctly.
The second responder generates a response PPDU containing an RAC (not strictly necessary), BA, Data, BAR. The BA indicates that the data is received correctly.
The initiator generates a PPDU containing the retried data to responder 1 and grants to both responders. The grant to the first responder is allocated because the initiator has transmitted a BAR and expects to receive a BA in response. The grant to the second responder is allocated because the initiator did not correctly receive all of the uplink data and is scheduling an opportunity for the second responder to retransmit.
The responder 1 then sends the BA, and the responder 2 successfully transmits data and BAR.

The initiator needs to respond to the BAR and sends an aggregate containing the BA. Because this is sent to a single receiver, this aggregate need not contain a MRAD. Strictly, the IAC is not necessary, and the response could be sent in a non-aggregate PPDU.

7.1.11.5.2 Aperiodic Traffic

The MRA may also be used with aperiodic traffic. Aperiodic traffic has no implicit RDR at the start of the sequence. In this case, the initiator sends an MRA containing one or more IACs containing a scheduled response period long enough to contain an RAC and any expected BA MPDUs.
The initiator advertises its willingness to allow aperiodic traffic in an MRA response by including an RDL value and specifying an RDTID that permits aperiodic traffic.
The responder can also, at this point, generate an RDR to request additional response duration for its data transmissions to the initiator.
Note, this sequence is exactly the same as for a LongNAV bidirectional single-receiver sequence – except for the absence of the initial RTS/CTS.

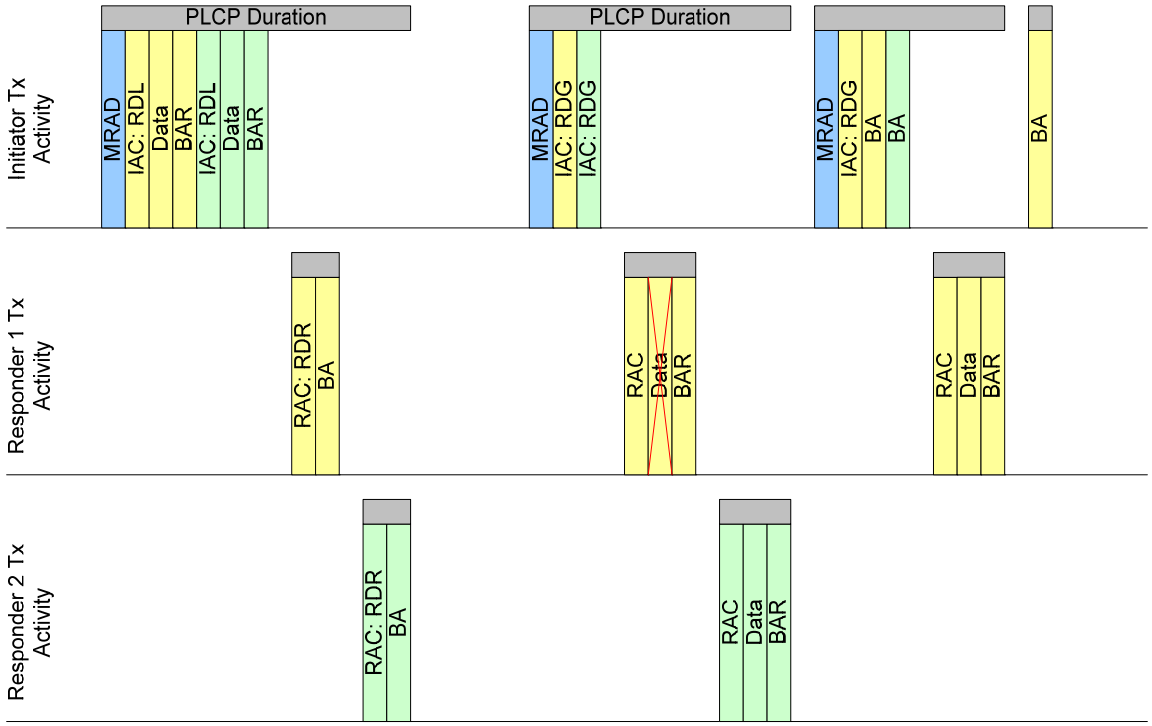


Figure 31 – MRMRA Example for Aperiodic Traffic

7.2 MAC Header Compression

7.2.1 Introduction (informative)

Frame aggregates can contain multiple Data MPDUs with similar MAC Headers. MAC Header Compression reduces the overhead associated with the MAC Headers of these MPDUs. The

efficiency gain of MAC Header Compression depends on the number and size of the Data MPDUs in the frame aggregate.
A transmitting STA or AP includes a MAC Header (MHDR) MPDU in an aggregate frame. This MPDU then provides MAC header information to multiple Compressed Header Data (CHDATA) MPDUs. Both MPDUs include a Header Identity (HID) to relate a CHDATA MPDU and its corresponding MHDR MPDUs.

7.2.2 Functional description

MAC Header compression is limited to the boundaries of a frame aggregate, i.e. only Data MPDUs within a single aggregate are compressed using MAC Header compression. As a consequence MAC Header (MHDR) MPDUs and Compressed Header Data (CHDATA) MPDUs shall not be transmitted as non-aggregate MPDUs.
The following requirements apply to a transmitting and receiving STA or AP in an 802.11n network:

- It is optional for a transmitting STA to use MAC Header compression in a frame aggregate.
- It is mandatory for a receiving STA to be able to receive frames with compressed MAC Headers in an aggregate.

7.2.2.1 Transmitting side

When the transmitting STA uses MAC Header compression it includes a MAC Header (MHDR) MPDU, identified by a Header Identity (HID), in the aggregate.
The transmitting STA shall include one or multiple Compressed Header Data (CHDATA) MPDUs, each having the same HID as the corresponding MHDR MPDU.
The transmitting STA may include multiple copies of MHDR MPDUs to increase robustness in case of loss of a MHDR MPDU. These MHDR MPDUs may be located anywhere within the sequence of CHDATA MPDUs subject to the constraints specified in section 7.2.3.
The transmitting STA performs MAC Header compression per destination and per Traffic Class separately.
Retransmissions of CHDATA MPDUs may be as either uncompressed Data or CHDATA MPDUs.

7.2.2.2 Receiving side

A receiving STA determines whether MAC Header compression has been used based on the Type and Subtype definitions in the received MPDUs.
A receiving STA shall use the fields from the MAC Header (MHDR) MPDU to replace the missing fields of the Compressed Header Data (CHDATA) MPDUs, having the same HID value, to recreate full Data MPDUs.
MHDR MPDUs are not acknowledged by the receiving STA or AP.
The acknowledgement policy of CHDATA MPDUs is defined in the QoS field as present in the corresponding MHDR MPDU.
Retransmissions of corrupt or non-received CHDATA MPDUs by the transmitting STA or AP are supported by the presence of sequence numbers in these MPDUs.
CHDATA MPDUs for which no corresponding MHDR MPDU is found in the frame aggregate, shall be discarded.

7.2.3 Ordering constraints for MHDR and CHDATA MPDUs

These constraints are designed so that only a single MHDR needs to be stored by a receiver in order to fully reconstruct an aggregate containing CHDATA MPDUs.
A CHDATA MPDU for a particular HID shall be preceded by at least one MHDR MPDU for the same HID. A MHDR for a different HID shall not occur between a CHDATA and its matching MHDR MPDU(s).
Within an aggregate, MHDR MPDUs that are not identical shall have different HID values.

7.2.4 Multiple receiver aggregates

The Header Identity (HID) ensures a unique link between a single MAC Header MPDU and a CHDATA MPDU. For each receiver address in a Multi Receiver Aggregate, a MHDR MPDU is included in the aggregate with a unique HID. The subsequent CHDATA MPDUs include the HID for the corresponding MHDR MPDU.

7.3 Multirate support

The Supported MCS Set (6.3.1) and Basic MCS Set (6.3.2) indicate MCS values for HT format PPDUs only.
The existing basic rate and supported rate sets indicate rate selections for legacy format PPDUs only.
Rate selection for control frames:
A control frame that needs to be received by legacy devices (i.e. to set the NAV) shall be selected from basic rates in the BSS's basic rate set, and transmitted using the legacy PPDU format.
Multicast/broadcast frames in pure mode and control frames that need to be received by HT devices only shall be selected from the Basic MCS set.
Control frames that need to be received by a single HT device only may be sent at any rate supported by the intended receiver. An example of this is a single Ack/BlockAck by a responder within a sequence protected by LongNAV.
Under pairwise and single-ended spoofing rules, it is not necessary for any third part HT STA to be able to receive MPDUs, so aggregate PPDUs may use any MCS supported by the intended receiver(s).

7.4 Quality of Service (QoS)

7.4.1 802.11e Requirements

An HT device shall support the mandatory features of an 802.11e QSTA.
An HT device shall support the use of the Block ACK mechanism introduced in 802.11e.

Note: this document extends the 802.11e Block Ack by defining a compressed Block Ack frame format, and rules for implicit Block Ack request in an aggregate. These are defined in section 6.1.6 and 6.3.5.

The default TXOP limit is 3.008 ms for Best Effort (BE) traffic. This value needs to be validated by simulation in order to determine the effect of this value on bandwidth differentiation between Best Effort and Video (VI) traffic. This value should be considered to be a place-holder.
A TXOP limit of zero is considered to allow an Aggregate exchange sequence containing the transfer of at most a single MSDU or MMPDU, which may be fragmented and acknowledged using the block Ack protocol.

7.4.2 Periodic RDR

7.4.3 Introduction (Informative)

Periodic RDR is a new type of traffic specification (TSPEC) that results in the HC seeing a periodic reverse direction request for data associated with that TSPEC.
In a conventional periodic uplink TXOP, the HC delivers a polled TXOP for a known duration during a known service period to a QSTA. The QSTA can then use this TXOP for any purpose it chooses. It can send MPDUs to the HC or any other QSTA using DLP. It is responsible for retrying any transmission failures.
A periodic RDR is much more constrained. The HC will deliver an RDG of duration sufficient to collect the first transmissions of data during a known service period. The HC also allows in the scheduling of the service period enough time to send a subsequent RDG to collect retransmissions of failed data.
The HC can service multiple QSTAs during the same service period. In practice it is likely to group together into a composite service period RDG for QSTAs with similar TSPECs (i.e. service interval and delay limit) and possibly PHY parameters. How it does this is outside the scope of the standard.

7.4.4 Meaning of Service Period with Periodic RDR

When a STA creates a TSPEC using periodic RDR, the AP shall create a service period during which the STA shall be awake to receive data and RDGs.
The AP may add TSPEC requests from multiple STAs into a composite service period during which it will service the STA. This allows it to use MRMRA.
Note, as STAs are added to or removed from the composite service period, the AP may need to update the service period notification by transmitting an updated service period element to the STAs.
The AP shall ensure that any IAC transmissions intended to service the periodic RDR traffic stream lie wholly within the service period.
A STA that receives data during an MRMRA sequence may go to sleep as soon as it has received and acknowledged all data sent to it by the initiator, and has received acknowledgement for all data it has sent.

7.5 Enhanced Block acknowledgment

7.5.1 Introduction (Informative)

The 802.11e Block Ack mechanism is enhanced with the usage of implicit Block Ack Request when using frame aggregation.
An aggregated frame need not include a BAR frame and the recipient interprets the receiving aggregation frame as though it includes a BAR frame. This is the Implicit BAR mechanism.
A compressed variant of the 802.11e Block Ack frame is also defined. The compression parameters are defined during the Block Ack set up by an ADDBA request/response frame

exchange based on whether to use fragmentation or not and the maximum number of MSDUs in the Block Ack window.
The size of the compressed BA frame is defined by the BA agreement, and is constant throughout the lifetime of the BA agreement.

7.5.2 Implicit BAR mechanism

The originator may omit the inclusion of a BAR frame in an aggregated frame.

An HT STA

1. that is permitted to generate a response PPDU as defined in section 7.1.4, and
2. receives directed data frames for a particular TID sent under BA acknowledgement policy as part of an SRA or MRA, and
3. For which it receives no matching BAR in the aggregate

interprets this as an implicit BAR, and shall prepare a BA for the TID for transmission in its response PPDU.

In legacy 802.11e, the Starting Sequence number in the BAR frame causes the recipient to indicate complete buffered MSDUs to its higher layer and release those buffers. The Starting Sequence number also defines the Starting Sequence number for the response BA.

For the Implicit BAR mechanism, as the BAR frame is omitted in the aggregate, these functions are determined as defined below:

Condition	Receive Buffer Update	BAR starting sequence number
Aggregate received and the first data MPDU for this TID is received correctly	Update according to the first data MPDU's sequence number	Use this MPDU's sequence number
Aggregate received and the first data MPDU for this TID is not received correctly	Do not update the receive buffer	Use the seqence number of the first data MPDU received correctly. In an aggregate MPDUs are strictly ordered by sequence number
A BAR MPDU is received	Update according to the starting sequence number contained in the BAR	Use the starting sequence number in the BAR MPDU

For the multiple TID aggregate, the first data MPDU of a TID is distinguished if the last data MPDU of the previous TID has no error. The first data MPDU for the directed part in the MRA case can be also distinguished in the same way.

The reception buffer management for the Implicit BAR mechanism is explained in the following figures.

The basic operation of buffer management is shown in **Figure 32**. When the recipient, here as a responder receives the aggregated frame, the first MPDU is received correctly. Therefore the Starting Sequence number will be updated at the recipient. By the update of the Starting Sequence number, all the complete MSDUs with lower sequence numbers than the Starting Sequence number and also the complete sequential MSDUs after the Starting Sequence number will be indicated to the higher layer as in 802.11e. All the buffers held by preceding MPDUs are released.

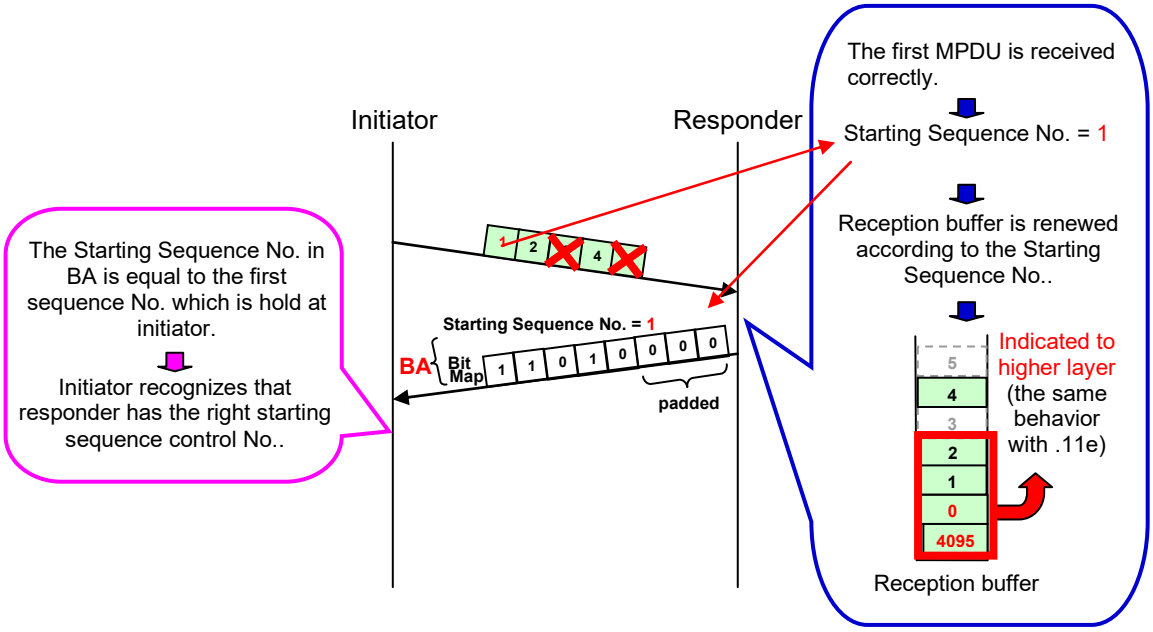


Figure 32 - Basic buffer management in Implicit BAR

When the first MPDU is not correctly received, the buffer management will operate as Figure 33 or Figure 34. Since the first MPDU has an error, the recipient cannot acquire the correct Starting Sequence number from the aggregation. The recipient uses sequence number of the first correct MPDU to generate the BA frame. In this case, completed MPDUs are not indicated to the higher layer and buffers are not released.

When the BA frame is received at the originator, it compares the Starting Sequence number indicated in the BA frame with the first sequence number it is holding, i.e., the real Starting Sequence number. If the values are not equal, then the originator recognizes that the recipient has failed to acquire the Starting Sequence number from the aggregation. To inform the recipient of the correct Starting Sequence number, the initiator may transmit a BAR frame as in Figure 33. Or the retransmission of the failed MPDUs in the next aggregation includes the retransmission of the right Starting Sequence number. This is shown in Figure 34. Additional MPDUs may be included in the next aggregation.

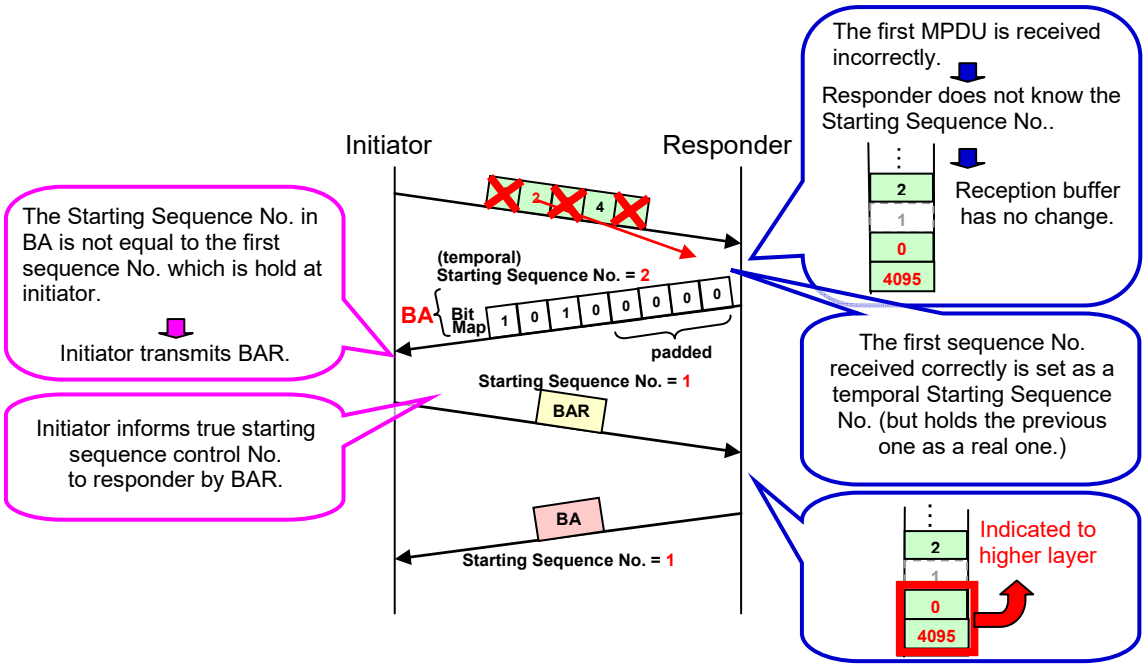


Figure 33 - Buffer management when the first MPDU is incorrect in Implicit BAR (Case 1)

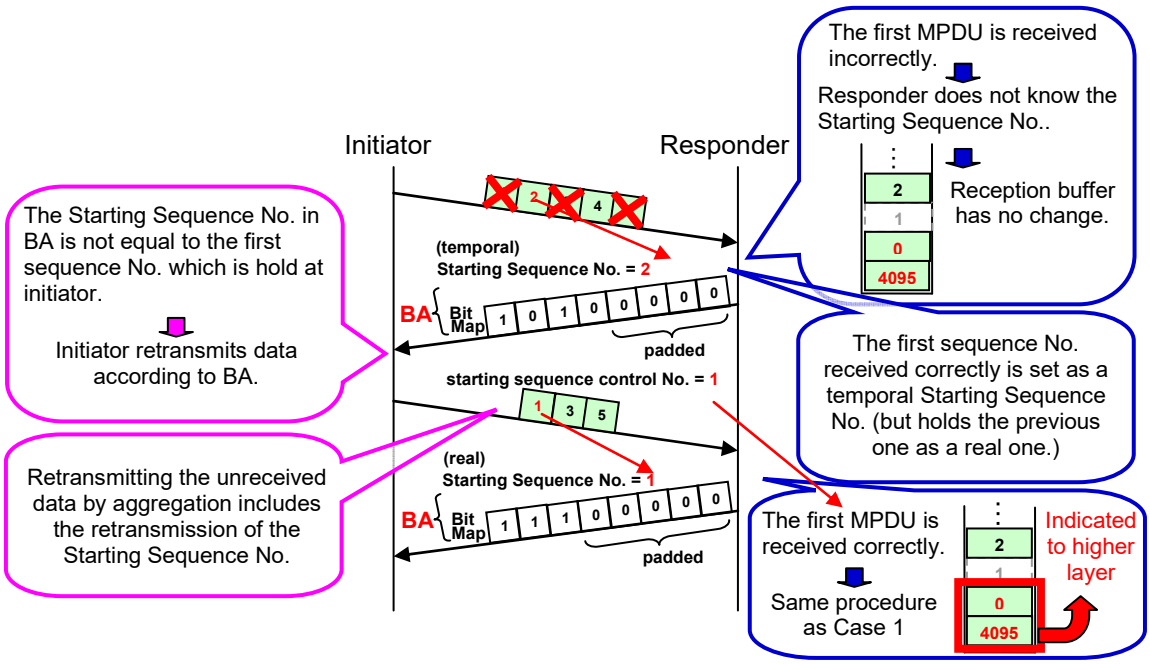


Figure 34 - Buffer management when the first MPDU is incorrect in Implicit BAR (Case 2)

When the originator does not receive the BA frame after the transmission of the aggregated MPDU, it shall transmit a BAR frame and obtain a BA response before transmitting any more data frames using this BA agreement. The recipient then transmits a BA frame according to the Starting Sequence number indicated by the BAR frame and the buffer status.

When all the MPDUs in the aggregate fail at the recipient, the recipient does not send a BA frame. From the originator's point of view, this is indistinguishable from the lost BA case described above, and is handled in the same way.

7.5.3 Maintaining Acknowledgement State for a BA agreement

The recipient shall maintain the acknowledgement bitmap of size defined by the combination of Non-Frag and NumMSDU_M1 defined in section 6.1.6.2 for each BA agreement.

This bitmap for the matching agreement shall be reset whenever a first data frame for that agreement is received in an aggregate. A STA retains this acknowledgement state in case its BA is not received by the originator, and the originator has to retry. However, once the originator has successfully received the BA, it may transmit data in an aggregate, and the receipt of a new aggregate indicates to the recipient that the previous acknowledgement state is no longer required for the matching TID.

7.5.4 Compressed BA

A compressed variant of the 802.11e legacy BA frame may be used according to parameters negotiated during the Block Ack setup.

The Non-Frag and NumMSDU_M1 parameters define the use of fragmentation and the maximum number of MSDUs acknowledged in the BA bitmap.
The compressed BA frame format is defined in section 6.1.6.

A STA transmitting an ADDBA MPDU shall first determine if its peer is a HT STA from an exchange of capabilities and shall not set the Non-Frag field to 1 if its peer is not a HT STA.

A HT STA shall always transmit a compressed format BA MPDU to a HT STA. A HT STA shall never transmit a compressed format BA MPDU to a non-HT STA.

The maximum value of the NumMSDU_M1 subfield is determined from the BA agreement as defined in 6.3.5. The number of MSDUs whose acknowledgment status is reported in the current BA MPDU is equal to (NumMSDU_M1+1).

The value of the Non-Frag subfield is shall match the BA agreement. This same value is present in all BA MPDUs transmitted for this agreement.

The length of the BA bitmap is determined from the ADDBA Buffer Size subfield and the Non-Frag subfield as defined in 6.3.5, and is constant for the lifetime of the BA agreement.

An HT STA shall not transmit fragmented MSDUs on a BA agreement for which Non-Frag has the value 1.

8 MAC sublayer management

8.1 Coexistence management

8.1.1 Operating Modes

The operating modes of an HT BSS are defined in the following table. These modes determine the format of beacon transmission and the kind of protection used for HT transmissions.

Topology	BSS Mode Name	Definition of Mode
Infrastructure	Pure	No co-channel legacy devices are present
	Mixed Capable	No co-channel legacy devices are present, but the HT AP is able to accept association from a legacy device that discovers this BSS (e.g. through receiving its legacy beacon) and decides to associate.
	Managed Mixed	Legacy and HT devices are present co-channel. Access to the medium is managed by an AP. Note: In this mode, time is divided between contention periods for HT and legacy stations (by selectively setting the NAV)
	Unmanaged Mixed	BSS contains both legacy and HT devices or there are co-channel legacy devices of a different BSS present
	20 MHz-Base Managed Mixed	BSS contains both legacy and HT devices. There may be legacy devices of overlapping BSS in either or both of the channels. Legacy and HT devices associate to the AP's BSS in the control channel. The AP manages the generation of 40 MHz mode periods, during which 40 MHz HT devices are allowed to access in 40 MHz.
Independent (ad-hoc)	Discoverable IBSS	No co-channel legacy devices are present. The IBSS may be discovered by a legacy device.
	Mixed IBSS	There are legacy devices co-channel

8.1.2 Receiver capabilities wrt legacy operation in the extension channel

An HT AP or IBSS STA should be able to detect activity in the extension channel without being able to receive legacy MPDUs in it. On detecting a possible legacy MPDU, the STA should enable itself to receive in that extension channel in order to determine if a legacy BSS has arrived on the extension channel. It is not necessary that this capability be available simultaneous with operation on the control channel.

8.1.3 Infrastructure BSS Operation

8.1.3.1 At the AP

Table 14 – AP Mode Definitions

AP Mode	Beacon Transmission	Indicated STA operation
Pure	Pure	Pure
Mixed Capable	Legacy	Mixed
Managed Mixed	Dual (see below)	Pure
Unmanaged Mixed	Legacy	Mixed
20 MHz-Base Managed Mixed	Legacy mandatory in the control channel, optional in the extension channel	Mixed

The AP controls the operating mode of its STA through an information element in the beacon. In pure mode, the AP will ignore any probe requests sent by legacy STAs and any legacy beacons sent by overlapping co-channel legacy devices. The AP transmits its beacon using a HT PPDU type. The beacon contains an HT management element that requires pure mode operation of its STAs.

> Note, use of pure mode at the AP is only suitable for managed installations where it is not permitted that legacy APs may share the same channels as HT APs.

In the mixed-capable mode, the AP transmits its beacon using a legacy PPDU. The beacons contain an HT management element that indicates this is an HT AP [4].
In a mixed-capable BSS, an HT STA shall operate in pure mode except when communicating with a legacy STA using DLP. An HT STA shall switch to legacy mode for communication with a legacy STA using DLP.
In the mixed-capable mode, the AP may receive association and probe requests from legacy STAs. It shall respond to a legacy probe request with a legacy probe response. It shall respond to a legacy association request with a legacy association response. The AP may choose to accept or deny the association request. If the AP sends an association response with status = OK, it shall enter either managed, unmanaged or 20 MHz-base managed mixed mode operation, and stay there while it has any associated legacy STA. The AP monitors its channel for a legacy co-channel BSS, and if it detects one, transitions into mixed mode, or attempts to find an alternative channel using the procedures defined in section 8.2.
In the unmanaged mixed mode, the AP transmits its beacon using the legacy PPDU type. The beacon contains an HT management element that requires mixed mode operation of its STAs. STAs in an unmanaged mixed mode BSS may use legacy or HT transmissions.

[4] This will stop other co-channel mixed-capable APs from considering the AP to be a legacy AP.

In the 20 MHz-base managed mixed mode, the AP transmits its beacon using the legacy PPDU type. The beacon is transmitted on the control channel but for the purpose of reserving the extension channel, an AP may send beacons also in the extension channel. Beacons in the extension channel may cause legacy STAs to attempt association but the HT AP shall deny association or may ignore those requests. The beacon contains an HT management element that requires mixed mode operation of its STAs.

8.1.3.2 At the STA

An infrastructure HT STA supports three possible modes of operation: legacy, mixed and pure. [5]
In pure mode, there are no overlapping legacy STAs. Protection of HT frames from legacy devices is not required.
In mixed mode, the HT STA is operating in the presence of legacy STAs co-channel on the control channel and/or the extension channel These STAs may be part of the same BSS, or may be associated with an overlapping legacy BSS. In mixed mode, use of a legacy protection mechanism is required (see section 7.1.6). An HT STA shall operate in pure mode except for communicating with a legacy STA on a Direct Link. An HT STA shall switch to legacy mode for communication with a legacy STA on a Direct Link. A 40 MHz capable HT STA, whose permitted width set is 20 and 40 MHz, shall switch to 20 MHz mode for communication with a 20 MHz HT STA.[6]
In 20 MHz-base managed mixed mode, a 40 MHz capable HT STA is enabled to communicate in 40 MHz mode during the 40 MHz period. In case a 40 MHz capable HT STA wishes to communicate with a 20 MHz mode HT STA or with a legacy STA, it shall do so in the 20 MHz period.
In legacy mode, the HT STA operates exactly as a legacy device, with the exception that it may switch from legacy mode to mixed or pure modes while scanning in order to detect an HT AP.

8.1.3.3 Managed Mixed Mode Operation (Informative)

This section describes a coexistence mechanism that permits a legacy network and an 802.11 HT network to operate co-channel without requiring the 802.11 HT devices to be able to communicate directly with the legacy devices.
The mode is operated by an HT AP in managed mixed mode that has both legacy and HT STAs associated with it. Its HT STAs operate in pure mode.
In this mode, the AP manages coexistence of legacy devices in its control channel with HT devices occupying the HT channel, which may be 20 or 40MHz wide depending on the operating mode of the AP.
The case of a 40MHz channel where both 20MHz halves are overlapped by legacy devices is managed using the 20 MHz-base managed mixed mode which is described in section 8.1.3.4.

8.1.3.3.1 Assumptions

It is assumed that the legacy devices cannot receive HT PPDUs and cannot interpret MAC duration values. The managed mode avoids the need for individual STA to provide legacy protection, and thereby potentially reduces overhead.

[5] The STA cannot distinguish between the AP's managed mixed mode and pure mode.

[6] A 20 MHz HT STA is the one whose permitted width set is limited to 20 MHz.

8.1.3.3.2 Mechanism

The AP provides managed mixed mode operation.
This AP divides time between legacy and HT phases of operation. In the legacy phase it ensures that the NAV of all HT devices is set, and *vice versa*.
One way to do this is for the AP to periodically transmit a legacy CTS-to-self to start the HT phase and an HT CTS to start the legacy phase. Alternatively, the AP may combine setting the NAV with transmission of a beacon that contains a contention-free period element, and transmission of a matching CF-end at the end of the phase.
The legacy beacon transmitted by an HT AP in managed mixed mode contains an element which causes any HT STA scanning for a network to ignore the duplicate legacy beacon.
HT devices in pure mode should ignore legacy PPDUs they may receive, so they do not receive NAV reservations intended for legacy devices.

8.1.3.4 20 MHz-Base Managed Mixed Mode Operation

This section describes a coexistence mechanism where a BSS operates in 20 MHz mode and switches to a 40 MHz capable phase under the control of an HT AP. Both legacy and HT STAs may be associated in the BSS. An HT STA may be either a 40 MHz capable HT STA or a 20 MHz HT STA. STAs associate with the HT AP on the control channel.
The 20 MHz-base managed mixed mode allows overlapping BSS containing legacy STAs in both control and extension channels. Due to the protection of the 40 MHz period in both channels, it is tolerant of overlapping BSSs. If the overlapping BSS is a QBSS, access to the channel may experience longer latencies in the order of the TXOP duration of the QBSS. Similarly, channel access delays will be long when operating with an 802.11b BSS, which may use substantially lower data rates.
The AP, before it selects a control channel and extension channel and operating mode may take into account the likely performance benefits of using the 20MHz base mode operation. The AP should try to avoid overlapping 802.11e or 802.11b STA, particularly on the extension channel.

8.1.3.4.1 Assumptions

Legacy STAs and HT STAs may coexist in the 20 MHz-base managed mixed mode. Two types of STAs may coexist with the 40 MHz capable HT STAs in the 20 MHz-base managed mixed mode: legacy STAs and 20 MHz HT STAs. It is assumed that the legacy STAs cannot receive HT PPDUs and cannot interpret MAC duration values. The 20 MHz HT STAs are assumed not to be able to receive 40 MHz mode HT PPDUs and to interpret their MAC duration values. An HT AP in the 20 MHz-base managed mode provides legacy protection against legacy STAs and 20 MHz HT STAs.

8.1.3.4.2 Mechanism

The 40 MHz capable HT AP may operate in the 20 MHz-base managed mixed mode. It divides time into 20 MHz and 40 MHz periods. In a 20 MHz period, it ensures that the NAV of 40 MHz mode operation in 40 MHz capable HT STAs is set. In the 40 MHz period, the NAV of legacy STAs and 20 MHz HT STAs is set.

The basic period is the one in which operation is strictly in the 20 MHz control channel. To start a 40 MHz period, the HT AP first reserves the control channel by setting NAV of legacy and 20 MHz HT STAs with a legacy Beacon frame or an ICB frame. The transmission rate of the SCB frame shall be selected from the BSSBasicRateSet. Due to the range of the Duration/ID field in the MAC header, the ICB frame cannot be used to start a 40 MHz period longer than 32767 μs.

Then, the HT AP shifts to the extension channel for its reservation. The extension channel is reserved by a transmission of a CTS-to-self or a legacy Beacon frame after an appropriate channel access is performed in the extension channel.

The Beacon frame sent to reserve the control channel includes the channel extension indication information element. When a 40 MHz capable HT STA associated with the HT AP receives the Beacon frame, it extends its channel bandwidth to 40 MHz according to the information in the Extension Channel Offset information element. The 40 MHz capable HT STAs act in the same way when receiving an ICB control frame.

After setting the NAV in the extension channel is completed, the HT AP resets the NAV of 40 MHz capable HT STAs by sending a CF-end frame. Thereby the 40 MHz period is started, during which HT STAs can communicate in 40 MHz mode using EDCA.

The AP can also support operation using HCCA by polling HT STA during the 40 MHz period before allowing the 40MHz devices to contend using EDCA.

To end the 40 MHz period, the HT AP first sets the NAV of 40 MHz mode operation in 40 MHz capable HT STAs by transmitting an DCB frame. The 40 MHz capable HT STAs will switch back to the control channel in 20 MHz. Then the AP resets the NAV in the extension channel by transmitting a CF-End frame. The AP may reset the NAV in the control channel by a CF-End frame but it may also continue the CFP that was set in the last Beacon frame on the control channel. At this point the AP and all its STAs are operating on the control channel using 20MHz channel width. The process may e.g. be repeated periodically, such as related to the beacon interval. The superframe thus created is divided into a phase for communication on the 40 MHz channel and phases for communication on the 20 MHz channels. One cycle of the process is illustrated in Figure 35.

A 40 MHz capable HT STA that is attempting an EDCA channel access during a 40MHz period either freezes the backoff counter during 20MHz operation and resumes the interrupted backoff during the next 40MHz period or selects a new random backoff at the start of the 40MHz period. Likewise, if it is attempting a channel access during a 20MHz period, it either freezes the backoff counter during 40MHz operation and resumes the interrupted backoff during the next 20MHz period or it selects a new random backoff at the start of the 20MHz period.

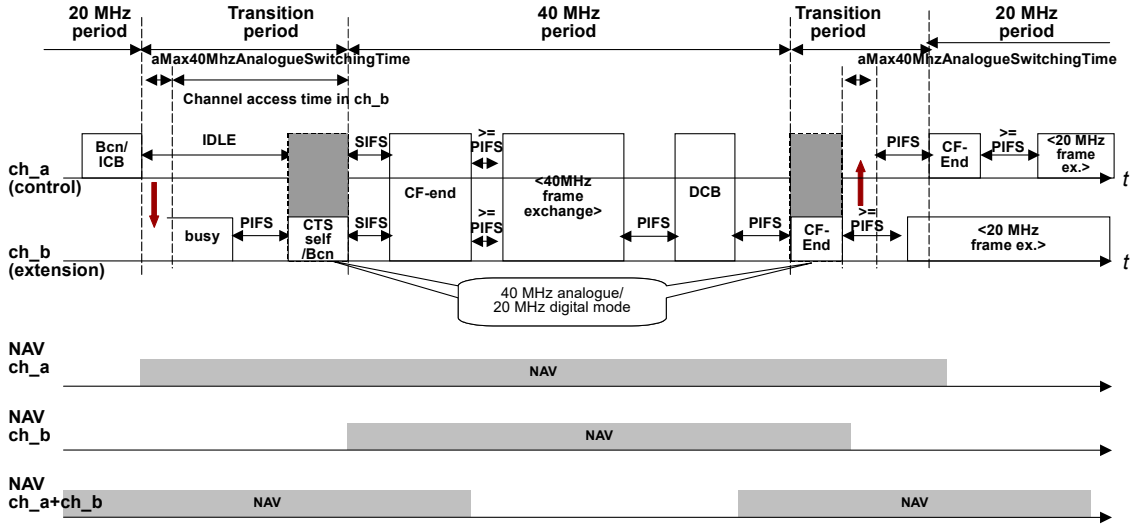


Figure 35 – 20 MHz-Base Managed Mixed Mode

8.1.3.4.3 Operation at HT AP

To start the 40 MHz period, the HT AP transmits a legacy Beacon frame or an ICB frame in the control channel to acquire the control channel and define a CAP to block the channel by setting the NAV of the legacy STAs and 20 MHz HT STAs on the control channel. The CFP of the Beacon or Duration field of the ICB frame will be set to cover the 40 MHz phase plus the transition periods between 20 and 40 MHz operation.

The HT AP announces the "extension channel access timeout" value in its Beacon and Probe Response frames to limit the maximum transition time to the 40 MHz period. When the HT AP transmits the Beacon frame or ICB frame on the control channel, it starts a timer of a duration, which is "Extension Channel Access Timeout" minus the duration of the Beacon or ICB frame. The "Extension Channel Access Timeout" is the maximum time, after which the STAs shall have received the Beacon or CTS-to-self on the extension channel. One reason why the AP might not be able to send the Beacon or CTS-to-self frame within "Extension Channel Access Timeout" time could be a busy medium on the extension channel.

After setting the NAV in the control channel, the AP may switch to 40 MHz analogue and 20 MHz digital mode and shall listen to both the control channel and the extension channel. As the control channel is supposed to be reserved by the previous operation, this phase is mainly given to wait for the extension channel to be idle. However, it shall be noted that while waiting for the extension channel to become idle, the control channel will be left without any activity and the STAs which did not receive the Beacon or ICB frame in the control channel may interfere with the reservation. The purpose of the mode transition to 40 MHz analogue and 20 MHz digital is to omit an analogue channel switch at the beginning of the transmission phase in the 40 MHz channel. If the extension channel becomes idle the AP shall transmit a CTS-to-self or legacy Beacon in the extension channel a PIFS after it has become idle. The CTS-to-self or Beacon shall block the channel by setting the NAV of the legacy STAs on the extension channel. The NAV shall cover the intended duration of the 40 MHz period and the transition times between 20 MHz period and 40 MHz period.

If the "extension channel access timeout" timer at the AP expires while attempting to transmit CTS-to-self or Beacon, the AP shall give up switching to the 40 MHz bandwidth and may try

again at a later time. The AP shall furthermore transmit a CF-end frame on the 20 MHz control channel in order to re-set the NAV of the legacy STAs on the control channel. HT STAs will have started an "Extension Channel Access Timeout" timer themselves after having received the first Beacon or ICB frame on the control channel and do therefore not have to be notified by the AP about the expiration of the timer.
The NAV setting in the extension channel is done through CTS-to-self or Beacon. CTS-to-self would give less overhead, however, the Beacon in the extension channel may avoid other BSSs being created, since other STAs will detect the presence of the HT BSS. Furthermore, the duration of the NAV that can be signaled by a CTS-to-self is limited. Therefore, for long periods setting the NAV by a Beacon frame will be required. According to this analysis, implementers may decide to send Beacons in the extension channel.
If the "Extension Channel Access Timeout" timer has not yet expired, the AP shall transmit a 40 MHz mode CF-end to signal to the HT STAs that the 40 MHz channel is available. The CF-end frame shall be transmitted at least aMax40MhzAnalogueSwitchingTime after the end of the first Beacon or ICB frame on the control channel. aMax40MhzAnalogueSwitchingTime is the maximum allowed time for the STAs and the AP to carry out an analogue channel switch from 40 MHz to 20 MHz and vice versa. The AP has to wait at least aMax40MhzAnalogueSwitchingTime before starting the 40 MHz phase in order to account for the STA with the slowest possible switching time. If aMax40MhzAnalogueSwitchingTime time has already expired since the end of the Beacon or ICB frame on the control channel, the AP shall send the CF-end or QoS CF-Poll frame a SIFS after the CTS-to-self or Beacon on the extension channel.
After the 40 MHz frame exchanges, the AP shall set the NAV of the 40 MHz HT STAs until the intended end of the 20 MHz period by means of a DCB frame. Then the AP shall free the extension channel and the control channel for communication in 20 MHz mode by transmitting CF-End frames in both channels. The CF-End frame in the control channel is not required if the AP wishes to continue the CFP. The CF-End frame on the control channel shall not be sent earlier than aMax40MhzAnalogueSwitchingTime after the CF-End on the extension channel. The first CF-End frame in the extension channel may be transmitted in 40 MHz analogue and 20 MHz digital mode in order to avoid an additional analogue channel switch. After the first CF-End frame the AP switches to the control channel in 20 MHz analogue mode and transmits the second CF-End frame.

The ratio between the 40 MHz and 20 MHz period should be adjusted according to types of traffic and priority. Whether frames are sent in the 40 MHz or 20 MHz period shall be scheduled depending on their traffic types.
Note that in scenarios with heavy interference between control and extension channel the AP may choose to switch to a different 40 MHz channel, on which no time-sharing with legacy STAs might be required. The selection of channels affects the performance and efficiency for the 20 MHz-base managed mixed mode. Not only the initial channel selection but also monitoring the channels while operating in 20MHz-base mixed mode is necessary to cope with condition changes.

8.1.3.4.4 Operation at 40 MHz capable HT STAs

A 40 MHz capable HT STA is allowed to operate in 20 MHz mode on the control channel or in 40 MHz mode on both channels. It is not allowed to operate in 20 MHz mode on the extension channel.
40 MHz capable HT STAs shall store the "extension channel access timeout" value contained in Beacon or Probe Response frames sent by an HT AP.

When a HT STA receives a Beacon or ICB frame with channel extension indication information element set, it shall start an associated timer of duration "Extension Channel Access Timeout" and wait in the 40 MHz analogue mode for the HT AP to reset its NAV by a CF-End frame. Furthermore, upon reception of the Beacon or ICB frame in the control channel, a 40 MHz capable HT STA shall start the timer included in the Duration/ID field of the ICB frame or the CFP Parameter Set element in the Beacon frame. If the HT STA operates in the 20 MHz mode, it shall shift to 40 MHz analogue mode by the reception of the Beacon or the ICB frame on the control channel. A HT STA shall switch to 40 MHz analogue mode within at least aMax40MhzAnalogueSwitchingTime time.
In case the "extension channel access timeout" timer expires before receiving the CTS-to-self or Beacon on the extension channel, a 40 MHz capable HT STA leaves the waiting state. It shall switch back to 20 MHz mode on the control channel depending on the operation mode of the AP.

When a 40 MHz capable HT STA receives a CF-end frame, its NAV for the 40 MHz channel will be reset and it becomes free to access in 40 MHz mode. The 40 MHz capable HT STA switches back to 20 MHz mode on the control channel if the timer runs out before it receives the DCB frame which indicates the end of the 40 MHz period.
The HT STA may send frames not only in EDCA but also in HCCA, which is enabled by QoS CF-Poll to it. A 40 MH capable HT STA which is granted TXOP may only transmit frames to the AP or other 40 MHz capable HT STAs.
By the reception of a DCB frame from the AP on the 40 MHz channel, a 40 MHz capable HT STA sets its NAV for the 40 MHz channel. It may switch back to 20 MHz mode if it wishes to communicate in the 20 MHz period when receiving the DCB frame.

A 40 MHz capable HT STA that is attempting an EDCA channel access during a 40MHz period freezes the backoff counter during 20MHz operation and resumes the interrupted backoff during the next 40MHz period. Likewise, if it is attempting a channel access during a 20MHz period, it freezes the backoff counter during 40MHz operation and resumes the interrupted backoff during the next 20MHz period.

8.1.3.4.5 Operation at legacy STAs and 20 MHz HT STAs

The NAVs of Legacy STAs and 20 MHz HT STAs are set either by a Beacon or ICB frame in the control channel or by the CTS-to-self or Beacon frame in the extension channel. Their NAVs are reset when they receive a CF-End frame in their operating channel.

8.1.4 Ad-hoc (IBSS) Operation

8.1.4.1 Effect of operating mode as an HT IBSS (Informative)

In IBSS operation, an HT IBSS is always discoverable by the legacy IBSS. A legacy device wishing to join the IBSS will discover it and share in the beaconing. This will alert HT devices to its presence based on omission of HT beacon elements.

8.1.4.2 Rules of operation of an HT IBSS

Each STA maintains a table for STAs it can see according to the mode of operation of the STA. Stale entries in this table are removed by a timeout process. A STA can be classified as legacy if it is seen to transmit a legacy beacon without the HT elements (note, this also includes legacy

APs). Otherwise, its HT IBSS operating parameters are discovered from the HT element in the beacon. This element includes the field: IBSS mode.
The mode of operation of an individual STA is defined in the following table, along with the rules for operation.

Table 15 – IBSS Mode Definitions

HT IBSS Mode	Description	Condition	Beacon Tx
Discoverable IBSS	Only HT IBSS on this channel	No legacy devices in STA table	Legacy format
Mixed IBSS	Mixture of HT IBSS and	One or more legacy devices are present in STA table.	Legacy format

8.1.4.3 Rules for IBSS Use of Protection Mechanism

When operating in Mixed IBSS mode, an HT IBSS STA acting as the initiator of a frame exchange sequence shall use a legacy protection mechanism.
Otherwise, the STA should use HT frame formats.

8.2 Channel management

8.2.1 Channel Management at the AP

The AP transmits a beacon containing the following elements:

Table 16 – Beacon Elements for HT Channel Management

Element	Description	Values
Permitted Width Set	This is the set of channel widths that may be used within the BSS.	{20} or {20, 40}
Control Channel Number	Channel number for the control channel	
Extension Channel Offset	Describes the location of the extension channel relative to the control channel.	1 = 1 channel above, –1 = 1 channel below, 0 = not present

An HT AP shall receive the beacons of other HT BSSs on its channel. If one of these indicates a subset of its current permitted width set, the AP shall reduce its permitted width set so that it matches or shall select a new operating channel.

HT BSSs shall only overlap HT BSSs with the same control channel and extension channel location. [7]

8.2.2 Channel Management at the STA

All transmissions must use a width in the permitted width set.
In unmanaged mixed mode, all transmissions of unaggregated control and broadcast frames must be receivable by legacy devices operating on the control channel.
A STA shall only transmit an HT PPDU to a peer STA using a width that it knows the peer to support and that is within the permitted width set.

8.2.3 Channel Selection Methods

8.2.3.1 Introduction (Informative)

The best coexistence is achieved when all BSSs occupy separate channels and when legacy and HT devices do not overlap.
The HT AP should prefer to operate a pure mode BSS (i.e. no overlapping legacy BSS) because overlapping BSS requires a mixed mode operation that may entail higher overhead.
The mixed-capable AP can detect co-channel legacy BSSs arriving on its control channel without the need for scanning. This may cause it to enter a mixed mode of operation.
The AP may also detect co-channel legacy BSS arriving on a non-control channel. It may do this by scanning during an 802.11h "quiet" period, or by getting one of its STAs to scan for it, again using 802.11h signaling [8]. Alternatively, the AP may be able to receive legacy MPDUs on either overlapped channel.
On detecting the arrival of a co-channel legacy BSS on the non-control channel, the AP has to either adjust its permitted channel widths to avoid the legacy BSS, select a new set of channel parameters or switch to 20 MHz Base Managed Mixed Mode.
The AP may also scan periodically all channels, and evaluate whether there is a better channel (as defined in **Error! Reference source not found.**) to operate on. If it discovers a better channel, it may instruct its STAs to change channel parameters using signaling in its beacon.
Note: If 802.11h-like DFS becomes mandatory in all geographies, this behavior should be consistent with those requirements. If DFS is not mandatory, operation of this automatic channel selection may be disabled to permit fully manually administered installations.

8.2.3.2 Rules

The HT AP shall either be configured manually or configure itself automatically as described in this section.
The HT AP shall scan its environment before selecting channel parameters. If the AP is capable of mixed mode operation, it shall use mixed mode to scan.
If the selected channel parameters overlap an existing HT BSS, then they must use the same control channel and extension channel offset and permitted width set.

[7] This, together with the requirement for the same basic width set ensures that any pure-mode control frames sent in one BSS can be received in any overlapping HT BSS.

[8] Note, the assumption that 802.11h will be used (perhaps in modified form) by TGn.

The AP shall reselect new channel parameters if an HT BSS that does not have the same control channel, extension channel offset and permitted width set starts operating on an overlapped channel.
Given a choice of 40MHz channel selections permitted by the previous rules, the AP should prefer a channel selection that aligns its 40MHz channel with the starting (i.e. lowest frequency) pair of adjacent 20MHz channels within any band or sub-band plus zero or more times 40MHz.

Note: The purpose of this recommendation is to avoid introducing 20MHz "holes" between adjacent operating 40MHz BSS, thus making them unavailable for 40MHz operation.

8.3 TRMS Power-Saving Management

MIMO systems can reduce power consumption by only enabling a single receive chain during periods of inactivity. The STA may be raised to full receive capability following the reception of a unicast MPDU contained in a PPDU that can be received on a single receive chain.
This section defines the TRMS protocol. The mechanism is referred to as Timed Receive Mode Switching (TRMS) since it relies on a timer in the STA to trigger the transition to reduced receive mode and corresponding timers in the AP and STA peers to track the state change. The timers are reset when an MPDU is transmitted by the TRMS STA.
As a special case, an operating mode is supported where the STA reduces receive capability after each sequence. This mode could take advantage of any protection mode where a SISO frame exchange already prefixes each sequence by providing power savings at no protocol overhead.
A STA shall support TRMS operation in other STAs with a zero hold-on time. It is an option of a STA to support TRMS operation in other STAs with a non-zero hold-on time.
When a STA that does not support TRMS operation in other STAs with a non-zero hold-on time transmits to a STA that is in TRMS enabled mode, it will assume that the STA uses a zero hold-on time. This will cause it to prefix any exchange to that STA with a SISO packet. This may or may not (depending on the exchange sequence used and the type of protection used) create a performance reduction compared to support of a non-zero hold-on timer.

8.3.1 BSS Association, DLP Setup and Mode Changes

8.3.1.1 Association

When associating, a STA that wishes to operate using TRMS includes the TRMS IE in its association or reassociation request.
If the STA wishes to update the TRMS parameters under which it is operating after association, then it sends a directed TRMS Update frame to the AP with the new parameters.

8.3.1.2 DLP Setup

A STA that supports TRMS operation in its DLP peers includes a HT capability element with an asserted TRMS Support bit in the DLP request or response frame. A STA initiating a DLP session that wishes to operate using TRMS includes the TRMS IE in the DLP request frame.
A STA that receives a DLP request with the TRMS Support bit set and that wishes to operate using TRMS includes a TRMS IE in the DLP response.

If a new DLP session results in the STA changing TRMS operating parameters it should send a TRMS Update to the AP and its other DLP peers.

8.3.1.3 Mode Changes

A STA may change its TRMS operating parameters. When doing this it sends a directed TRMS Update to the AP and each of its DLP peers. The STA should send the updates before the new parameters take effect if the parameters decrease the hold-on time or enable TRMS.
A STA may activate or deactivate TRMS operation at any time.

8.3.1.4 IBSS Operation

A STA that wishes to operate using TRMS broadcasts a TRMS Update frame and then enables TRMS operation. Once enabled, the STA should include the TRMS IE in future beacons.
Like IBSS power-saving, a STA only has potentially stale information about a peer's TRMS mode. A STA can refresh this information by performing a capability exchange using directed Probe MPDUs. It may infer that its knowledge is stale if it fails to transmit to a STA on the assumption that is is TRMS-disabled, but successfully transmits when it assumes it to be TRMS-enabled.

8.3.2 TRMS Operation

A STA using TRMS maintains a running timer that is reset each time a frame is transmitted. If the timer reaches the hold-on time then the STA switches to reduced receive operation unless a frame is being received, in which case the switch will be delayed until immediately after the frame has been received. The STA returns to full receive operation after transmitting any frame.

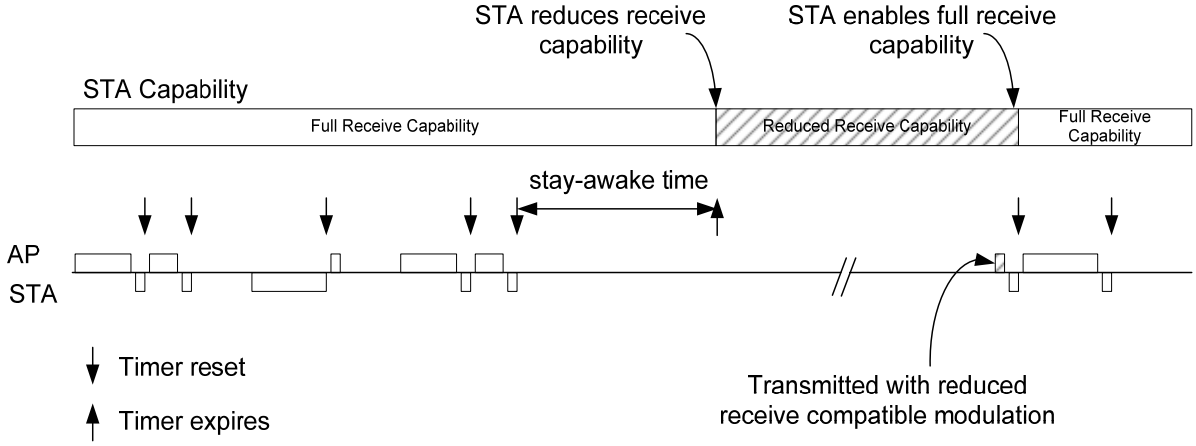


Figure 36 – Data transfer example with timed receive mode switching

An AP maintains a running timer for each associated STA that is operating using TRMS. The AP resets the timer for a particular TRMS using STA each time it receives a frame from that STA (whether or not the frame is addressed to the AP). If the timer reaches the hold-on time then the TRMS STA is assumed to have switched to reduced receive operation.
A STA that is a DLP or IBSS peer to other STAs using TRMS maintains a running timer for each such peer. Each time this STA receives a frame that is addressed to it and that was transmitted by such a peer, it resets the corresponding timer. It may also reset the timer if it receives a frame from such a peer that is not addressed to it. If the timer reaches the hold-on time then the peer TRMS STA is assumed to have switched to reduced receive mode.

Frames sent to a TRMS STA that is assumed to be operating with reduced receive capability shall be transmitted using a modulation compatible with the TRMS STA's reduced receive capability.

8.3.3 Performance Impact of using TRMS mode (Informative)

In TRMS mode with a non-zero timer value, this value can adjust the compromise between staying MIMO-enabled longer than necessary and the overhead of switching into a MIMO-enabled state. This is only relevant for bursty traffic. For periodic traffic where there is a single exchange sequence per period, there is no benefit in using a non-zero value of the timer. The overhead of entering a MIMO_enabled state depends on the details of the frame exchange sequence and protection used.

Sequence Type	Overhead
LongNAV IAC+RAC+Data/Data/BA+RAC/BA	In the case of a basic sequence, this overhead is zero because the initial packet is already a SISO packet.
LongNAV ABF (forward)	The overhead is still zero
LongNAV with legacy devices	The overhead is zero because the legacy initial MPDUs are suitable SISO packets
LongNAV ABF (both directions)	The ABF sounding for the reverse-link cannot occur in the first packet, so ABF trained reverse-link data can first be sent in the 4th PPDU. However this creates no additional overhead, because the operation of the RDL/RDR/RDG protocol does not permit reverse-direction data until the 4th PPDU, except in the case of periodic RDR. Periodic RDR is there to support applications like VoIP, and I don't expect to see VoIP handsets doing ABF transmissions.
Pairwise Spoofing Untrained IAC/Data/Data/BAR+RAC/BA	This sequence will not work with a MIMO MCS for the first MPDU. It makes little sense to transmit a SISO MCS for the first packet. Also there are practical difficulties knowing what MCS to use because any training established when the receiver is MIMO-enabled will not apply when it is MIMO-disabled. The alternative is to add an initial IAC/RAC exchange. This brings the following advantages: • Collisions are detected cheaply • Data is protected from hidden nodes • Data can be trained (fresh MCS and/or/

	ABF) These benefits are not useful if there are small numbers of contending stations, there are no hidden nodes, and training information is still fresh. In that case the additional IAC/RAC is a real overhead.
Pairwise Spoofing Trained forward-direction data	No overhead as the sequence has to be: IAC(TRQ)+RAC(HT-LTFs)+IAC/Data/… The initial IAC can be SISO.
Pairwise Spoofing reversed direction trained data	Same analysis as the LongNAV case. There is no additional overhead.

8.3.4 Calibration

8.3.4.1 Introduction

Differences in transmit and receive chains in a STA or AP destroy the inherent reciprocity of the over-the-air TDD channel. Calibration is necessary in order to remove differences in transmit and receive chains and enforce reciprocity in the observed baseband-to-baseband channels. A simple over-the-air calibration procedure effectively corrects for these differences. The calibration procedure ensures that the observed channel matrices in the two directions of the link are transposes of each other and thus renders the resultant channel reciprocal. Thus, if they are able to do so, STAs calibrate upon association. Calibration works for any square or non-square dimensionality, i.e., N_{Tx} and N_{Rx} greater than one and less than five.

8.3.4.2 Procedure

The Calibration Training frame exchange sequence, illustrated in Figure 37 operates as follows:

- STA A transmits a Calibration Training MPDU (Position field = 0).
- In response, STA B transmits a Calibration Training MPDU (Position field = 1) after a SIFS interval, which includes the HT-LTF that allows STA A to estimate the channel matrices.
- In response, STA A transmits a Calibration Training MPDU (Position field = 2) after a SIFS interval, which includes the HT-LTF that allows a receiving STA B to estimate the channel matrices.

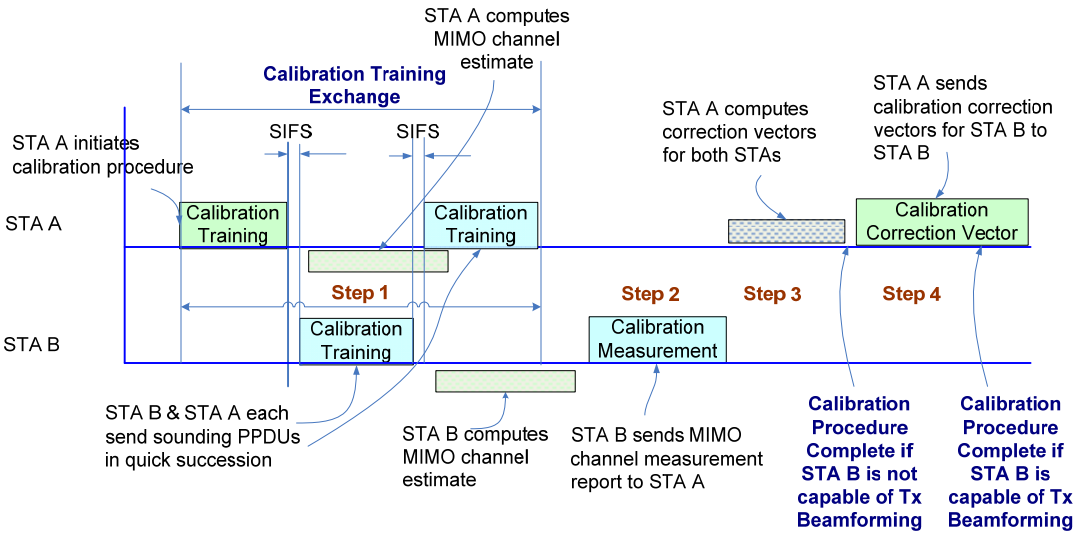


Figure 37 - Calibration Procedure

The Calibration Training MPDUs are carried in non-aggregated sounding PPDUs. The last two steps must occur over a sufficiently short time interval so that the channel does not change over the interval. When used in a mixed BSS, this frame exchange sequence must be protected with an appropriate protection mechanism.

The remaining message exchange in the calibration procedure is not time critical.
When the channel measurements become available at STA B, STA B sends one or more Calibration Measurement action frames that contains the MIMO Channel Measurement Report. STA B sets the Calibration Complete bit to '1' in case it is not capable of transmit beamforming. In this case STA A will not send Calibration Correction frame.

In case that both STAs are capable of transmit beamforming, once STA A has the MIMO Channel Measurement Report from STA B, along with its local channel measurement measured from the Calibration Training frame transmitted by STA B, STA A can compute the correction vectors required for both itself and STA B. The Calibration Correction frame is transmitted to complete the calibration procedure by STA A. The message includes the MIMO Correction Vector for STA B.

9 MAC Background (Informative)

This section contains informative MAC background material on scheduling, channel selection and the protection mechanisms.

9.1 Scheduling

9.1.1 Introduction

A data transfer sequence may contain multiple MSDUs carried in one or more aggregate PPDUs. The best throughput performance is gained when the length of the aggregate is maximized and the number of PPDUs is minimized. However, this maximum throughput is potentially gained at the expense of increased transport delay.
Apart from the effects of power-saving and 802.11e prioritization, a legacy 802.11 STA normally transmits MSDUs in the order that they are received at its MAC SAP.
Keeping this constraint on 802.11 HT would make throughput strongly dependent on the application traffic pattern at the MAC SAP. Therefore we allow a STA to reorder and delay transmission in the hope of building longer aggregates and using the medium more effectively.
Note: these schedulers are specific to single-receiver aggregates.

9.1.2 Scheduler 1

A STA keeps a queue per access category of Data MPDUs for transmission.
The STA starts its channel access attempt as soon as any queue is non-empty.
When a TXOP is gained, the STA selects one of it's transmit queues according to HCF rules.
The MPDU at the head of this queue defines the receiver address. The STA then scans it's transmit queues for more MPDUs addressed to the same receiver. These are available for transmission during the aggregate exchange.
For a sequence using embedded training, the rate may be modified when rate feedback is received. After an initial IAC/RAC training exchange, the STA can calculate how many of the MPDUs addressed to this receiver can fit in the available TXOP and transmits one or more aggregates containing these MPDUs.

9.1.3 Scheduler 2

In scheduler 1, the identity of the receiver is selected by HCF channel access, and is always defined by an MPDU at the top of an access category transmit queue. Furthermore, when an access category is not empty, a channel access attempt is active.
Scheduler 2 changes these two properties.
It delays EDCA channel access until:

- There is an aggregate to send exceeding some threshold length, or
- An MPDU has been delayed more than a certain threshold time in a transmit queue.

Once channel access is gained, the STA selects a receiver address as follows:

- If any MPDUs have exceeded the delay threshold, select the oldest one, or
- Select the receiver with the most buffered data

The effect of this algorithm is that an application flow (i.e. point-point link) that is saturated sees low delay and high throughput. Aggregate throughput is maximized because the scheduler prefers the longest aggregate.
An application flow that is low-rate and isochronous sees additional delay, because it never exceeds the length threshold, so it gets transmitted only when its delay threshold has been exceeded. [9]

9.1.4 Retransmissions and Allocation of Time

The strategy that maximizes throughput is to maximize the duration of aggregates, and minimize the number of PPDUs.
According to this strategy, a STA with a fully backlogged transmit queue will generate an aggregate that, together with the expected response, just fills the TXOP.
Any MPDUs that were not correctly received in the TXOP will have to wait for a subsequent channel access for retransmission.
An alternative strategy can be used that trades throughput for average channel access delay. This can be particularly useful with data for which a traffic specification (TSPEC) exists that specifies a delay bound.
In this case, the STA monitors or predicts the PER [10] From this, it can predict the number of failures in an aggregate, and leave enough room for retransmissions.
One way of allocating enough room is to require that the aggregate ($N_{aggregate}$) and its retries ($N_{retries}$) are expected to fit within the TXOP with some defined probability – known here as a confidence limit.
For any defined aggregate length and PER, the space to reserve for retries can be determined by inspecting a cumulative distribution function (Prob errors<N vs N) at the required confidence limit. If the packet errors are independent, a cumulative binomial distribution can be used.

9.1.5 Selecting RDL/RDG values (Informative)

9.1.5.1 Scheduling decisions for RDL/RDG

In a bi-directional data sequence, the initiator can control the amount of time spent by the responder in response (aggregate) PPDUs through use of the RDG element. There are two scheduling decisions to be made:

- How much of the time to offer to the responder (RDL and RDG)
- How to use the offered time at the responder

The bi-directional data sequence is designed to allow trained data to flow in both directions, sharing the overhead of channel access and training exchange.
The mechanism was designed specifically to allow a TCP link to "piggyback" TCP ack collection onto TCP data transmission. In the ideal mode of operation, the TCP data downlink for a half congestion window of data collect the TCP ack uplink for the previous half congestion window's worth of data.
In this mode, the transmissions by the responder are short TCP ACK MSDUs.

[9] Note, do not confuse this threshold with an MSDU lifetime limit. The delay threshold is the delay below which no channel access attempt need be made. The MSDU lifetime limit is the delay after which delivery is useless, which is hopefully a bigger number!

[10] How it does this is not relevant here.

Assuming that the Block ACK protocol is being used to allow the use of aggregate PPDUs, there is some limited opportunity for the responder to select Data MPDUs from its transmit buffer, subject to Block ACK windowing constraints so as to use the offered response duration.

9.1.5.2 Simple RDL/RDG heuristic

This section contains an initial recommended heuristic.
The initiator will normally try and fill the TXOP with its own data. Aggregate throughput is maximized by maximizing the length of the data in a TXOP. This is best achieved with known data rather than potential response data.
If there is an excess of data buffered for the chosen receiver address, the initiator does not advertise any reverse direction (RDL=0).
If there is room in a TXOP after delivering all its own data, the initiator can send an aggregate containing a non-zero RDL indicating the remaining TXOP duration. If the responder does not respond with a request (i.e. RDR=0), the sequence terminates.
If, on winning a TXOP, the initiator has less data than will clearly fill the TXOP, it can advertise some fraction of the TXOP in its initial IAC, thereby allowing possible reverse direction data flow early in the sequence, and maximizing throughput.
The responder uses any offered response duration for its data of the same user priority or access category, without any intentional reordering of data packets within the access category. This recommendation is based on the understanding that if there are no large packets in the responder's buffer, any TCP ACKs will get sent. But if there are, the responder's TCP ACKs can be sent as part of a later data exchange sequence it initiates, and that good efficiency will be achieved due to the presence of the large packets.
This proposed mechanism does not seek to achieve low TCP RTT, but only good utilization of the medium.

9.1.6 Scheduling heuristics for MRA

When creating an MRA, the HT AP may select receiver addresses so as to optimize throughput.
Typically, if clients are sending small amounts of data, the benefits of MRA aggregation should outweigh any reduction in transmit rate.
Above a certain data duration, this may no longer hold true. The HT AP can, in principle consider sending to an HT STA separately using its known transmit parameter versus including it in an existing MRA in order to determine which is the more effective use of the channel.
When using spoofing protection, the length of spoofed length field in the PLCP header imposes a limit of the total duration of a MRA and its multiple immediate ACKs/BAs plus data.

9.2 Method to determine Best channel and channel width

The method described here selects the channel parameters that maximize aggregate throughput.
It also has the side effect of generally avoiding operation with co-channel legacy devices.
The HT AP keeps a database of BSS and load information per 20 MHz channel. In each such channel it records its observations of load [11], and the number of legacy and HT BSSs.
Alternatively it can record load for legacy and HT traffic separately if it is capable of making that distinction.
Each HT BSS that permits 40MHz operation will have an entry for both 20-MHz channels it overlaps. The HT BSS load measurement is applied to both the overlapped channels.

[11] Load is defined as the percentage of time the channel is busy (averaged over some observation interval).

The device monitors or estimates its own load. If its load is not known, the device assumes a load of 100% for itself.
From any configuration of channel parameters, BSSs and loads, the device can calculate the following performance metric. It sums the product of load and rate over all channels to determine an aggregate throughput. For channels with only legacy BSSs, it uses the highest legacy rate. For channels with only HT devices, it uses the highest HT rate. For channels with a mixture, it first de-rates the HT highest rate to allow for use of a protection mechanism and then averages this with the legacy rate weighting the average according to the number of BSSs of each type. Alternatively, if it is collecting the load per BSS type, it sums the product of the appropriate rate (legacy or "mixed mode" HT rate) and load.
In considering a channel parameter configuration to evaluate the performance metric, the device shall not allow a configuration if it:

- Has an HT BSS with a different value of control channel or extension channel offset
- Has an HT BSS with a different permitted width set

9.2.1 First time Channel Selection

When starting operation, the device performs the following steps:

- It scans all channels to build the database.
- It considers adding its load into each channel in turn to find the position and width and evaluates the performance metric. As it adds its own load into a channel, it caps the load at 100%.
- It selects the combination of position and width that results in the highest value of the performance metric.

9.2.2 Subsequent Channel Selection

In order to update its channel parameters, the device first removes its estimated current load from its current channel(s) and then performs the same algorithm as described in the previous section.

9.3 Background to Protection Mechanisms

9.3.1 Introduction to Protection Mechanisms

Protection of HT exchanges from legacy devices (if present) and other HT devices is necessary for two reasons:

- Legacy devices do not understand HT PPDU formats
- On a link between two HT STAs that is adapted to the channel between those STAs, the probability that a 3rd party STA will be able to reliably receive the MPDUs is smaller than in the legacy case. This is particularly true when closed-loop adaptation is used.

There are two broad classes of protection defined in this document.
The first is called "Spoofing". It uses the legacy PLCP rate and length fields to indicate a duration that may or may not correspond to the duration implied by the HT signal field.

Spoofing is effective because legacy devices respect the duration contained in the legacy signal field, even if they cannot correctly decode the data. The signal field is transmitted using a robust modulation and omnidirectional antenna pattern.
The second class of protection relies on mechanisms to set the NAV using MAC layer signaling. In the presence of legacy devices, this requires that this signaling is carried in legacy receivable MPDUs.

9.3.2 Effect of Spoofing

The Spoofing Mechanism is an effective way to protect media from nodes that can not decode the "Duration" field in the MAC Header, which could be the majority of third party nodes in an 802.11n environment, particularly if Closed-Loop MIMO is being used.
The "Duration" field is purposed to set the NAV for third party nodes that the packet is not intended to, in order to prevent these nodes from interrupting the response from the recipient of that packet. However, when certain PHY features are used in 802.11n, the MAC Header which contains the "Duration" Field will not be decodable to all legacy nodes and most of the third party nodes. This is because the MAC Header will be coded with the optimized rate and mode for the recipient of the packet, which makes it difficult for third party nodes to correctly decode the MAC Header.
The spoofing mechanism solves this problem by virtually setting the NAV Duration in the PHY header where the information is coded at the lowest rate that every recipient can decode. By utilizing this Spoofed NAV, the system can assure higher protection of the media from legacy and third party nodes, without adding unwanted overhead of sending legacy packets for media protection.
An HT STA interprets a spoofed PLCP duration as one source of NAV setting, in addition to any MPDUs that it correctly receives in the PPDU. It can receive after the current PPDU regardless of the spoofed duration.
A legacy STA interprets a spoofed PLCP duration as a PPDU of that duration, and remains in the receiving state for the whole duration. It is therefore not capable of acquiring new PPDUs until this period has finished.

9.3.3 Pairwise spoofing

The Pairwise Spoofing mechanism is a way to maintain protection from hidden nodes, by disallowing third party nodes to initiate transmission during the duration of the intended response of a packet.
The mechanism takes after the "Virtual Carrier Sense" with the "Network Allocation Vector (NAV)" specified in the 802.11 standard. Each packet will contain a duration field (or a Spoofed NAV field: See section 7.1.7.1 for details), and nodes that receive packets destined to another node will assume channel occupancy for the time specified in that duration field.

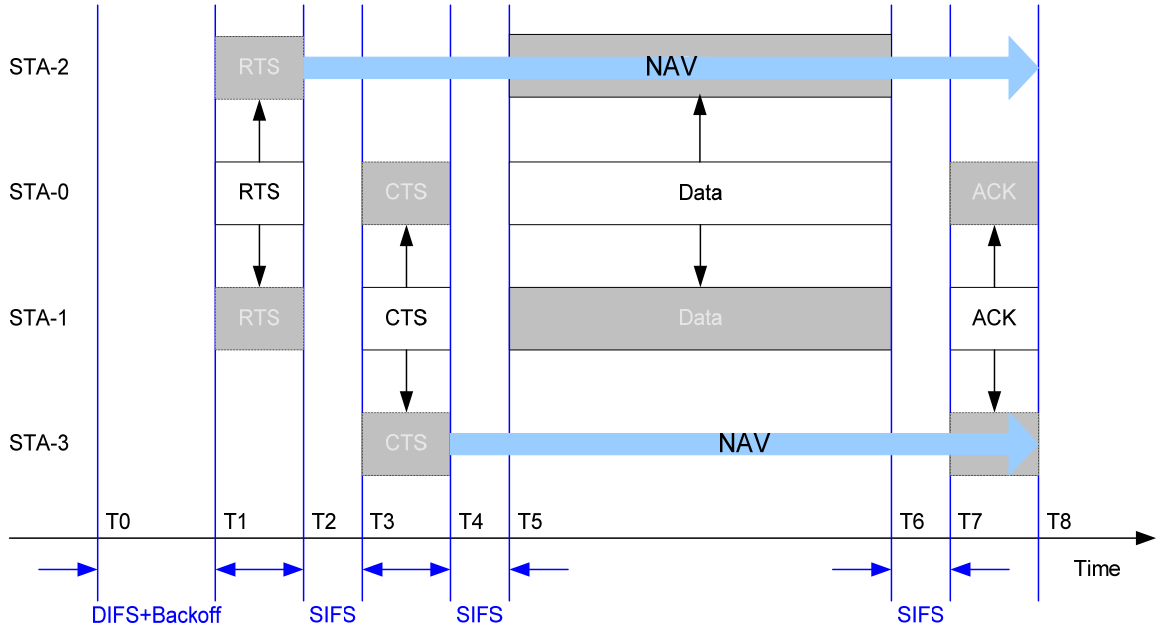


Figure 38 – Standard Virtual Carrier Sense

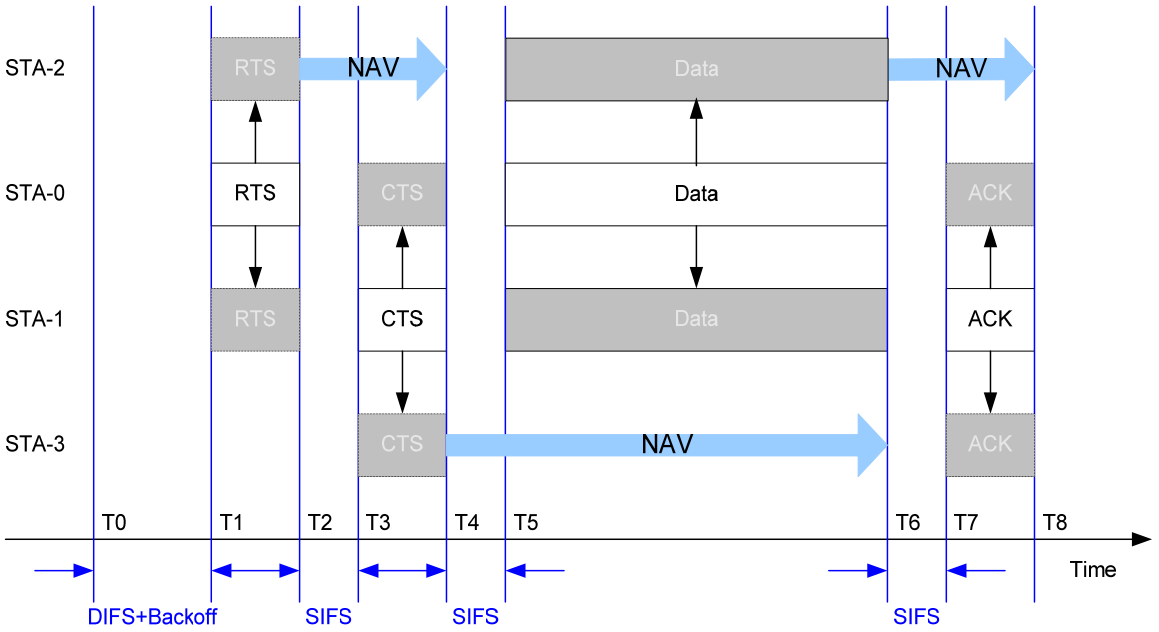


Figure 39 – Pairwise Spoofing Mechanism

With the Pairwise Spoofing method, the duration will be set to protect the intended response packet. In the standard 802.11 MAC, when using RTS / CTS, the duration fields in the RTS / CTS / DATA packets point at the end of the ACK. By using this original method, because the Duration must be fixed when sending RTS, the mode and rate in which the data packet would be sent at must be fixed before sending RTS. However, in realistic environments where advanced techniques like MIMO are used, channel conditions vary rapidly and optimization of MIMO modes by the RTS / CTS transaction will become essential. By shortening the NAV duration so that it protects only the following packet, the rate and mode that the DATA packet is sent at can be selected by the recipient of the packet, allowing more flexible adaptation, without sacrificing

the needed protection from hidden nodes. Furthermore, the shortening of the NAV can contribute to less waste of medium time when the RTS / CTS packets fail to be received by their counterparts.

9.3.4 Selecting between Pairwise Spoofing and LongNAV

The selection of protection mechanism is made by the initiator.
Use of Pairwise Spoofing is indicated by the FPD element in an IAC. This causes the RAC to generate a PPDU with a spoofed duration.
The initiator should operate under LongNAV rules if the NAV has already been set (e.g. by a QoS CF-poll sent by the HC).
In unmanaged mixed mode, protection from co-channel 802.11b devices shall be provided using LongNAV.
Otherwise the choice between Pairwise Spoofing and LongNAV can be made on the basis of trading flexibility of scheduling with performance.

9.3.5 HT PPDU Format with Legacy Protection

An HT PPDU Format looks like the following (See section 11.1.1.3 for a detailed description of the HT PPDU Format.)

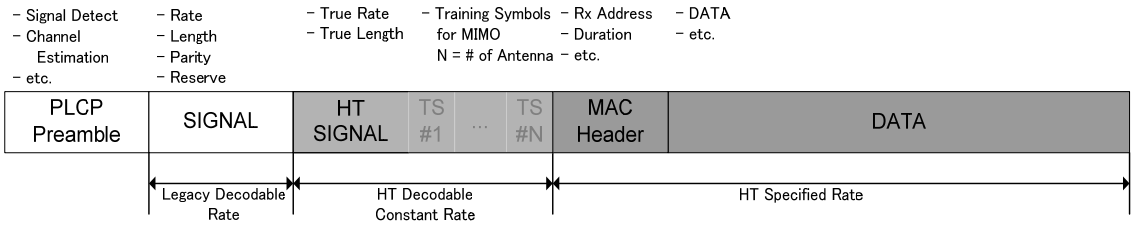


Figure 40 – PPDU Format with Legacy Protection

The PLCP Preamble is followed by a legacy SIGNAL Field. This SIGNAL Field is coded at a constant legacy rate that all nodes can decode at (for 802.11a nodes, 6[Mbps]). The legacy SIGNAL Field contains Rate and Length Fields used for spoofing the legacy nodes.
The legacy SIGNAL Field is followed by an HT SIGNAL Field which contains the True Rate and True Length. This field is coded at a rate and mode that all HT nodes can decode.
The HT SIGNAL Field is followed by the MAC Header. In the MAC Header, the Duration Field is set for setting the NAV for third party nodes that receive this packet. However, from the MAC Header on, the mode and rate is optimized for the intend recipient of the packet, so third party nodes will very unlikely be able to decode the MAC Header properly which will destroy protection by the NAV mechanism.

9.3.5.1 Interpretation by a Legacy Node

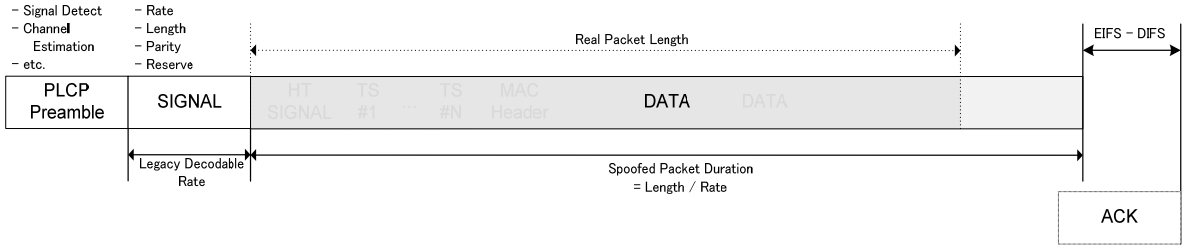


Figure 41 – When Received by a Legacy Node

When this HT PPDU is received by a legacy node, it decodes the legacy SIGNAL Field and tries to receive the packet for Length / Rate. When it finishes its reception, it finds that the CRC has failed for the data frame and defers for the EIFS period (instead of DIFS). After the media has been idle for that EIFS period, it restarts its contending mechanism. Therefore, an HT node who requires protection from legacy nodes, will set its Length and Rate Field so that the;
Length / Rate = NAV Duration – (EIFS – DIFS)
There could be a number of combinations of Lengths and Rates to achieve the same "Length / Rate", so the initiator will have the liberty to choose among the combinations. The maximum duration that an initiator can set by spoofing is determined by the maximum length of the packet and the minimum rate that the original protocol can withstand.
The signal energy level will go down during the reception of the packet, because the actual packet length is less than the spoofed length. But, according to the IEEE 802.11a MAC Specification (See IEEE Std 802.11a – 1999, page37, Figure 125), legacy nodes are required to honour the Length and Rate field, and must keep quiet for the Length / Rate. This requirement is also a part of the compliant test plan specified by the WiFi Alliance.

9.3.5.2 Interpretation by an HT Node

As seen in Figure 40, in an HT PPDU, a legacy SIGNAL Field is followed by an HT SIGNAL Field coded at unified rate. This HT SIGNAL Field will contain the True Rate and the True Length that HT nodes will use. However, after decoding the legacy SIGNAL Field, an HT node will not know whether the SIGNAL Field is spoofed and is followed by an HT SIGNAL Field (sent at unified rate), or the packet is a legacy "un-spoofed" packet and is followed by a MAC Header coded at a rate specified in the SIGNAL Field.

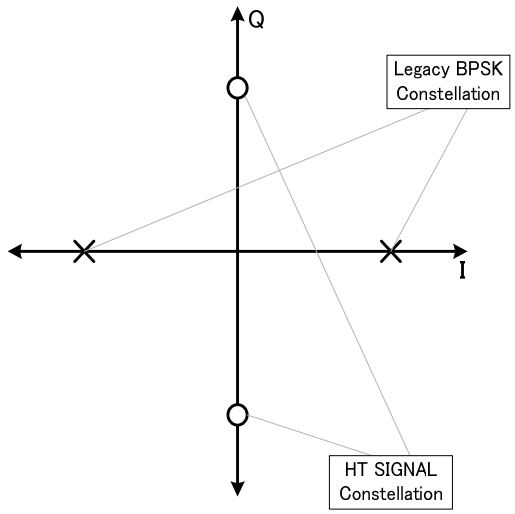


Figure 42 – HT Signal Field BPSK Constellation

Therefore, an HT SIGNAL Field will be transmitted with a 90[degrees] rotation, for BPSK mapping, from the I-Q plane indicated in the legacy PLCP Preamble. An HT device will decode the BPSK constellation specified in the PLCP Preamble at the rate indicated in the SIGNAL Field, and at the same time will decode, at the unified rate for HT SIGNAL Field, the BPSK constellation 90[degrees] from its original.
After decoding the two constellations in parallel for a given length of the HT SIGNAL Field, an HT node will check its CRC for the HT SIGNAL Field. If the CRC passes, it will recognize that

this is an HT PPDU and continue to decode the MAC Header using the True Rate and True Length specified in the HT SIGNAL Field. If the CRC fails, it will recognize that it is a legacy PPDU and interpret the legacy MAC Header that it has been receiving.

9.3.5.2.1 MPDU Not Decodable

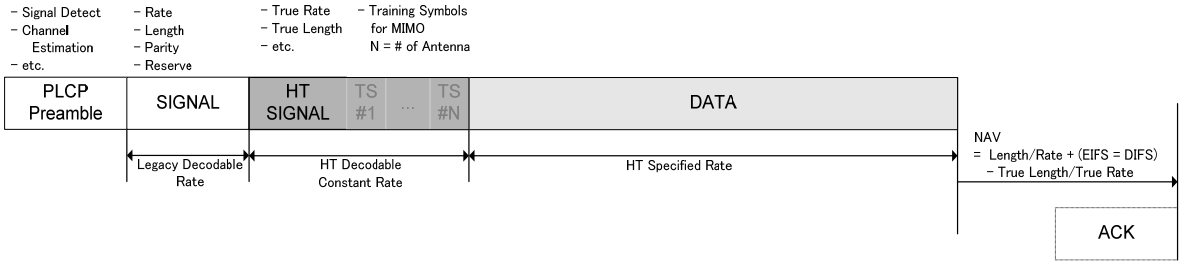


Figure 43 – When Received by an HT Node but MPDU is not Decodable

After decoding the HT SIGNAL Field, an HT node will attempt to decode the following MPDU with the specified rate. When the MPDU's FCS is not correct, it will discard the MPDU and will calculate the NAV duration that it should defer. Because the spoofed NAV Duration is (EIFS – DIFS) less than the intended Duration, to prevent unfairness to legacy nodes, the NAV will be set according to the rules defined in section **Error! Reference source not found.**.

9.3.5.2.2 MPDU Decodable

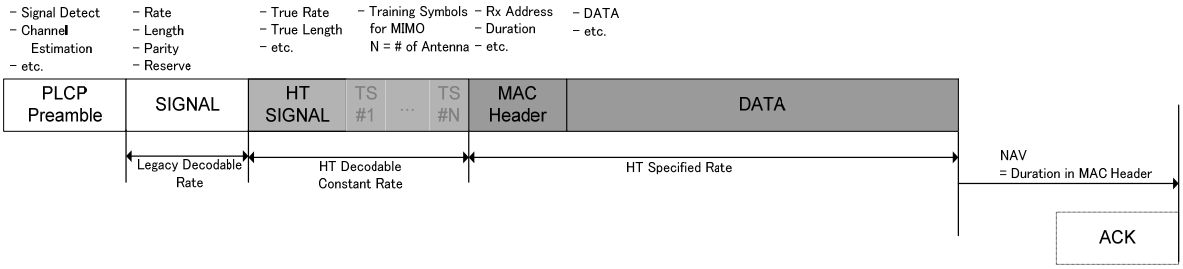


Figure 44 – When Received by an HT Node and the MPDU is Decodable

When the MPDU is decodable, an HT node will continue on with its MAC operation. If the Receive Address in the MAC Header does not match itself, it will set its NAV to the Duration indicated following existing 802.11 rules.

9.3.6 Protection mechanisms for MRA

9.3.6.1 How to choose between LongNAV and spoofing in MRA

The choice of protection mechanism is the choice of the initiator. If using a polled TXOP, or during the CFP, NAV protection has already been provided, and spoofing is not necessary. However if the NAV has not been set, it is more effective to use spoofing protection.
There is only one possible performance advantage for LongNAV, if the final (or only) response PPDU is shorter than the grant, the initiator does not need to wait until the end of the scheduled response period, but can initiate additional transmissions a SIFS after the end of the actual response.
The recommended choice of protection mechanism is listed below:
Use LongNAV:

- During Polled TXOPs (zero overhead)
- For long aggregates
- Where maximum flexibility in scheduling decisions must be kept
- For .11b coexistence

Use spoofing:

- For short aggregates
- In IBSS or ad-hoc mesh networks

9.3.6.2 Protection from Third Party HT Nodes

The protection of multi-responses from third party HT nodes is achieved using either single-ended spoofing or LongNAV.
In both cases, all STAs that can hear the initiator will have a non-zero NAV during the response periods, preventing them from transmission.
A response period for a receiver might not be used up, for some reason. This leaves the WM idle, possibly for more than a SIFS time. Although this unused TXOP is wasted, it does not create any problem, since all nodes that can hear the initiator will have their NAV set during the responses.

9.3.6.3 Protection from Hidden Nodes

In the case of MRA transmission from an HT AP, all nodes within the BSS should be able to hear and honor the spoofed NAV in the aggregate header, which is always transmitted at the most robust rate, there are no nodes hidden to the AP.

9.3.6.4 Protection from Legacy Nodes

Protection of the MRA from legacy nodes may be achieved via the spoofed NAV in the PLCP header of a MRA aggregate. Although not recognizing an HT MPDU, a legacy client should still be able to recognize such a spoofed NAV from the legacy PLCP header of the MRA aggregate and therefore enters a receiving state for the indicated duration, thereby stopping it from initiating transmissions during the indicated duration.
Alternatively the initiator transmits a legacy compatible PPDU that sets the NAV.
In both cases, the level of protection afforded by this method is the same as CTS-to-self.

1

10 PHY-SAP service specification

The PHY service is provided to the MAC entry at an STA through the PHY service access point (PHY-SAP), as shown in Figure 11 of [1] and described in Clause 12 of that reference.

The amendments to the PHY-SAP parameters TXVECTOR and RXVECTOR are described in Sections 10.1.1 and 10.1.2 respectively.

10.1 PHY Specific Service Parameter List

The architecture of the IEEE 802.11 MAC is intended to be PHY independent. Some PHY implementations require PHY-dependent MAC state machines running in the MAC sublayer in order to meet certain PMD requirements. The PHY-dependent MAC state machine resides in a sublayer defined as the MAC sublayer management entity (MLME). In certain PMD implementations, the MLME may need to interact with the PLME as part of the normal PHY SAP primitives. These interactions are defined by the PLME parameter list currently defined in the PHY service primitives as TXVECTOR and RXVECTOR. The list of these parameters, and the values they may represent, are defined in the specific PHY specifications for each PMD. This subclause addresses the TXVECTOR and RXVECTOR for the MIMO-OFDM High Throughput (HT) PHY. The service parameters for TXVECTOR and RXVECTOR shall follow subclause 10.1.1 and 10.1.2.

10.1.1 TXVECTOR

The TXVECTOR is an array of fields sent from the MAC to the PHY instructing the PHY how to transmit the PPDU.

Parameter	**Associated Primitive**	***Value***
PPDU FORMAT	PHY-TXSTART.request (TXVECTOR)	*Legacy 20 in 20* *HT 20 in 20* *Legacy upper 20 in 40* *Legacy lower 20 in 40* *HT upper 20 in 40* *HT lower 20 in 40* *HT 40 in 40*
L-LENGTH	PHY-TXSTART.request (TXVECTOR)	*1 – 4095*
L-DATARATE	PHY-TXSTART.request (TXVECTOR)	*6, 9, 12, 18, 24, 36, 48, and 54 (Support of 6, 12, and 24 data rates are mandatory.)*
L-SERVICE	PHY-TXSTART.request (TXVECTOR)	*Scrambler initialization;* *7 null bits + 9* *reserved null bits*
TX_PWR_LEVEL	PHY-TXSTART.request (TXVECTOR)	*1 – 8*
HT-LENGTH	PHY-TXSTART.request (TXVECTOR)	*1 – 262143*

MCS	PHY-TXSTART.request (TXVECTOR)	*0 – 63*
ADV CODING	PHY-TXSTART.request (TXVECTOR)	*0, 1*
SOUNDING	PHY-TXSTART.request (TXVECTOR)	*0, 1*
NLTF	PHY-TXSTART.request (TXVECTOR)	*1 – 4*
AGGREGATE	PHY-TXSTART.request (TXVECTOR)	*0, 1*
SCRAMBLER-INIT	PHY-TXSTART.request (TXVECTOR)	*0, 1, 2, 3*
SHORT-GI	PHY-TXSTART.request (TXVECTOR)	*0, 1*
BEAMFORMED	PHY-TXSTART.request (TXVECTOR)	*0, 1*
BEAMFORM MATRIX	PHY-TXSTART.request (TXVECTOR)	$N_{TX} \times N_{SS} \times N_{SD}$ *complex matrices*

Table 17: Table of parameters for the TX Vector

10.1.1.1 TXVECTOR PPDU FORMAT

The PPDU FORMAT field specifies whether the PPDU is to be sent using a legacy PPDU format or using a high-throughput format. Legacy is interpreted to be 802.11g if operating in the 2.4 GHz band and 802.11a if operating in the 5 GHz band. Furthermore, it specifies the channel (i.e, the full 20MHz channel, the upper 20MHz sub-channel in the 40MHz channel, the lower 20MHz sub-channel in the 40MHz channel, or the full 40MHz channel) at which the PPDU is to be sent.

The value of this field affects the interpretation of several other TXVECTOR fields.

10.1.1.2 TXVECTOR L-LENGTH

The allowed values for the L-LENGTH parameter are in the range of 1–4095.

In legacy format this parameter is used to indicate the number of octets in the MPDU which the MAC is currently requesting the PHY to transmit. This value is used by the PHY to determine the number of octet transfers that will occur between the MAC and the PHY after receiving a request to start the transmission. In high-throughput format this parameter (along with L-DATARATE) is used to spoof the length of the current PPDU and possibly additional PPDUs.

10.1.1.3 TXVECTOR L-DATARATE

In legacy format the L-DATARATE parameter describes the bit rate at which the PLCP shall transmit the PSDU. Its value can be any of the rates: 6, 9, 12, 18, 24, 36, 48, and 54. Data rates of 6, 12, and 24 shall be supported; other rates may also be supported. In high-throughput format

this parameter (along with L-LENGTH) is use to spoof the length of the current PPDU and possibly additional PPDUs. For spoofing purposes, L-DATARATE in the L-SIG is set to 6Mbps.

10.1.1.4 TXVECTOR L-SERVICE

In legacy format the L-SERVICE parameter consists of 7 null bits used for the scrambler initialization and 9 null bits reserved for future use. In HT format this parameter is Null.

10.1.1.5 TXVECTOR TXPWR_LEVEL

The allowed values for the TXPWR_LEVEL parameter are in the range from 1–8. This parameter is used to indicate which of the available TxPowerLevel attributes defined in the MIB shall be used for the current transmission. In high-throughput format the transmit power is the aggregate of the power from all transmit antennas.

10.1.1.6 TXVECTOR HT-LENGTH

The allowed values for this parameter are 1 – 264143. In HT format this parameter indicates the number of octets in the PSDU to be sent in the PPDU. In legacy format this parameter is Null.

10.1.1.7 TXVECTOR MCS

The allowed values for this parameter are 0 – 63. In HT format this parameter indicates the modulation and coding scheme. The number is an index into the MCS table given in Table 26 and Table 29. In legacy format this parameter is Null.

10.1.1.8 TXVECTOR ADV CODING

In HT format this parameter indicates whether advanced coding is used. A value of 0 indicates that advanced coding is not used. In legacy format this parameter is Null. A value of 1 indicates LDPC coding is used.

10.1.1.9 TXVECTOR SOUNDING

In HT format this parameter indicates whether this PPDU is a sounding PPDU. A value of 0 indicates that it is not a sounding PPDU. A value of 1 indicates that it is a sounding PPDU. In legacy format this parameter is Null.

10.1.1.10 TXVECTOR NLTF

In high-throughout format this parameter specifies the number of long training fields. In legacy format this parameter is Null.

10.1.1.11 TXVECTOR AGGREGATE

In high-throughout format this parameter indicates whether the PPDU contains an aggregate PSDU. A value of 0 indicates that the PSDU is not an aggregate and a value of 1 indicates that the PSDU is an aggregate. In legacy format this parameter is Null.

10.1.1.12 TXVECTOR SCRAMBLER-INIT

In high-throughout format this is the scrambler initialization field. It is a two-bit field, which serves as a pointer to one of four 7-bit long scrambler initialization sequences. In legacy format this parameter is Null.

The 2 bits are mapped to the following four scrambler initialization sequences:

b'00 → step 15 b'0111100
b'01 → step 84 b'0111101
b'10 → step 103 b'0111110
b'11 → step 127 b'0111111

Table 18: Scrambler Initialization Sequence

10.1.1.13 TXVECTOR SHORT-GI

In high-throughout format this parameter indicates whether the data is sent with a short guard interval (400 ns). A value of zero indicates a full GI (i.e. 800ns) and a value of 1 indicates a short GI (i.e. 400 ns). In legacy format this parameter is Null.

10.1.1.14 TXVECTOR BEAMFORMED

In high-throughout format this parameter indicates whether a beamforming matrix is applied to the data. A value of zero indicates that no beamforming is applied and a value of 1 indicates that beamforming is applied. In legacy format this parameter is Null.

10.1.1.15 TXVECTOR BEAMFORM MATRIX

In high-throughout format this field specifies the complex matrix $N_{TX} \times N_{SS} \times N_{SD}$ which is used to beamform the PPDU. This field is only valid when the BEAMFORMED parameter is 1. In legacy format this parameter is Null.

10.1.2 RXVECTOR

The RXVECTOR is an array of fields sent from the PHY to the MAC providing information regarding the received PPDU.

Parameter	Associated Primitive	*Value*
PPDU FORMAT	PHY-RXSTART.indicate (RXVECTOR)	*Legacy 20 in 20* *HT 20 in 20* *Legacy upper 20 in 40* *Legacy lower 20 in 40* *HT upper 20 in 40* *HT lower 20 in 40* *HT 40 in 40*
L-LENGTH	PHY-RXSTART.indicate (RXVECTOR)	*1 – 4095*
RSSI	PHY-RXSTART.indicate (RXVECTOR)	*0 – RSSI maximum*
L-DATARATE	PHY-RXSTART.indicate (RXVECTOR)	*6, 9, 12, 18, 24, 36, 48, and 54*
BW	PHY-RXSTART.indicate (RXVECTOR)	*20, 40*
HT-LENGTH	PHY-RXSTART.indicate (RXVECTOR)	*1 – 262143*
MCS	PHY-RXSTART.indicate (RXVECTOR)	*0 – 63*
ADV CODING	PHY-RXSTART.indicate (RXVECTOR)	*0, 1*
SOUNDING	PHY-RXSTART.indicate (RXVECTOR)	*0, 1*
NLTF	PHY-RXSTART.indicate (RXVECTOR)	*1 – 4*
AGGREGATE	PHY-RXSTART.indicate (RXVECTOR)	*0, 1*
SHORT-GI	PHY-RXSTART.indicate (RXVECTOR)	*0, 1*

Table 19: Table of parameters for the RX Vector

10.1.2.1 RXVECTOR PPDU FORMAT

The PPDU FORMAT field specifies whether the received PPDU has a legacy PPDU format or a high-throughput format. Legacy is interpreted to be 802.11g if operating in the 2.4 GHz band and 802.11a if operating in the 5 GHz band. Furthermore, it specifies the channel (i.e, the full 20MHz channel, the upper 20MHz sub-channel in the 40MHz channel, the lower 20MHz sub-channel in the 40MHz channel, or the full 40MHz channel) at which the received PPDU was sent.

10.1.2.2 RXVECTOR L-LENGTH

The allowed values for the L-LENGTH parameter are in the range of 1–4095. This parameter is used to indicate the value of the L-LENGTH field in received PPDU. If this was a legacy PPDU then this parameter is the actual length of the PPDU. If this is a HT PPDU then this parameter, along with the L-DATARATE parameter, is used to calculate the spoofed duration.

10.1.2.3 RXVECTOR RSSI

The allowed values for the receive signal strength indicator (RSSI) parameter are in the range from 0 through RSSI maximum. This parameter is a measure by the PHY sub-layer of the energy observed at the antenna used to receive the current PPDU. RSSI shall be measured during the reception of the PLCP preamble. RSSI is intended to be used in a relative manner, and it shall be a monotonically increasing function of the received power. The value of RSSI is the maximum of the RSSI values from each receive antenna.

10.1.2.4 RXVECTOR L-DATARATE

The allowed values for L-DATARATE are: 6, 9, 12, 18, 24, 36, 48, and 54. This parameter is used to indicate the value of the L-DATARATE field in the received PPDU. If this is a legacy PPDU this is the data rate at which the data was transmitted. If this is an HT PPDU then this parameter, along with L-LENGTH, is used to calculate the spoof duration.

10.1.2.5 RXVECTOR BW

In high-throughput format this parameter specifies whether the received PPDU is a 20 MHz bandwidth or 40 MHz bandwidth HT PPDU. In legacy format this parameter is Null.

10.1.2.6 RXVECTOR HT-LENGTH

The allowed values for this parameter are 1 – 264143. In HT format this parameter indicates the value contained in the HT-LENGTH field. This parameter specifies the number of octets to be transferred from the PHY to the MAC sub-layer. In legacy format this parameter is Null.

10.1.2.7 RXVECTOR MCS

The allowed values for this parameter are 0 – 63. The number is an index into the MCS table given in Table 26 and Table 29. In high-throughput format this parameter indicates the value of the MCS field in the PPDU HT-SIG. In legacy format this parameter is Null.

10.1.2.8 RXVECTOR ADV CODING

In HT format this parameter indicates whether advanced coding was used in the PPDU. A value of 0 indicates that advanced coding is not used. In legacy format this parameter is Null. A value of 1 indicates that LDPC coding is used.

10.1.2.9 RXVECTOR SOUNDING

In HT format this parameter indicates whether this PPDU is a sounding PPDU. A value of 0 indicates that it is not a sounding PPDU. A value of 1 indicates that it is a sounding PPDU. In legacy format this parameter is Null.

10.1.2.10 RXVECTOR NLTF

In high-throughout format this parameter specifies the number of long training fields. In legacy format this parameter is Null.

10.1.2.11 RXVECTOR AGGREGATE

In high-throughout format this parameter indicates whether the PPDU contains an aggregate PSDU. A value of 0 indicates that the PSDU is not an aggregate and a value of 1 indicates that the PSDU is an aggregate. This value is used by the MAC sub-layer to indicate whether de-aggregation is required. In legacy format this parameter is Null.

10.1.2.12 RXVECTOR SHORT-GI

In high-throughout format this parameter indicates whether the data was sent with a short guard interval (400 ns). A value of zero indicates a full GI (800 ns) and a value of 1 indicates a short GI (400 ns). In legacy format this parameter is Null.

11 MIMO-OFDM HT PHY specification

11.1 Overview

11.1.1 Introduction (Informative)

With the aim to increase the maximum PHY layer data rate, we have introduced two primary techniques that are mandatory in our proposal:

- Transmission of Multiple Spatial Streams via Multiple Transmit Antennas. Our recommendation is to make 2 spatial streams mandatory, with the option to scale to 4 spatial streams.

- Extended Bandwidth signaling. 20MHz channelization shall be mandatory worldwide. Our recommendation is to support 40MHz channelization as mandatory in regulatory domains that will permit wider bandwidth operation

We have also introduced several secondary techniques for increasing throughput. These include:

- A reduced guard interval of 400ns (i.e. half-GI), channel dispersion permitting (e.g. TGn channel model B). Support for half-GI is mandatory in 20MHz.

- Higher channel coding rate of 7/8. Support for 7/8 is mandatory in 20MHz.

- Support for 256-QAM is optional in our proposal.

Due to the susceptibility of a wireless link to multipath fading, especially in a MIMO channel, our proposal introduces two optional techniques for improving the robustness of the link between the transmitter and the receiver, with the aim to keep the client (station) configuration low-cost and low-power:

- Transmitter beamforming, with the understanding that Access Points will be able to support more than 2 antennas, while the client stations would be restricted to 2 antennas.

- Advanced channel coding. Our proposal includes LDPC as an option. LDPC codes exhibit a larger minimum free distance compared to the existing 802.11a convolutional code, and hence are able to better exploit the frequency and spatial diversity in the channel.

11.1.1.1 Frequency Bands of Operation

Our proposal supports operation of devices in 2.4GHz, 5GHz, and 4.9GHz. Operation in 4.9GHz shall conform to 802.11j rules.

11.1.1.2 Options for throughput enhancement

The throughput calculation for 802.11a, at its highest data rate, can be factorized as:

$$\text{Data Rate} = \underbrace{\frac{20M \text{ time samples}}{\text{second}}}_{\text{channel spacing}} \cdot \underbrace{\frac{48 \text{ freq tones}}{64 \text{ freq tones}}}_{\text{guard band overhead}} \cdot \underbrace{\frac{6 \text{ coded bits}}{\text{freq tone}}}_{\text{constellation size}} \cdot \underbrace{\frac{3 \text{ info bits}}{4 \text{ coded bits}}}_{\text{coding rate}} \cdot \underbrace{\frac{64 \text{ freq tones}}{80 \text{ times samples}}}_{\text{guard interval overhead}}$$

$$= 54M \text{ info bits/second}$$

Besides spatial multiplexing, there are four other options for increasing data rate: (a) increasing the channel bandwidth, (b) increasing the constellation size, (c) increasing the coding rate, and (d) reducing the guard interval (GI) overhead. These options are summarized in Table 20.

802.11a/g	**Our Proposal**	**Requirement**	***Throughput Scaling Factor***
Channel BW = 20MHz Number of data subcarriers = 48	Channel BW = 20MHz Number of data subcarriers = 48	Mandatory in device	*1x*
	Channel BW = 40MHz Number of data subcarriers = 108	Mandatory in device	*2.25x*
	Channel BW larger than 40MHz	Not supported in device	
Number of Transmit Antennas = 1	Number of Transmit Antennas = 2	Mandatory in device	*2x*
	Number of Transmit Antennas > 2	Optional in device (e.g. 3 and 4)	*3x or 4x*
Maximum Constellation Size = 64QAM	64-QAM	Mandatory in device	*1x*
	greater than 64QAM (i.e. 128 or 256 QAM)	256 QAM is optional in device	*1.16x (128-QAM)* *1.33x (256-QAM)*
GI = 800ns Tsymbol = 3200ns	GI / Tsymbol = 800ns/3200ns	Mandatory in device	*1x*
	GI / Tsymbol = 400ns/3200ns	Mandatory in device	*1.11x*
	GI / Tsymbol = 800ns / 6400ns	Not supported in device	*1.11x*
	GI / Tsymbol = 400ns / 6400ns	Not supported in device	*1.176x*
Coding Rate	1/2, 2/3, 3/4	Mandatory in device	*1x*
	7/8 (advanced coding option)	Mandatory in device	*1.167x*

Table 20: Options for Throughput Enhancement

Of the options presented in Table 20 for throughput enhancement, our proposal focuses upon using spatial multiplexing (via multiple transmit antennas) and bandwidth expansion (from 20 MHz to 40 MHz) as the primary means for increasing the physical layer data rates. This is because both spatial multiplexing and channel bandwidth increase data rate by integer increments, whereas the rest of the techniques increase the data rate by fractional increments. The use of a shorter Guard Interval (GI=400ns) and 7/8-th rate coding is also mandatory for the data portion of the PPDU. With 4 spatial streams in 40MHz, the maximum PHY data rate in our proposal corresponds to 630Mbps, which is more than 10x the current data rate of 54Mbps. Note that the High Throughput (HT) preamble shall always use the full GI, i.e. 800us.

The basic mandatory configuration in our PHY proposal for throughput enhancement is 2 spatial streams (i.e. 2 transmit antennas) across a 20MHz channel bandwidth. The transmitter datapath is shown in Figure 45.

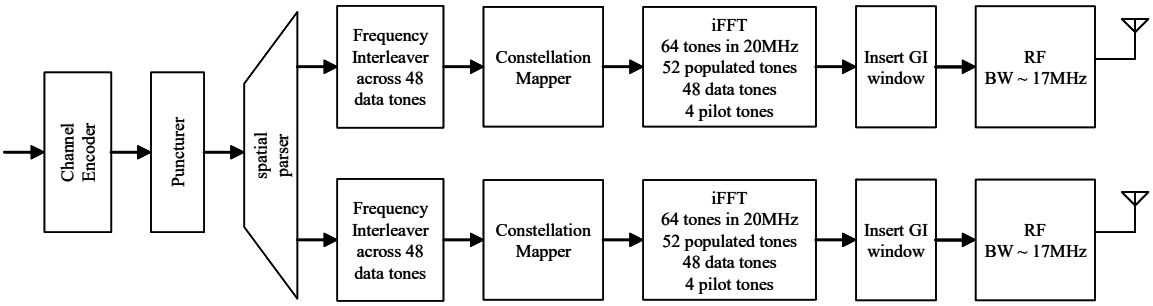


Figure 45: Transmitter Datapath for 2-antenna MIMO in 20MHz

Our proposal also includes provisions for 40MHz signaling in regulatory domains that permit it. The transmitter datapath for such a configuration is shown in Figure 46:

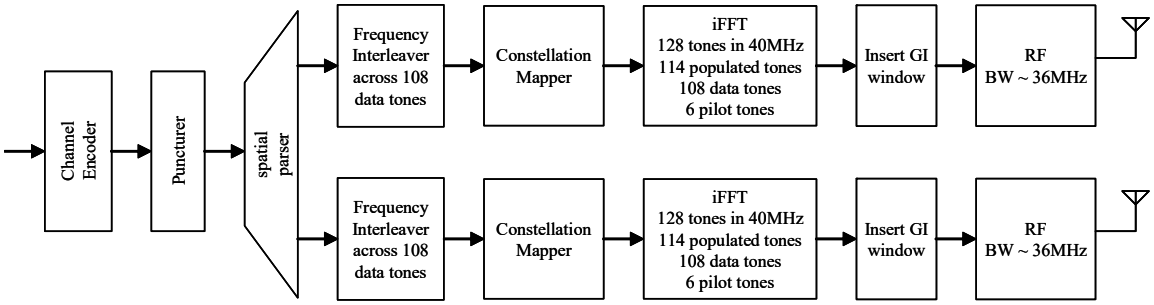


Figure 46: Transmitter datapath for 2-antenna MIMO in 40MHz[12]

11.1.1.3 PPDU Format for the Basic (Mandatory) MIMO Transmission

The PPDU format for transmission using 2 antennas in a 20MHz channelization is shown in Figure 47:

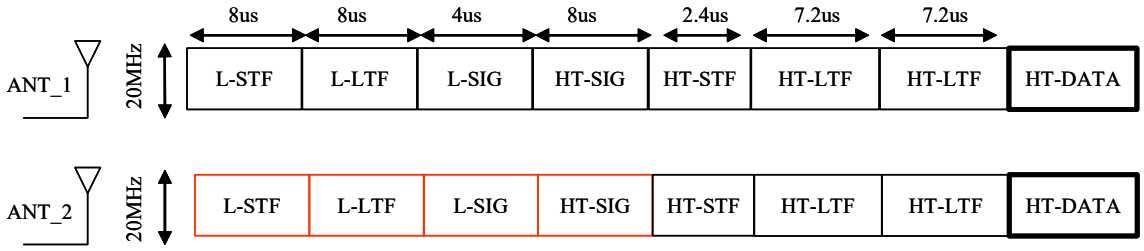


Figure 47: PPDU Format for 2x20 Mandatory Basic MIMO Transmission

In the 40MHz bandwidth, there are two 20MHz wide sub-channels. The upper sub-channel spans from 0 to +63, and the lower sub-channel spans from -64 to -1, as shown in Figure 48. A 40MHz-HT device receiving a 20MHz-only transmission shall indicate to the MAC which sub-channel is currently active.

[12] 40MHz operation by an HT device shall only be permitted in those regulatory domain that allow it.

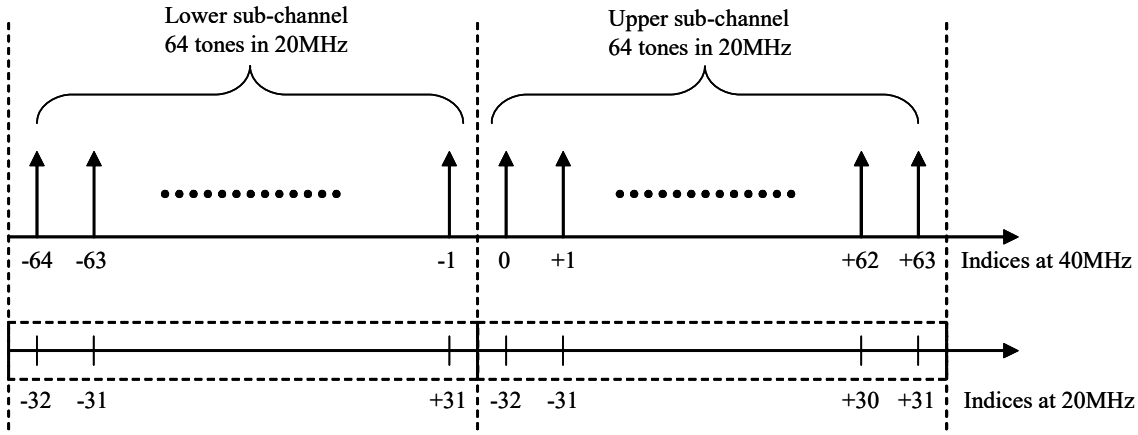


Figure 48: Upper and Lower sub-channel specification in 40MHz

The PPDU format from two transmit antennas, each 40MHz wide, is shown in Figure 49.

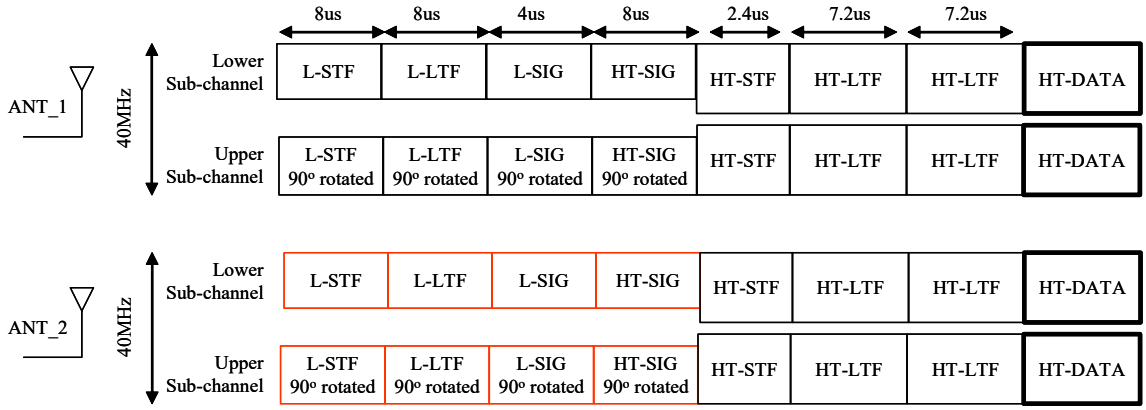


Figure 49: PPDU Format for 2x40 Mandatory Basic MIMO Transmission[13]

The legacy part of the HT preamble (i.e. L-STF to L-SIG) and HT-SIG may either be transmitted from one antenna or both antennas. If transmitted from both antennas, the mapping of this single spatial stream to two antennas has to be such that beamforming in the far-field is mitigated. One method for achieving this is to use a cyclical delay diversity (CDD) mapping. The optional part of the PPDU, corresponding to the CDD format, is shown in red color in Figure 47 and Figure 49.

11.1.1.4 HT PCLP Preamble for Synchronization purposes and Legacy Interoperability

The functions performed by the preamble include:

1. Start-of-Packet (SoP) detection
2. AGC
3. Coarse Frequency Offset Estimation
4. Coarse Timing Offset Estimation

[13] 40MHz operation by an HT device shall only be permitted in those regulatory domain that allow it.

5. Fine Timing Offset Estimation
6. Fine Frequency Offset Estimation
7. Channel Estimation

The High Throughput (HT) preamble in our proposal is a concatenation of the legacy preamble (identical to 802.11a/g) and a HT-specific preamble, as shown in Figure 50:

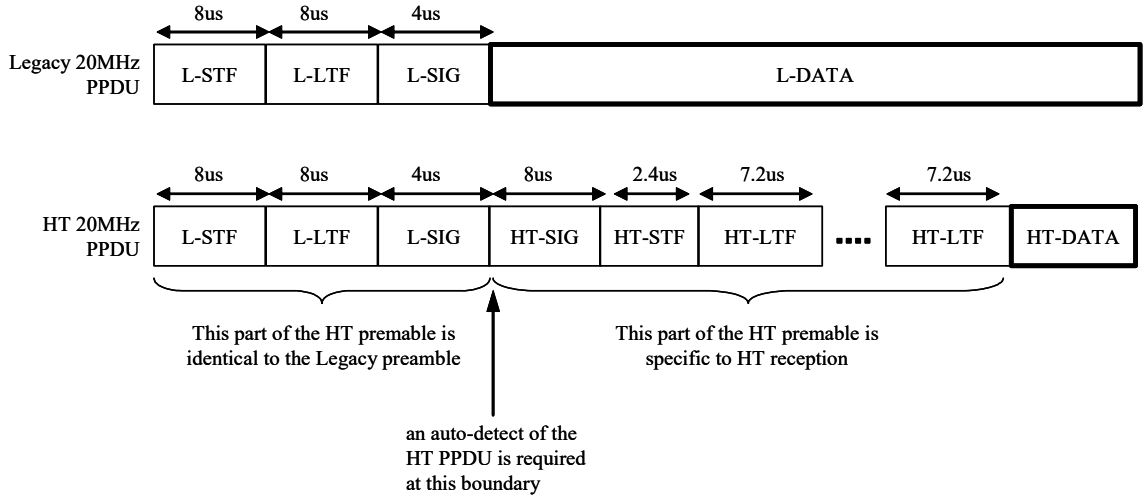


Figure 50: HT Preamble format

Concatenation of the legacy preamble allows PHY layer interoperability with 802.11a and ERP-802.11g modems.

A HT receiver does not know apriori whether it is receiving a legacy PPDU or a HT PPDU. Hence, a HT receiver has to perform an auto-detect after the legacy signal field (L-SIG) as to whether a HT-SIG is being received or legacy data (L-DATA). Our preamble format includes two provisions for auto-detecting the HT-SIG: (a) transmitting the HT-SIG as a BPSK signal on the quadrature axis instead of the in-phase axis, and (b) inverting the pilot polarity in going from the L-SIG to the HT-SIG.

The legacy part of the HT preamble allows a legacy receiver to properly decode the L-SIG. A HT receiver shall also perform functions 1 through 7 (SoP to Channel Estimation) to successfully decode the HT-SIG. It is anticipated that functions 1 through 4 (and perhaps function 5) are performed by the legacy short training field (L-STF), and functions 5 through 7 are performed by the legacy long training field (L-LTF).

Since a HT device (in our proposal) shall support multiple transmit antennas, the legacy part of the preamble can either be transmitted from one antenna or from multiple antennas. If more than one antenna is used for transmission, it is recommended that the signal be shifted in a cyclical fashion across the transmit antenna array. This technique introduces cyclical delay diversity (CDD), and it tends to reduce the “effective” length of the guard interval, as seen by a legacy device. Hence, the total shift across the transmit antenna array should be limited to 50ns.

In the 40MHz mode, the L-STF in the lower sub-channel is identical to legacy, whereas the L-STF in the upper sub-channel is a 90 degree phase rotation of the legacy STF. This rotation is performed to keep the PAPR of the STF low (see Figure 48 and Figure 49)

Since a legacy OFDM RX is not expected to "understand" any transmission after the L-SIG, the RATE and LENGTH fields in the L-SIG are "spoofed" such that a legacy OFDM RX can defer for the correct period of time.

The HT-SIG contains information needed to correctly de-map and decode the HT data field (i.e. the MIMO-OFDM symbols). It also contains fields for supporting advanced techniques that are described next.

Since the L-STF does not provide a sufficiently accurate power measurement for MIMO AGC purposes, a HT-STF has been introduced after the HT-SIG. The HT-STF is tone-interleaved symbol across the transmit antenna array. The HT short training symbol uses 24 tones (L-STF uses 12 tones in a 20 MHz channel). The choice of 24 tones permits accurate AGC measurements across 4 transmit antennas. In the 40 MHz channel, 48 tones are used in the HT-STF. The 24 (48) tones are interleaved across the antennas, i.e., 6 (12) tones are used per antenna if there are 4 transmitter antennas, and likewise 12 (24) are used per antenna if there are 2 transmitter antennas. The length of the HT-STS is 1.6µs. Since, the purpose of HT-STF is purely for AGC purposes, only one STS is needed. Together with the 0.8µs guard interval, the total length of the HT-STF is 2.4µs. The GI of the HT-STF, i.e. the first 0.8us, would contain echoes from the previous symbol. Hence, a clean measurement of power is made in the 1.6us portion of the HT-STF. It is anticipated that the RF transient settling after the AGC switch would occur in the GI of the HT-LTF that follows the HT-STF.

After the AGC has been "fine-tuned" for HT reception, channel estimation and fine frequency offset estimation are performed using the HT Long Training Symbols. The number of HT-LTFs is equal to the number of antennas in the basic MIMO mode. For example, with two transmit antennas, two HT-LTFs are sent. The HT-LTF is also tone interleaved to minimize the power fluctuation with respect to the AGC setting (obtained via the HT-STF). Therefore, one HT-LTF would cover half of the tones in the two antenna case. Remaining tones are covered in the second HT-LTF. The HT-LTF consists of two Long Training Symbols (LTS) as in 802.11a/g and a regular guard interval of 0.8 µs. The total length of one HT-LTF is 7.2µs.

11.1.1.5 Duplicate 6Mbps mode

The lowest data rate in the 40MHz mode would be 13.5Mbps (= 6Mbps x 2.25), if independent information bits are mapped to the upper and lower sub-channels. A 6Mbps mode in 40MHz is obtained by duplicating the data from the lower sub-channel to the upper sub-channel. For the 6Mbps mode in 40MHz, the MPDU is encoded, punctured, interleaved, and constellation-mapped as a 6Mbps transmission in 20MHz. The duplication is performed just prior to the iFFT sub-carrier mapping, as shown in Figure 51.

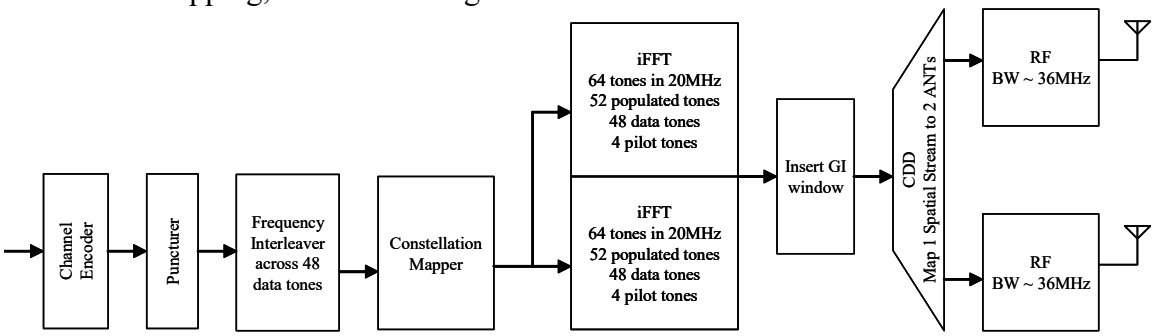


Figure 51: Transmitter configuration for 6Mbps in 40MHz[14]

Since the 6Mbps duplicate mode comprises of 1 spatial stream, only 1 HT-LTF is needed. Also note that there is no tone-fill past the HT-SIG, unlike the regular 40MHz PPDU format show in Figure 49. The PPDU format for the 6Mbps duplicate mode is shown in Figure 52:

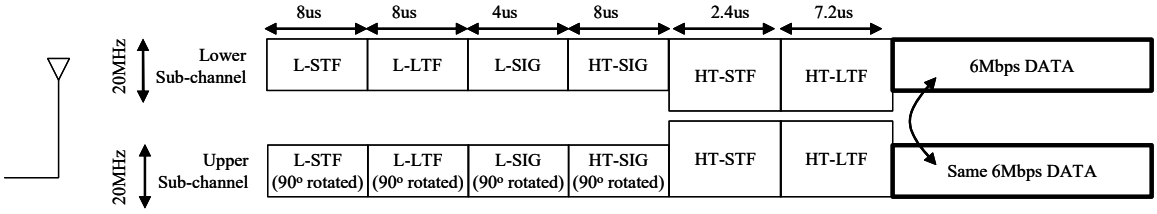


Figure 52: PPDU format for 6Mbps duplicate mode

Pilot locations in the duplicate mode follow the pilot locations in the 20MHz mode. The duplicate 6Mbps mode (in 40MHz) is signaled as mode 32 in the MCS set (see Table 26).

11.1.1.6 Support for Advanced Techniques

All advanced techniques are optional in our proposal.

In our basic mandatory 20/40MHz MIMO mode, the number of data streams created by the spatial parser in Figure 45 is the same as the number of transmit antennas. And, in this basic configuration, each spatial stream is mapped to exactly one transmit antenna.

If Channel State Information (CSI) is available at the transmitter, we can perform spatial "shaping" using a steering matrix for each sub-carrier, as shown in Figure 53. In SVD style MIMO, the steering matrix is obtained via singular value decomposition (SVD) of the physical channel transfer function matrix, *H*. In SVD-MIMO, each spatial stream is "transmitted" on a "singular vector" obtained from the SVD of *H*. We can further perform spatial water-filling by selecting an optimum power level, modulation level, and coding rate for each spatial stream.

Our proposal introduces two classes of transmit beamforming: basic beamforming with MIMO (BF-MIMO), and advanced beamforming with MIMO (ABF-MIMO). In basic transmit beamforming, only spatial steering is performed. Support for basic beamforming is mandatory in a receiver since there is minimal implemenation overhead. In advanced beamforming, each spatial stream may have a unique MCS and a unique power level. We also support bi-directional beamforming in the ABF-MIMO mode.

[14] 40MHz operation by an HT device shall only be permitted in those regulatory domain that allow it.

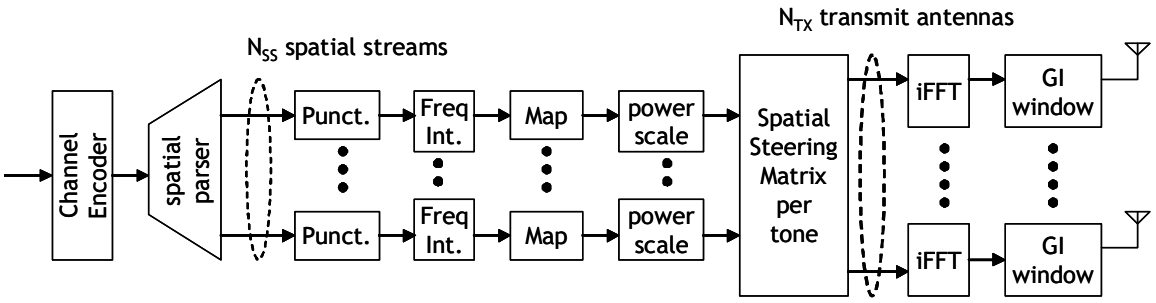


Figure 53: Transmitter Datapath with option to perform spatial shaping

We have also introduced an advanced coding option, i.e. LDPC coding. The use of LDPC coding is signaled in the HT-SIG.

11.1.1.7 Feedback Mechanisms

Our proposal consists of several feedback mechanisms, supported at the MAC layer, as shown in Table 21.

Nature of Feedback	Device Requirement	Type of Feedback	Description
Implicit Feedback	Mandatory in device	Binary Response packet	• This only provides information to the TX whether its choice of the modulation-coding mode was supported by the channel. • Does not provide any information to help perform spatial shaping • Used when transmitting to legacy devices
Implicit Feedback	Optional	Binary Response packet sent as a sounding packet	• Allows the Initiator to estimate the physical channel in the reverse link. If reciprocity can be assumed, and if the number of TX and RX antennas are identical at the recipient, the TX can use the derived CSI to perform spatial shaping (e.g. in the style of ABF-MIMO)
Explicit Feedback	Mandatory	RR (Rate Recommendation) Response packet	• The response packet contains a rate recommendation to the TX • Since spatial shaping at the initiator is not desired, HT-LTFs are not appended
Explicit Feedback	Mandatory	RR Response packet sent as a sounding packet	• The response packet contains a rate recommendation to the TX • Response packet does not contain forward link Channel State Information (CSIT) • However, the sounding packet allows reverse link channel estimation. And, assuming reciprocity and equal number of TX RX antennas at the recipient, the initiator can perform spatial shaping.

Table 21: Feedback Mechanisms

11.1.1.8 Comparison among MIMO transmission modes

Our proposal supports three classes of MIMO Transmissions Modes, as outlined in Table 22:

MIMO Modes	Characteristic	Extent of TX adaptation
Basic-MIMO (Mandatory)	• Typically, $N_{SS} = N_{TX}$ • Each spatial stream is mapped to one antenna	• No adaptation is performed in the spatial domain • MCS selection (adaptation) is based upon either: (1) binary feedback (i.e. ACK/noACK), or (2) recommendation from the RX, or (3) availability of CSI at the TX • All spatial streams shall use the same MCS and transmit at same power
Basic-MIMO with TX Beamforming (Optional)	• Typically, $N_{SS} \leq N_{TX}$ • All spatial streams are spatially shaped via a steering matrix • The steering matrix is defined for each tone in the OFDM modulation • All spatial streams have identical MCS and power levels	• CSI is required at the TX • Sounding packet from the recipient to the initiator is required to estimate the CSI at the transmitter • RF calibration is required since reciprocity is used to perform TX beamforming
Advanced MIMO with TX Beamforming (Optional)	• Typically, $N_{SS} \leq N_{TX}$ • All spatial streams are spatially shaped via a steering matrix • The steering matrix is defined for each tone in the OFDM modulation • Power levels for each spatial stream do not have to be identical • MCS for each spatial stream does not have to be identical • Bi-Directional Advanced Beamforming MIMO is supported via UD decomposition	• CSI is required at the TX • Sounding packet from the recipient to the initiator is required to estimate the CSI at the transmitter • RF calibration is required since reciprocity is used to perform TX beamforming

Table 22: MIMO Transmission options

11.1.1.9 "Modulation-Coding Scheme" (MCS) set for basic MIMO mode

Our proposal augments the 802.11a MCS set through the use of multiple spatial streams and bandwidth extension. The complete set is shown in Table 26. Our proposal recommends a mandatory data of 243Mbps using two spatial streams in regulatory domains that permit 40MHz operation. And, with an eye towards the future, our proposal supports scalability to 4 spatial streams, offering data rates in excess of 600Mbps.

11.1.2 Timing related parameters

Parameter	Value for 20 MHz Channel	Value for 40 MHz Channel
N_{SD} : Number of data subcarriers	48	108
N_{SP} : Number of pilot subcarriers	4	6
N_{SN} : Number of center null subcarriers	1 (tone = 0)	3 (tones = -1,0,+1)
N_{SR} : Subcarrier range (index range)	26 (-26 … +26)	58 (-58 … +58)
Δ_F : Subcarrier frequency spacing	0.3125 MHz (= 20 MHz / 64)	0.3125 MHz (= 40 MHz / 128)
T_{FFT} : IFFT/FFT period	3.2 μsec	3.2 μsec
T_{GI} : GI duration	0.8μsec	0.8μsec
$T_{ShortGI}$: Short GI duration	0.4μsec	0.4μsec
T_{GI2} : Legacy LongTraining symbol GI duration	1.6μsec	1.6μsec
T_{SYM} : Symbol interval	4μsec	4μsec
T_{LONG} : Long training field duration	8μsec	8μsec
$T_{HT-LONG}$:HT Long training field duration	7.2μsec	7.2μsec
T_{SHORT} : Short training field duration	8μsec	8μsec
$T_{HT-SHORT}$: HT Short training field duration	2.4μsec	2.4μsec
T_S : Nyquist sampling interval	50nsec	25nsec

Table 23: Timing Related Parameters

11.1.3 Mathematical conventions in signal descriptions

The transmitted signals will be described using the complex baseband signal notation. The transmitted signal at RF, $r_{RF}(t)$, is related to the complex baseband signal by the following relationship:

$$r_{RF}(t) = \mathrm{Re}\left\{ r(t)\exp\left(j2\pi f_c t\right)\right\} \quad (1)$$

where

$\mathrm{Re}\{\cdot\}$ represents the real part of a complex variable;

f_c denotes the center frequency of the carrier.

The transmitted baseband signal consists of several subframes. The timing boundaries for the various subframes are shown in Figure 54:

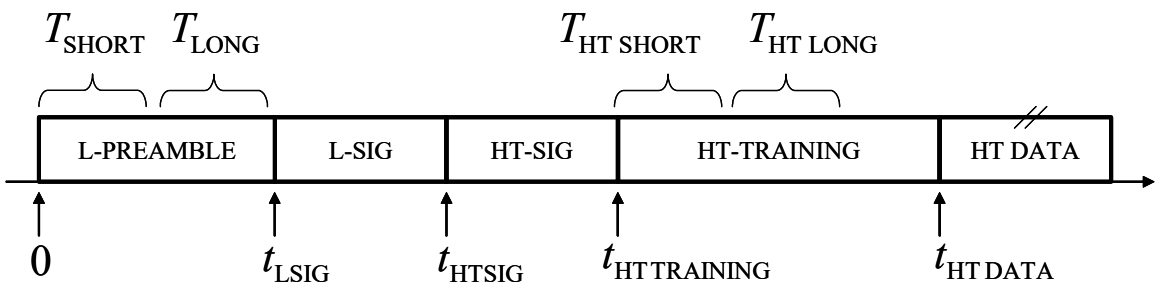


Figure 54: Timing boundaries for subframes in a HT transmission

The signal transmitted on antenna i_{Tx} is given by:

$$\begin{aligned} r_{PACKET}^{(i_{Tx})}(t) \;=\; & r_{PREAMBLE}^{(i_{Tx})}(t) + r_{LSIG}^{(i_{Tx})}(t - t_{LSIG}) \\ & + r_{HTSIG}^{(i_{Tx})}(t - t_{HTSIG}) \\ & + r_{HT\ TRAINING}^{(i_{Tx})}(t - t_{HT\ TRAINING}) \\ & + r_{HT\ DATA}^{(i_{Tx})}(t - t_{HT\ DATA}) \end{aligned} \qquad (2)$$

The subframes of (2) are described in 11.2.1.3.1, 11.2.1.3.2, 11.2.1.4.1, 11.2.1.4.2 and 11.2.1.5. The time offsets $t_{SUBFRAME}$ determines the starting time of the corresponding subframe. Per convention, it is assumed that the signal in equation (2) is causal, i.e. signal value is zero when its argument is zero.

All subframes of the signal are constructed using an inverse Fourier transform:

$$r_{SUBFRAME}^{(i_{Tx})}(t) \;=\; w_{SUBFRAME}(t) \sum_{k=-N_{SR}}^{N_{SR}} C_k \exp\left(j2\pi k \Delta_F t\right) \qquad (3)$$

The PHY related parameters are listed in Table 23.

11.2 PLCP sublayer

11.2.1 Basic MIMO Mode

As described in Table 22, the basic MIMO mode does not involve any spatial steering. The basic MIMO mode can operate in both 20MHz and 40MHz channel bandwidths.

11.2.1.1 Basic MIMO PPDU format

The HT Physical layer defines a PPDU format that provides interoperability and coexistence with 802.11a stations and with ERP-OFDM stations of 802.11g at the link level without protection mechanisms. Figure 55 shows the PPDU format for the basic MIMO mode:

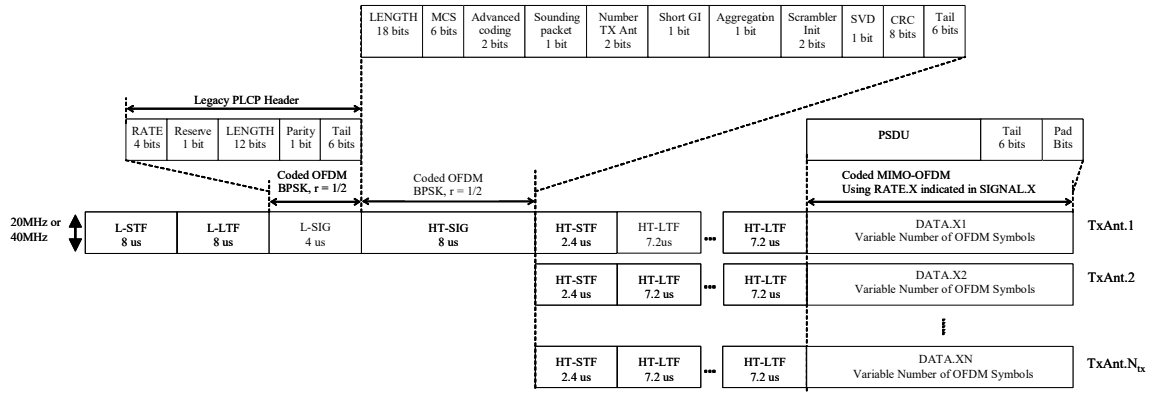


Figure 55: PPDU Format for N_{TX} antennas

L-STF is the Legacy Short Training Field. Details of the L-STF symbols can be found in clause 11.2.1.3.1. The HT-STF is specific for HT. Details can be found in 11.2.1.3.3

L-LTF stands for Legacy Long Training Field. Details of the L-LTF symbols can be found in clause 11.2.1.3.2**.** The HT-LTFs are specific for HT and the details can be found in clause 11.2.1.3.4

L-SIG is based on the Legacy SIGNAL field as defined in clause 17.3.4 in [1]. The modulation and coding procedure also follows the steps of 17.3.4 in [1]. The RATE and LENGTH fields in L-SIG are defined in accordance with "spoofing" procedures described in 11.2.1.4.1.

The HT-SIG field is the signal field dedicated to the HT mode. The structure of this field is the same as the L-SIG. Since our proposal supports interoperability between HT devices and legacy OFDM devices at the PHY layer, a HT receiver will not a-priori know whether a legacy packet is being received or a HT packet. That determination can only be made after the L-SIG. Our proposal supports two mechanisms for signaling a HT transmission: (a) send the HT-SIG as a BPSK signal on the quadrature axis, instead of on the in-phase axis, and (b) invert the polarity of the pilots in the HT-SIG with respect to the pilots in the L-SIG. These two mechanism allow an auto-detect of a HT transmission. Note that the pilots in the HT-SIG are not transmitted on the Quadrature axis. The definition of the contents of HT-SIG is specified in 11.2.1.4.2.

The L-STF, L-LTF, L-SIG, and HT-SIG should be transmitted from the antenna array such that an omni-directional beam pattern is created in the far field. This reduces the probability of creating hidden nodes in the BSS. One way to create an omni-beam is to transmit these fields from strictly one antenna. However, if the transmitter is power-amplifier limited, array gain can be derived by transmitting on multiple antennas. Since transmitting the same signal from multiple antennas invariably creates beamforming, one way to mitigate beamforming is to phase shift (i.e. time delay) the signal across the transmit antenna array. For an OFDM symbol, the delay can be introduced cyclically. Since delaying the same signal introduces frequency selectivity at the receiver, such a transmission is commonly referred to as "cyclical delay diversity" or CDD. This transmission format is shown in Figure 56:

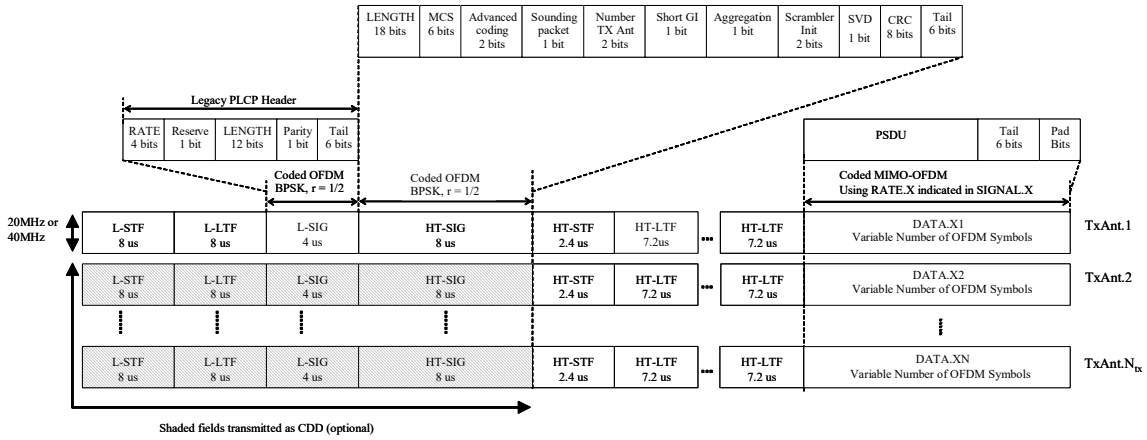


Figure 56: PPDU format leveraging CDD transmission

11.2.1.2 Legacy Protection

All MIMO transmissions (whether basic or advanced) are protected against legacy OFDM devices at the physical layer, and protected against legacy CCK/DSSS devices at the MAC layer. The legacy short training field, the legacy long training field and the legacy signal field are transmitted such that a legacy OFDM receiver (802.11g or 802.11a) is able to "understand" these fields. In the basic MIMO mode, the transmitter selects the Modulation Coding Scheme either based on conventional rate selection heuristics or via implicit or explicit feedback. Thus, the RATE and LENGTH fields in the legacy signal field can be "spoofed" to allow the legacy OFDM receivers to correctly set their NAV.

RTS/CTS or CTS-to-self mechanism may be used to protect against legacy CCK/DSSS devices in an 802.11n BSS.

11.2.1.3 Training symbols specification

The training fields are used for synchronization purposes. As shown in Figure 55 and Figure 56, at the start of a packet, a legacy short training field is transmitted .The legacy short training field (L-STF) contains ten identical short training symbols. The legacy short training field may be transmitted from a single antenna, or from multiple antennas using CDD (refer to the next section).The legacy long training field follows the legacy short training field. Mapping of the legacy long training field (L-LTF), the L-SIG and the HT-SIG to the antenna array shall be identical to that of the legacy short training field (L-STF). After the HT-SIG, a HT-STF is transmitted across all the antennas to "fine tune" the AGC for MIMO reception. The HT-LTFs subsequently follow the HT-STF. The number of HT LTFs is equal to the number of spatial streams (which equals the number of transmit antennas in the basic MIMO mode). Data transmission follows the HT-LTFs.

11.2.1.3.1 Legacy Short Training Field

The legacy short training symbol in our proposal is identical to the 802.11a short training symbol. The legacy short training OFDM symbol in the 20MHz mode is:

$$S_{-26,26} = \sqrt{13/6}\{0,0,1+j,0,0,0,-1-j,0,0,0,1+j,0,0,0,-1-j,0,0,0,-1-j,0,0,0,1+j,0,0,0,$$
$$0,0,0,0,-1-j,0,0,0,-1-j,0,0,0,1+j,0,0,0,1+j,0,0,0,1+j,0,0,0,1+j,0,0\}. \quad (4)$$

The normalization factor accounts for the fact that 12 tones out of 52 are used in the STF, and also because QPSK is used as the constellation on the 12 tones.
The legacy short training OFDM symbol in 40MHz mode is:

$$S_{-58,58} = \sqrt{13/6}\{0,0,1+j,0,0,0,-1-j,0,0,0,1+j,0,0,0,-1-j,0,0,0,-1-j,0,0,0,1+j,0,0,0,$$
$$0,0,0,0,-1-j,0,0,0,-1-j,0,0,0,1+j,0,0,0,1+j,0,0,0,1+j,0,0,0,1+j,0,0,0,0,0,$$
$$0,0,0,0,0,0,0,0,0,0,0,j-1,0,0,0,-j+1,0,0,0,j-1,0,0,0,-j+1,0,0,0,-j+1,0,0,0,j-1,$$
$$0,0,0,0,0,0,0,-j+1,0,0,0,-j+1,0,0,0,j-1,0,0,0,j-1,0,0,0,j-1,0,0,0,j-1,0,0\} \quad (5)$$

The tones in the upper subchannel (see Figure 48) are obtained by phase rotating the 802.11a short OFDM training symbol tones by 90 degrees, whereas the tones in the lower subchannel are identical to the 802.11a short training OFDM symbol tones. The 90º rotation helps keep the PAPR of the STF in 40MHz comparable to that in 20MHz. A scaling factor of $\sqrt{13/6}$ ensures that the legacy short training symbols have the same power as the L-LTF, L-SIG, and HT-SIG fields, all of which have 104 tones in 40MHz. There are 24 tones in the L-STF (with QPSK constellation) in 40MHz.

Let N_{TX} be the total number of transmit antennas. We normalize the transmit power across N_{TX} transmit antennas to be N_{TX}.

The short training symbols can be mapped to the transmit antenna array in several ways. L-STF should be transmitted omni-directionally to avoid beamforming, else hidden nodes can be potentially created in the BSS. One simple way is to transmit the L-STF on just one antenna. The L-STF is then described as follows:

$$r_{SHORT}^{(i_{TX})}(t) = \begin{cases} \sqrt{N_{TX}}\, w_{T\,SHORT}(t) \sum_{k=-N_{SR}}^{N_{SR}} S_k \exp\{j2\pi k\Delta_F t\} & i_{Tx} = 1 \\ 0 & otherwise \end{cases} \quad (6)$$

Since the legacy short training field (L-STF) constitutes a single spatial stream, it is also possible to map L-STF to multiple transmit antennas. This may be particularly useful if the transmit power is limited by the power amplifier characteristics. However, mapping 1 spatial stream to multiple transmit antennas would invariably result in beamforming, which can be mitigated by introducing a linear phase shift across the transmit antenna array. In the case of OFDM, the delay is introduced cyclically, and is commonly knows as Cyclical Delay Diversity (CDD). When the STF is transmitted in a CDD fashion, the short training signal at the i_{TX} th transmit antenna is generated according to the following rule:

$$r_{SHORT}^{(i_{TX})}(t) = w_{TSHORT}(t) \sum_{k=-N_{SR}}^{N_{SR}} S_k \exp\left\{j2\pi k\Delta_F \left(t - (i_{TX}-1)T_{shift}\right)\right\} \quad (7)$$

The short training signal is repeated every 0.8μs. The entire short training field includes ten such periods, with a total duration of 8μs. The short training field at the i_{TX} th transmit antenna is cyclically shifted by $(i_{TX}-1)T_{shift}$, where T_{shift} is defined such that $(N_{TX}-1)\cdot T_{shift} = 50ns$ for both 20MHz and 40MHz. For example, with 4 transmit antennas, $T_{shift} = 16.67ns$

11.2.1.3.2 Legacy Long Training Field

The legacy long training OFDM symbol is identical to the 802.11a long training OFDM symbol. In the 20MHz mode, the long training OFDM symbol is given by

$$L_{-26,26} = \{1,1,-1,-1,1,1,-1,1,-1,1,1,1,1,1,1,-1,-1,1,1,-1,1,-1,1,1,1,1,1,0,$$
$$1,-1,-1,1,1,-1,1,-1,1,-1,-1,-1,-1,-1,1,1,1,-1,-1,1,-1,1,-1,1,1,1,1\} \quad (8)$$

Similar to the L-STF, the legacy long training field (L-LTF) can be mapped to either a single antenna, or to multiple antennas via CDD. However, to guarantee AGC accuracy, the antenna mapping for the L-LTF should be identical to that of the L-STF. If only one antenna is used to transmit the legacy long training field, then the long training signal is given by:

$$r_{LONG}^{(iTX)}(t) = \begin{cases} \sqrt{N_{TX}}\, w_{T\,LONG}(t) \sum_{k=-N_{SR}}^{N_{SR}} L_k \exp\left(j2\pi k \Delta_F \left(t - T_{G12}\right)\right) & i_{Tx} = 1 \\ 0 & otherwise \end{cases} \quad (9)$$

where $T_{G12} = 1.6\mu s$.

If the long training signal is generated in a CDD fashion, we get:

$$r_{LONG}(t) = w_{T\,LONG}(t) \sum_{k=-N_{SR}}^{N_{SR}} L_k \exp\left\{j2\pi k \Delta_F \left(t - \left(i_{TX} - 1\right) T_{Shift} - T_{G12}\right)\right\} \quad (10)$$

There are two legacy long training symbols in the legacy long training field. This is to ensure channel estimation accuracy, and also to permit fine frequency offset estimation. The total duration of the long training field (L-LTF) is $T_{LONG} = 1.6 + 2 \times 3.2 = 8\mu s$.
The legacy long training OFDM signal in the 40MHz mode is given by:

$$L_{-58,58} = \{1,1,-1,-1,1,1,-1,1,-1,1,1,1,1,1,1,-1,-1,1,1,-1,1,-1,1,1,1,1,1,0,$$
$$1,-1,-1,1,1,-1,1,-1,1,-1,-1,-1,-1,-1,1,1,1,-1,-1,1,-1,1,-1,1,1,1,1,1,0,0,0,0,0,0$$
$$0,0,0,0,0,0,0,j,j,-j,-j,j,j,-j,j,-j,j,j,j,j,j,j,-j,-j,j,j,-j,j,-j,j,j,j,j,0,$$
$$j,-j,-j,j,j,-j,j,-j,j,-j,-j,-j,-j,-j,j,j,-j,-j,j,-j,j,-j,j,-j,j,j,j\} \quad (11)$$

Tones in the upper subchannel (subcarriers +6 to +58) are obtained by phase rotating the 802.11a long OFDM training symbol tones by 90 degrees, whereas the tones within the lower subchannel (subcarriers -58 to -6) are identical to the 802.11a long OFDM training symbol tones.
It should be noted that neither the legacy fields (L-STF, L-LTF, LSIG) nor the HT-SIG undergo any phase rotation in the lower subchannel.

The subcarriers at ±32 in 40MHz, which are the DC subcarriers for the legacy 20MHz transmission, are both nulled in the L-LTF. Such an arrangement allows proper synchronization of the 20MHz legacy device.

The legacy long training field together with the short training field constitutes the legacy PLCP preamble. The legacy device can perform the necessary synchronization to decode the legacy signal field. Hence, we can write $r_{PREAMBLE}(t)$ in (2) as

$$r_{PREAMBLE}^{(iTX)}(t) = r_{SHORT}^{(iTX)}(t) + r_{LONG}^{(iTX)}(t - T_{SHORT}) \quad (12)$$

11.2.1.3.3 HT-Short Training Field

The HT-STF follows the HT signal field (see Figure 55), whose only purpose it to "fine tune" the AGC for HT MIMO reception. Note that the first AGC operation is performed on the L-STF. The HT short training symbol in the 20MHz mode is

$$HTS_{-26,26} = \sqrt{13/12}\begin{cases} \{0, 0, 1+j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 0\} & N_{Tx}=1 \\ \{0, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, \\ 0, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, \} & N_{Tx}=2 \\ \{0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, 0\} & N_{Tx}=3 \\ \{0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 0\} & N_{Tx}=4 \end{cases} \quad (13)$$

Unlike the legacy short training symbol which uses 12 tones, the HT-Short training Symbol uses 24 tones in 20MHz which are tone-interleaved across the transmit antennas. For example, with $N_{TX}=4$, each antenna would transmit 6 tones, which provide sufficient accuracy for estimating the received power from one transmit-antenna to one receive-antenna.

The HT short training symbol in the 40 MHz mode consists of 48 tones, and is given by:

$$HTS_{-58,58} = \sqrt{19/16}\begin{cases} \{0, 0, 1+j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, \\ 0, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, 0, 0, 0, 0, 0, 0, 0, \\ 0, 0, 0, 0, 0, 0, 0, 0, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, 1+j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 0\} & N_{Tx}=1 \\ \{0, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, \\ 0, 0, 1+j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, 0, 0, 0, 0, 0, 0, 0, \\ 0, 0, 0, 0, 0, 0, 0, 0, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, 1+j, 0, \\ 0, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 0\} & N_{Tx}=2 \\ \{0, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, 0, 0, 0, 0, 0, 0, 0, \\ 0, 0, 0, 0, 0, 0, 0, 0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 1+j, 0, -1-j, 0, -1-j, 0, -1-j, 0, 0\} & N_{Tx}=3 \\ \{0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, 0, 0, 0, 0, 0, 0, 0, \\ 0, 0, 0, 0, 0, 0, 0, 0, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, \\ 0, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 1+j, 0, 1+j, 0, 1+j, 0, 1+j, 0, -1-j, 0, -1-j, 0, 0\} & N_{Tx}=4 \end{cases} \quad (14)$$

In the duplicate 6Mbps mode in 40MHz, the HT-STF shown in (14) is used, except for a scaling factor change. The scaling factor for the 6Mbps duplicate mode is $\sqrt{13/12}$.

The HT short training signal is then generated by the following rule.

$$r^{(i_{Tx})}_{MIMOSHORT}(t) = w_{TMIMOSHORT}(t)\sqrt{N_{TX}}\sum_{m}\sum_{k=0}^{\left\lfloor \frac{11}{N_{TX}} \right\rfloor} MS_{2N_{TX}k+2(i_{Tx}-1)+m}\exp\{j2\pi(2N_{TX}k+2(i_{Tx}-1)+m)\Delta_F(t-T_{GI})\} \quad (15)$$

where $T_{GI} = 0.8\mu s$. The elements of m are $m \in \{-24, 2\}$ for 20MHz and $m \in \{-54, -30, 8, 34\}$ for 40MHz, respectively. The HT short training signal has a period of 1.6μs – corresponding to 24 tones in 20MHz, and 48 tones in 40MHz – which is deemed sufficient for power measurement purposes. Since the HT-STF is tone-interleaved at the transmitter, the power periodicity of the HT-STF is also $1.6\mu s$.

11.2.1.3.4 HT-Long Training Field

The HT long training fields are transmitted after the HT-STF. The number of HT long training fields is equal to the number of spatial streams. In the basic MIMO mode, the number of spatial streams is equal to the number of transmit antennas.

To minimize the power fluctuation between the HT-LTF and the AGC setting (which is derived from the HT-STF), the HT long training symbols are also tone-interleaved across the transmit antenna array. Since each antenna only transmits a sub-set of the tones, the HT long training sequences have been optimized to achieve the lowest peak-to-average power ratio (PAPR). There are a total of 9 sequences:

Table 24: List of HT Long Training Fields

Sequence Number	N_{TX}	BW	*Comment*
1	1	20MHz	*Same as 802.11a*
2	2	20MHz	*52 tones*
3	3	20MHz	*52 tones*
4	4	20MHz	*52 tones*
5	1	40MHz	*Duplicate format, 6Mbps in 40MHz (104 Tones)*
6	1	40MHz	*114 tones*
7	2	40MHz	*114 tones*
8	3	40MHz	*114 tones*
9	4	40MHz	*114 tones*

In 20MHz, for $N_{TX} = 1$, sequence 1 is identical to the 802.11a LT sequence, as shown in (8). In 20MHz, for $N_{TX} = 2$, sequence 2 is:

$$HTL_{-26,+26} = \{-1, 1, -1, 1, 1, 1, -1, -1, -1, -1, 1, 1, 1, 1, -1, 1, -1, 1, -1, -1, 1, 1, 1, -1, 1, 1, 0, -1, 1, 1, -1, -1, 1, -1, -1, 1, -1, -1, 1, -1, -1, 1, 1, 1, 1, -1, 1, 1, 1, 1, 1, 1\} \quad (16)$$

The tone-interleaving rule across the two antennas in 20MHz is grouped into 2 sets:

$$set_1 = [-26:2:-2], [+2:2:+26]$$
$$set_2 = [-25:2:-1], [+1:2:+25] \quad (17)$$

where the notation in the brackets should be read as [starting index : step value : ending index].

In 20MHz, for $N_{TX} = 3$, sequence 3 is:

$$HTL_{-26,+26} = \{\text{-1 -1 1 1 1 1 1 -1 -1 -1 1 -1 -1 -1 -1 -1 1 1 1 1 1 1 1 1 -1 -1 1}$$
$$\text{0 1 -1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 -1 1 1 -1 1 1 -1 1 -1 -1 1}\} \quad (18)$$

The tone-interleaving rule across three antennas in 20MHz is grouped into 3 sets:

$$set_1 = [-26:3:-2], [2:3:26]$$
$$set_2 = [-25:3:-1], [3:3:24] \quad (19)$$
$$set_3 = [-24:3:-3], [1:3:25]$$

In 20MHz, for $N_{TX} = 4$, sequence 4 is:

$$HTL_{-26,+26} = \{\text{-1 1 1 1 1 -1 -1 -1 1 -1 1 1 1 -1 1 1 -1 1 -1 -1 -1 1 1 1 -1 1}$$
$$\text{0 1 1 -1 1 -1 -1 1 -1 -1 -1 -1 1 -1 -1 -1 1 1 1 1 -1 1 1 -1 1 1 1}\} \quad (20)$$

The tone-interleaving rule across four antennas in 20MHz is grouped into 4 sets:

$$set_1 = [-26:4:-2], [3:4:23]$$
$$set_2 = [-25:4:-1], [4:4:24]$$
$$set_3 = [-24:4:-4], [1:4:25] \quad (21)$$
$$set_4 = [-23:4:-3], [2:4:26]$$

In the 40MHz mode, the subcarriers at ±32 are populated in the HT long training OFDM symbol. In 40MHz, for $N_{TX} = 1$, sequence 5 is identical to sequence 6. Sequence 5 is used in the duplicate 6Mbps mode, whereby only 104 tones are needed, whereas sequence 6 corresponds to the 114 tones. Note that all 114 tones are transmitted in the 6Mbps duplicate mode.

$$HTL_{-58,+58} = \{\text{-1 1 1 1 1 -1 1 1 -1 -1 -1 -1 1 1 1 1 1 -1 1 1 -1 1 -1 1 1 -1 -1}$$
$$\text{1 1 -1 -1 1 1 -1 -1 -1 -1 -1 1 -1 1 1 -1 -1 -1 1 -1 -1 1 -1 -1 1 1 -1 1 1 1}$$
$$\text{0 0 0 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 -1 1 -1 1 1 -1 -1 1-1 -1 1 -1 1}$$
$$\text{-1 1 -1 -1 -1 -1 1 1 -1 1 -1 -1 -1 1 -1 1 -1 1 -1 1 1 -1 1 -1 -1 1 -1 1 1 1}\} \quad (22)$$

In 40MHz, for $N_{TX} = 2$, sequence 7 is:

$$HTL_{-58,+58} = \{\text{-1 1 1 1 1 -1 1 1 -1 -1 -1 -1 1 1 -1 1 1 -1 1 1 -1 1 -1 1 1 -1}$$
$$\text{-1 1 1 -1 -1 -1 1 1 -1 -1 -1 1 -1 1 -1 1 1 1 -1 -1 1 -1 -1 -1 -1 -1 -1 1 -1 1 -1 1}$$
$$\text{0 0 0 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 -1 1 1 1 -1 1 -1 1 1 -1 -1 1 -1 1 1 1 -1 1 -1}$$
$$\text{1 -1 -1 -1 -1 1 1 -1 1 -1 -1 -1 -1 -1 1 -1 1 1 1 1 -1 1 1 -1 1 -1 1 1 1}\} \quad (23)$$

The tone-interleaving rule across two antennas in 40MHz is grouped into 2 sets:

$$set_1 = [-58:2:2], [2:2:58]$$
$$set_2 = [-57:2:-3], [3:2:57] \quad (24)$$

The tone-interleaving pattern in space and time for the HT long training fields is shown in Figure 57 for 2 antennas (2 spatial streams):

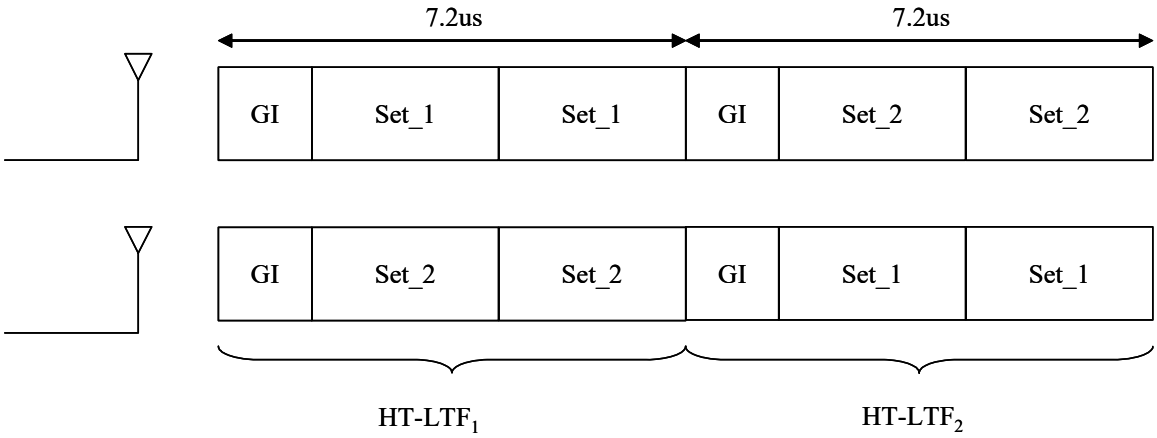


Figure 57: HT-LTF pattern for 2 antennas (2 spatial streams)

In 40MHz, for $N_{TX} = 3$, sequence 8 is:

$$HTL_{-58,+58} = \{-1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ -1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ -1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ 1$$
$$1\ 1\ 1\ 1\ 1\ -1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ 1\ 1\ 1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ 0 \quad (25)$$
$$0\ 0\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ -1\ -1\ -1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ -1\ -1$$
$$-1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ -1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ 1\ -1\ -1\ -1\}$$

The tone-interleaving rule across 3 antennas in 40MHz is grouped into 3 sets:

$$set_1 = [-58:3:-4], [2:3:56]$$
$$set_2 = [-57:3:-3], [3:3:57] \quad (26)$$
$$set_3 = [-56:3:-2], [4:3:58]$$

The tone-interleaving pattern in space and time for the HT long training fields is shown in Figure 58 for 3 antennas (3 spatial streams):

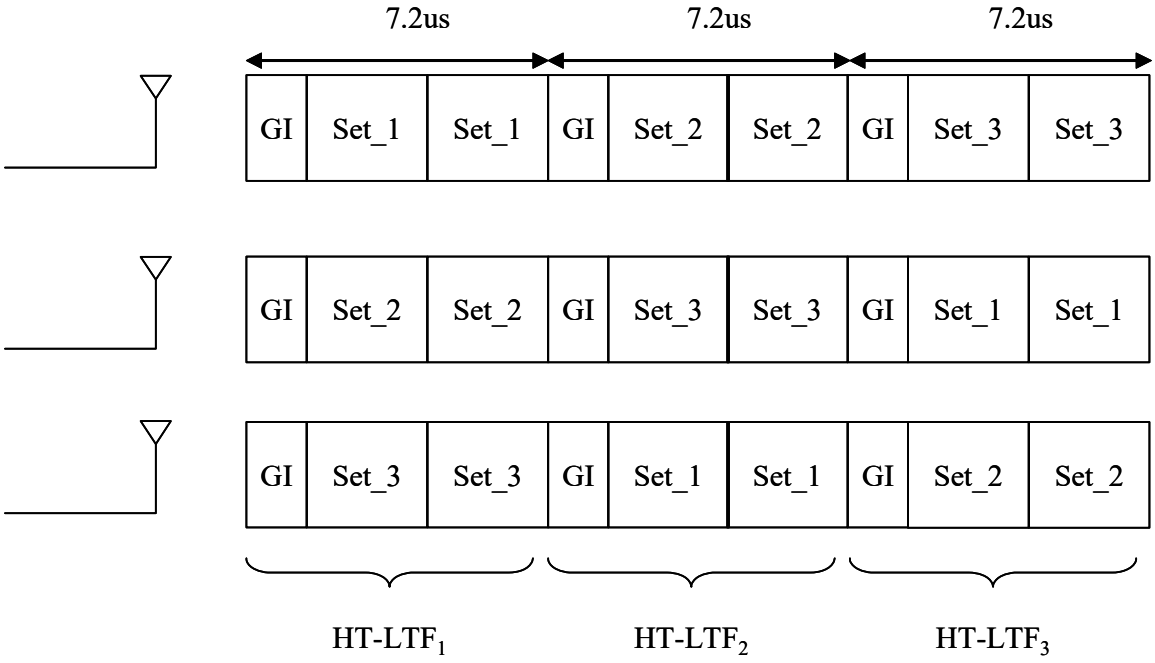


Figure 58: HT-LTF pattern for 3 antennas (3 spatial streams)

In 40MHz, for $N_{TX} = 4$, sequence 9 is:

$$\begin{aligned} HTL_{-58,+58} = \{&\text{-1 1 -1 -1 -1 1 1 -1 1 1 1 -1 1 1 1 1 -1 -1 -1 -1 1 -1 -1 -1 1 -1 1 -1} \\ &\text{1 -1 -1 -1 -1 1 1 1 -1 -1 -1 1 1 1 1 -1 1 -1 -1 1 -1 1 1 1 1 1 -1 -1 -1 1 0} \\ &\text{0 0 -1 1 1 -1 -1 -1 -1 -1 1 1 1 -1 1 1 -1 -1 -1 1 -1 1 -1 -1 -1 1 -1 1 -1} \\ &\text{-1 1 1 1 1 1 1 1 -1 -1 -1 -1 1 1 -1 1 1 1 -1 1 1 -1 -1 1 -1 1 -1 -1 1 1}\} \end{aligned} \quad (27)$$

The tone-interleaving rule across 4 antennas in 40MHz is grouped into 4 sets:

$$\begin{aligned} set_1 &= [-58:4:2], [5:4:57] \\ set_2 &= [-57:4:-5], [2:4:58] \\ set_3 &= [-56:4:-4], [3:4:55] \\ set_4 &= [-55:4:-3], [4:4:56] \end{aligned} \quad (28)$$

The tone-interleaving pattern in space and time for the HT long training fields is shown in Figure 59 for 4 antennas (4 spatial streams):

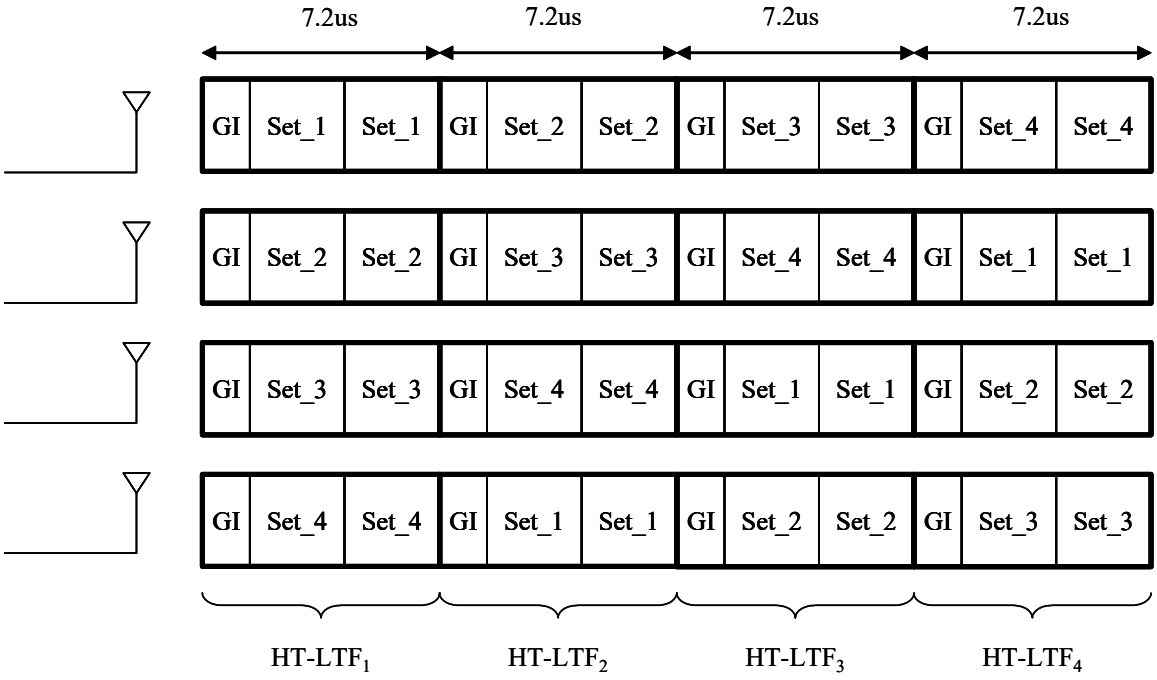


Figure 59: HT-LTF pattern for 4 antennas (4 spatial streams)

The HT long training signal is generated similarly to the HT short training signal.

$$r_{HT\,LONG}^{(i_{Tx})}(t) = w_{T\,HT\,LONG}(t)\sqrt{N_{TX}} \sum_{k\in set_(1+[(i_{TX}-1)+(n-1)]\bmod N_{TX})} HTL_k \exp\left\{j2\pi k\Delta_F\left(t - T_{GI} - (n-1)T_{HT\,LONG}\right)\right\}$$
$$for \quad (n-1)T_{HT\,LONG} \le t < nT_{HT\,LONG} \quad and \quad n = 1,\cdots,N_{TX} \qquad (29)$$
$$for \quad i_{TX} \in \{1,\ldots,N_{TX}\}$$

One HT long training field consists of two such long training symbol plus a regular 0.8μs guard interval. The total length of one long training field is 7.2μs, i.e., $T_{HT\,LONG} = 7.2\mu s$. Since only part of tones are transmitted on each antenna at any given long training symbol, it takes $N_{TX} \cdot T_{HT\,LONG}$ for all the antennas to cover all the tones. The duration of the HT-LTF is $N_{TX} \cdot T_{HT\,LONG}$.

The HT training fields in (2) for purposes of MIMO training can now be expressed as:

$$r_{HT\,TRAINING}^{(i_{Tx})}(t) = r_{HT\,SHORT}^{(i_{Tx})}(t) + r_{HT\,LONG}^{(i_{Tx})}\left(t - T_{HT\,SHORT}\right) \qquad (30)$$

If TX beamforming is performed on 1 spatial stream, i.e. $N_{ss} = 1$, *and* $N_{TX} > 1$, only one HT-LTF is needed after the HT-STF, and this HT-LTF is identical to the 802.11a sequence, as shown in (8). We should point out that the legacy part of the HT preamble is never beamformed, i.e. L-STF, L-HTF, L-SIG, and HT-SIG. And, the HT-STF and the HT-LTF are shaped by the same steering matrix as the HT-DATA, such that channel estimation is independent of the choice of the steering matrix used at the transmitter.

We should point out that the above discussion can be generalized for $N_{SS} \ge 2$ and $N_{TX} > N_{SS}$. This case corresponds to TX beamforming with multiple spatial streams, whereby the number of HT-LTFs is equal to NSS. And, as mentioned above, in addition to tone-interleaving, the HT-LTF also undergoes the same spatial steering as the HT-DATA fields.

11.2.1.4 Signal fields specification

The legacy long training field L-LTF shall be followed by the signal fields LSIG and HTSIG. The coding and modulation follow the steps of clause 17.3.4 in ref[2]: BPSK modulation is used for the L-SIG and HT-SIG, with code rate ½ resulting in 6 Mbit/s. Between the LSIG and the HTSIG symbol, the BPSK modulation shall be rotated counterclockwise by 90 degrees. This phase rotation can be used to "auto-detect" whether the frame being received is legacy or HT. The constellation bit encoding is depicted in Figure 60.

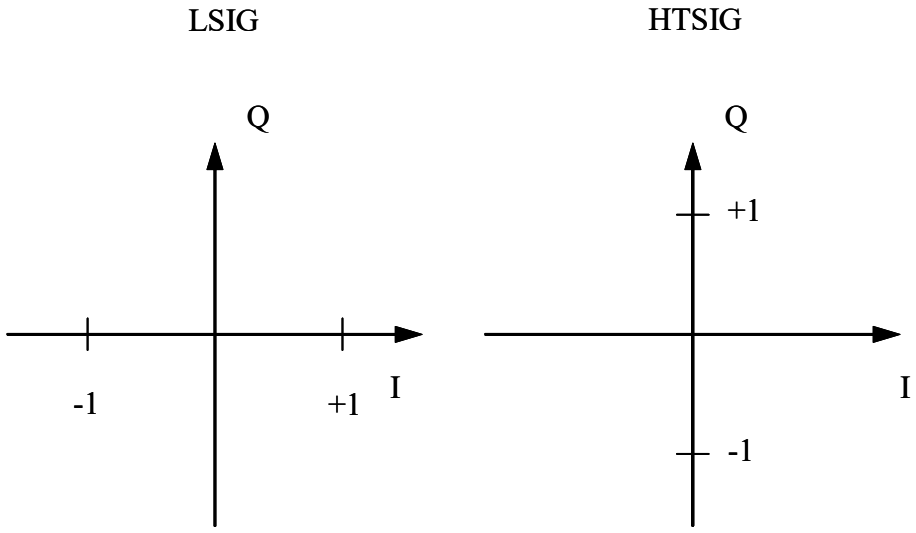


Figure 60: Constellation specification for LSIG and HTSIG

11.2.1.4.1 LSIG

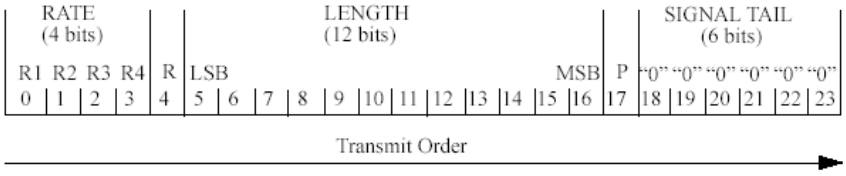


Figure 61: SIGNAL field bit assignment

In the 20MHz mode, the format of the LSIG field shall be the same as the SIGNAL field as defined in figure 111 in 17.3.4 in [1]. The combination of RATE and LENGTH is used to signal a station (OFDM mode only) the duration of a HT transmission. This duration does not have to be the actual duration of the HT frame. It can be used to reserve the medium beyond the actual transmission. The reservation of the medium through this kind of signaling is called "spoofing".

To prevent a legacy station from accessing the medium during a HT frame transmission, the "duration" calculated from the RATE and LENGTH fields shall be at least as large as the actual duration of the HT frame, which is measured from the end of the L-SIG field to the end of the DATA field. During an HT transmission, the RATE field is set to 6Mbps, whereas the LENGTH field is set based on the "duration" value.

In the 40MHz mode, the legacy signal field is duplicated such that tones are filled from -58 to -6 (lower-subchannel), and from +6 to +58 (upper-subchannel). The upper subchannel shall be phase rotated by 90 degrees, in the same way as the LTS as described in 11.2.1.3.2

The LSIG shall be transmitted in the same fashion as the legacy short and the legacy long training fields. If L-STF and L-LTF are transmitted from a single antenna, we index that antenna as $i_{Tx} = 1$ to be consistent with 11.2.1.3.2. Denote $LSIG_{-26,26}$ and $LSIG_{-58,58}$ as the LSIG OFDM symbol generated for 20 MHz and 40 MHz respectively. Then, for transmission over a single antenna, we have,

$$r_{LSIG}^{i_{Tx}}(t) = \begin{cases} \sqrt{N_{Tx}}\, w_{SYM}(t) \sum_{k=-N_{SR}}^{N_{SR}} LSIG_k \exp\left\{ j2\pi k \Delta_F \left(t - T_{GI} \right) \right\} & i_{Tx} = 1 \\ 0 & otherwise \end{cases} \quad (31)$$

Similarly, if the legacy training fields are transmitted across multiple transmit antennas using CDD, we have

$$r_{LSIG}^{(i_{Tx})} = w_{SYM}(t) \sum_{k=-N_{SR}}^{N_{SR}} LSIG_k \exp\left\{ j2\pi k \Delta_F \left(t - T_{GI} - \left(i_{Tx} - 1 \right) T_{shift} \right) \right\} \quad (32)$$

11.2.1.4.1.1 Data rate (RATE)

For a legacy station (11a or 11g), RATE has the original meaning as is defined in Table 80 of section 17.3.4.1 in [1]. For reference the table is repeated here:

Rate (Mbits/s)	R1–R4
6	1101
9	1111
12	0101
18	0111
24	1001
36	1011
48	0001
54	0011

Table 25: Content of the legacy SIGNAL field

11.2.1.4.1.2 PLCP length field (LENGTH)

The PLCP length field in HT is used to indicate to both legacy and HT stations a protection duration. It contains a value, which when combined with the rate field defines a duration. The number of bytes in the LENGTH field is calculated based on the perceived RATE of a legacy station defined in R1-R4 and the actual number of bytes that the MAC is currently requesting the

PHY to transmit, taking into account the number of transmit antennas (spatial streams), the bandwidth and the extra fields (symbols) in the preamble.

As indicated in 11.2.1.4.1, the calculated duration is allowed to be longer than the actual duration of the HT transmission. However the LENGTH value should be equal to or smaller than 2346 bytes to avoid mis-interpretation by existing legacy devices.

11.2.1.4.2 HTSIG

The LSIG field shall be followed by the HTSIG, which serves as the signal field for the HT PHY. HTSIG consists of two OFDM symbols: HTSIG1 and HTSIG2. HTSIG1 shall be transmitted first in time.

The encoding of HTSIG1 and HTSIG2 shall be the same as LSIG, with the exception that the BPSK modulation is rotated 90 degrees as described in clause 11.2.1.4.

In the 40MHz mode, the HTSIG1 and HTSIG2 are duplicated. The upper sub-channel shall be phase rotated by 90 degrees in the same way as the L-LTF, as described in 11.2.1.3.2 Figure 62 shows the bit assignment for the two symbols. They will be described in the next clauses.

HTSIG1 and HTSIG2 shall be transmitted in a similar fashion as the LSIG. Denote $HTSIG^{1,2}_{-26,26}$ and $HTSIG^{1,2}_{-58,58}$ as the first and second OFDM symbol of the following defined HTSIG as a data field. Then:

$$r^{(i_{Tx})}_{HTSIG}(t) = \begin{cases} \sqrt{N_{Tx}} \sum_{n=1}^{2} \left[w_{SYM}(t) \sum_{k=-N_{SR}}^{N_{SR}} HTSIG^{n}_{k} \exp\left\{ j2\pi k\Delta_F \left(t - T_{GI} - (n-1)T_{SYM} \right) \right\} \right] & i_{Tx} = 1 \\ 0 & otherwise \end{cases} \quad (33)$$

if the legacy training fields are sent only on the first antenna. And, if CDD is used to transmit the legacy signal fields, then

$$r^{(i_{Tx})}_{HTSIG}(t) = \sum_{n=1}^{2} \left[w_{SYM}(t) \sum_{k=-N_{SR}}^{N_{SR}} HTSIG^{n}_{k} \exp\left\{ j2\pi k\Delta_F \left(t - T_{GI} - (n-1)T_{SYM} - (i_{Tx}-1)T_{shift} \right) \right\} \right] \quad (34)$$

where $T_{SYM} = 4\mu s$ is the OFDM symbol duration.

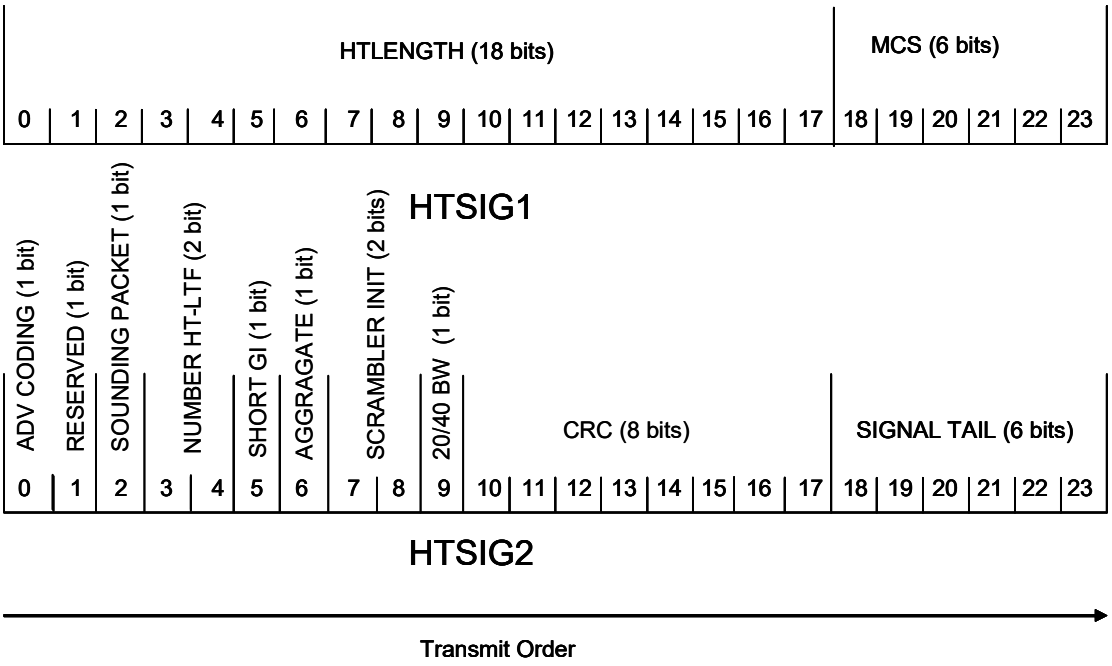


Figure 62: HT SIGNAL FIELD (HTSIG1 and HTSIG2) bit assignment

11.2.1.4.2.1 HTLENGTH field

The HTLENGTH field shall be an unsigned 18-bit integer that indicates the number of octets in the PSDU that the MAC is currently requesting the PHY to transmit. This value is used by the PHY to determine the number of octets that will occur between the MAC and the PHY after receiving a request to start transmission.

Bit 0 is the LSB (least significant bit) and shall be transmitted first in time to the channel encoder.

11.2.1.4.2.2 MCS field

The MCS field is a 6-bit field (bit 18-23 of HTSIG1) and defines the modulation and coding scheme, as indicated in Table 26.

Bits 18-23 in HT-SIG1 (MCS index)	Number of spatial streams	Modulation	Coding rate	GI = 800ns		GI = 400ns	
				Rate in 20MHz	Rate in 40MHz	Rate in 20MHz	Rate in 40MHz
0	1	BPSK	1/2	6	13.5	6.67	15
1	1	QPSK	1/2	12	27	13.33	30
2	1	QPSK	3/4	18	40.5	20	45
3	1	16-QAM	1/2	24	54	26.67	60
4	1	16-QAM	3/4	36	81	40	90
5	1	64-QAM	2/3	48	108	53.33	120
6	1	64-QAM	3/4	54	121.5	60	135
7	1	64-QAM	7/8	63	141.75	70	157.5

8	2	BPSK	1/2	12	27	13.33	30
9	2	QPSK	1/2	24	54	26.67	60
10	2	QPSK	3/4	36	81	40	90
11	2	16-QAM	1/2	48	108	53.33	120
12	2	16-QAM	3/4	72	162	80	180
13	2	64-QAM	2/3	96	216	106.67	240
14	2	64-QAM	3/4	108	243	120	270
15	2	64-QAM	7/8	126	283.5	140	315
16	3	BPSK	1/2	18	40.5	20	45
17	3	QPSK	1/2	36	81	40	90
18	3	QPSK	3/4	54	121.5	60	135
19	3	16-QAM	1/2	72	162	80	180
20	3	16-QAM	3/4	108	243	120	270
21	3	64-QAM	2/3	144	324	160	360
22	3	64-QAM	3/4	162	364.5	180	405
23	3	64-QAM	7/8	189	425.25	210	472.5
24	4	BPSK	1/2	24	54	26.67	60
25	4	QPSK	1/2	48	108	53.33	120
26	4	QPSK	3/4	72	162	80	180
27	4	16-QAM	1/2	96	216	106.67	240
28	4	16-QAM	3/4	144	324	160	360
29	4	64-QAM	2/3	192	432	213.33	480
30	4	64-QAM	3/4	216	486	240	540
31	4	64-QAM	7/8	252	567	280	630
32	1	BPSK	1/2		6		6.67

Table 26: MCS field content for the basic mode

MCS indices from 33 to 63 (31 in number) are reserved for the Advanced MIMO Mode.

11.2.1.4.2.3 ADVCODING

ADVCODING is a 1-bit field that indicates if advanced coding has been applied to the PPDU.

11.2.1.4.2.4 SOUNDING PACKET

This bit indicates that MIMO training is per antenna, implying that all transmit antennas are excited. If this bit is set, the packet can be used for antenna-to-antenna channel estimation necessary for closed loop beam forming. If NOT set, the packet cannot be used as a sounding packet either because the MIMO training excites fewer spatial streams than antennas, or because beam forming has been applied.

11.2.1.4.2.5 NUMBER (of) HT-LTF

These bits indicate the number of HT-LTFs in the current PPDU. This enumeration relates to the sounding packet, which is used to train the full spatial channel between the transmit and the receive antenna array. Note that the number of HT-LTFs is less than or equal to the number of transmit antennas.

11.2.1.4.2.6 SHORT GI

SHORT GI is a 1-bit field. If SHORT GI is 1, the Short Guard Interval is used in the data portion of the frame.

11.2.1.4.2.7 AGGREGATE

AGGREGATE is a 1-bit field. If AGGREGATE is 1, aggregation in the PPDU is used.
This field robustly transports the AGGREGATE parameter of the PHY DATA service, which is present in both TXVECTOR and RXVECTOR. The MAC uses this transported parameter to interpret the PSDU as either unaggregated or aggregated.

11.2.1.4.2.8 SCRAMBLERINIT

SCRAMBLERINIT is a 2-bit field that contains the initialization value of the scrambler. This value is used to initialize the scrambler at the receiver. The mapping of the 2-bit field to initialization sequences is described in Table 18.

11.2.1.4.2.9 20/40 BW

This bit indicates whether the current PPDU is signaled in 20MHz or 40MHz bandwidth. A value of 0 indicates 20MHz, whereas a value of 1 indicates 40MHz.

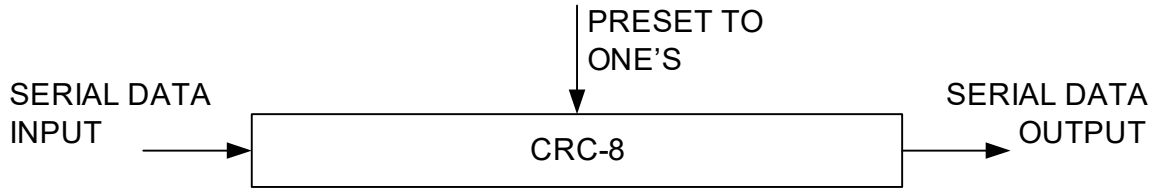
11.2.1.4.2.10 CRC

CRC is an 8-bit CRC value. The CRC is taken over LSIG, HTSIG1 and bit0-bit9 of HTSIG2. The CRC field shall be the one's complement of the remainder generated by the modulo 2 division of the protected bits by the polynomial

$$x^8 + x^2 + x^1 + x^0 \quad (35)$$

where the shift-register state shall be preset to all-ones. The protected bits shall be bit0-bit16 of LSIG, bit0-bit23 of HTSIG1, and bit0-bit9 of HTSIG2. The protected bits shall be processed by the CRC encoder in transmit order. A schematic of the processing is shown in Figure 63.
As an example, consider the following values for the LSIG, HTSIG1 and HTSIG2 fields:
1101 0001 1101 1101 0000 0000 [LSIG]
0111 1011 0100 1011 1010 0000 [HTSIG1]
0010 1101 10CC CCCC CC00 0000 [HTSIG2]

where b_0.....b_{23} [leftmost bit (b_0) is transmitted first in time]
Then, the protected bits that are input to the CRC encoder are
1101 0001 1101 1101 0, 0111 1011 0100 1011 1010 0000, 0010 1101 10
and the calculated CRC8 equals
00001101
so that HTSIG2 equals
0010 1101 1000 0011 0100 0000



1. Present shift register to all one's
2. Shift bit0-bit16 of LSIG, HTSIG1, and bit0-bit9 of HTSIG2 through the shift register
3. Take one's complement of the remainder
4. Transmit out serial x^7 first

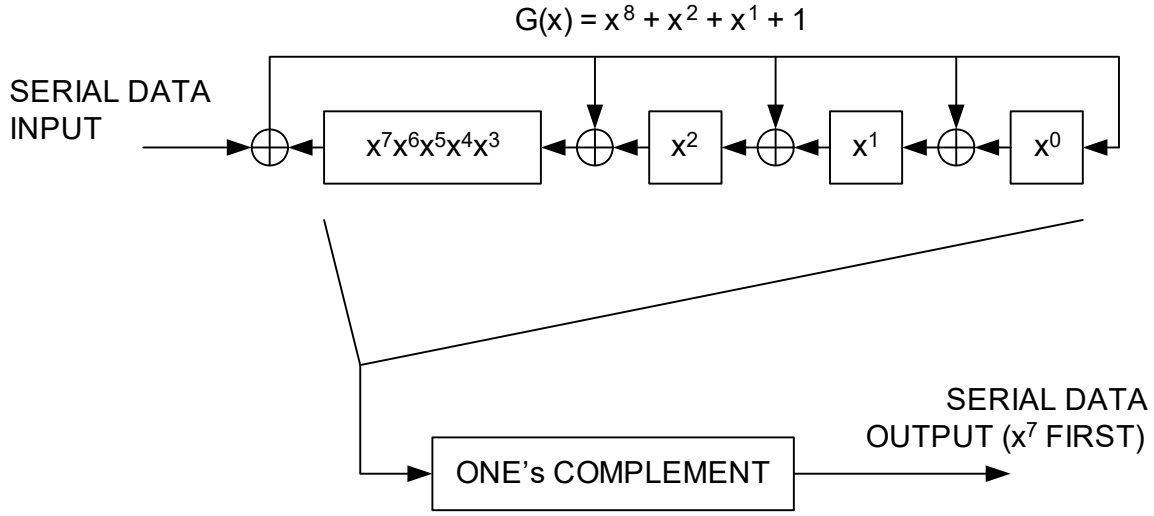


Figure 63: Generation of CRC bits

11.2.1.4.2.11 SIGNALTAIL

The bits 18-23 constitute the SIGNALTAIL field, and all 6 bits shall be set to zero.

11.2.1.5 Datafield

Contrary to 801.11a, the data field in our proposal does not contain a SERVICE field.

The data field contains the PSDU, the Tail bits and the PAD bits. All bits in the DATA field are scrambled, as described in 11.2.1.5.3. The PPDU is de-multiplexed across several spatial

streams, which are created after the parser. Figure 45. shows the de-multiplexing across 2 spatial streams. Our proposal supports up to 4 spatial streams.

11.2.1.5.1 PSDU tail bit field (TAIL)

The PSDU tail bit field shall be six bits of "0," which are required to return the convolutional encoder to the "zero state." This procedure improves the error probability of the convolutional decoder, which relies on future bits when decoding and which may be not be available past the end of the message. The PLCP tail bit field shall be produced by replacing six scrambled "zero" bits following the message end with six nonscrambled "zero" bits.

11.2.1.5.2 Pad bits (PAD)

The length of the message is extended so that it becomes a multiple of the number of data bits per OFDM symbol for each data (spatial) stream. The number of Pad bits is dependent how the bits are divided over the spatial streams.

11.2.1.5.3 PLCP Data Scrambler and Descrambler

The 11a Scrambler and Descrambler will be used as described in 17.3.5.4 in ref [1].
The DATA field, composed of PSDU, tail, and pad parts, shall be scrambled with the length-127 frame-synchronous scrambler.
The scrambler shall be initialized according to the SCRAMBLERINIT fields in HTSIG2.
As described in 11.2.1.5.1 the PLCP tail bit field shall be produced by replacing six scrambled "zero" bits following the message end with six nonscrambled "zero" bits.

11.2.1.5.4 Antenna streams / Transmitter architecture

The transmitter architecture with 2 spatial streams in the basic MIMO mode is shown in Figure 45. In the basic MIMO mode, the number of spatial streams equals the number of active transmit antennas. By active, we imply those antennas that transmit datafields.

The information bits first go through a convolutional encoder. The convolutional encoder is of rate ½ , from which other rates are derived by puncturing (see 11.2.1.5.6). The output of the puncturer feeds into the spatial parser, which creates several spatial streams, using round robin cycling. The frequency interleaver places the bits in the frequency domain and the constellation position according to the rules of the 802.11a interleaver. The interleaved bits are then mapped to constellation points. The resulting QAM symbols are fed as a block of data to the iFFT to create the time-domain signal. The pilot tones are also inserted in the frequency domain. The cyclical prefix is inserted in the time domain, and windowing of the OFDM symbols is also performed in the time domain. The resulting signal is then converted into analog format via a Digital-to-Analog (D/A) converter.

11.2.1.5.5 Encoder

The DATA field shall be coded with the convolutional encoder of coding rate R=1/2. The encoder is defined in clause 17.3.5.5 in [1] (i.e. the 11a encoder).
.

11.2.1.5.6 Puncturing of Coded Data

Coded bits at the output of the convolutional encoder are at rate 1/2. For lower rates, the puncturer can selectively puncture the coded bits to achieve the desired "effective" coding rate of 2/3, 3/4 or 7/8. For code rates of 2/3 and 3/4 the puncturing patterns are as defined in clause 17.3.5.5 in [1]. For a code rate of 7/8, the puncturing pattern is illustrated in Figure 64.

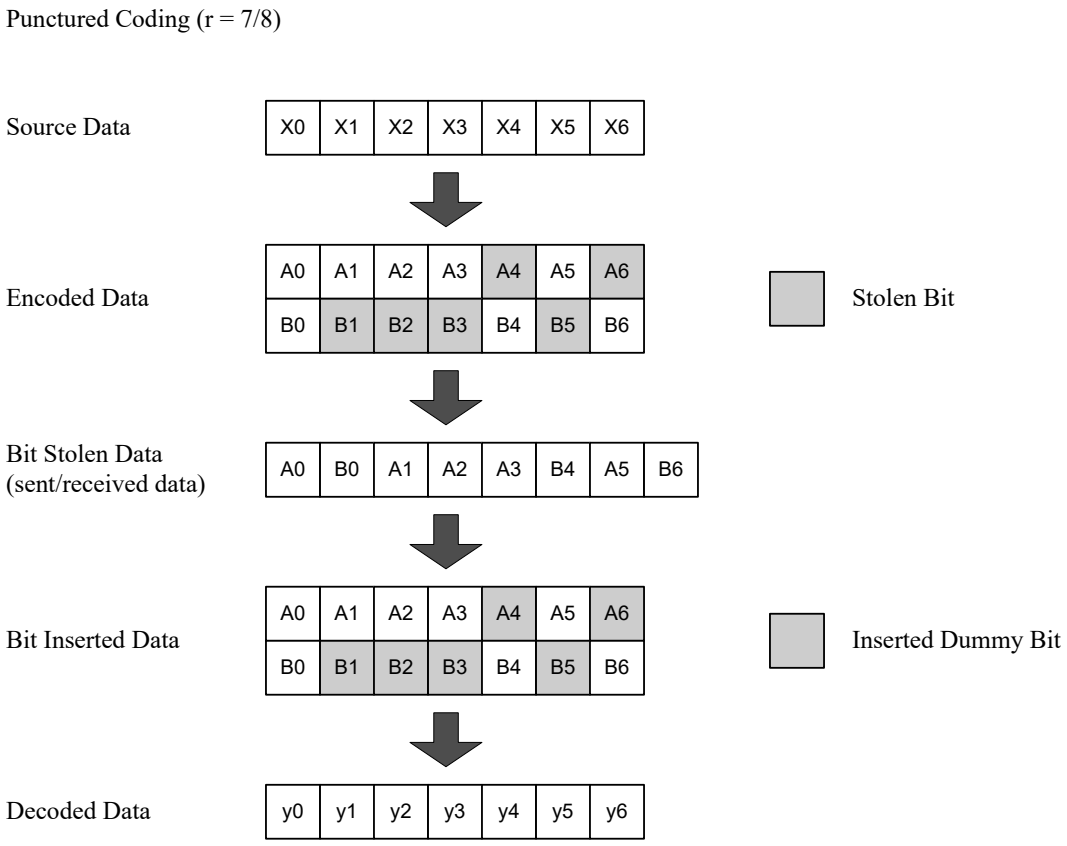


Figure 64: Puncturing pattern for code rate = 7/8

In the basic MIMO mode, all spatial streams are identically punctured.

11.2.1.5.7 Data interleaving

Coded and punctured bits are interleaved across spatial streams and frequency tones. There are two steps to the space-frequency interleaving: spatial stream parsing and frequency interleaving.

11.2.1.5.7.1 Spatial Parsing

Encoded and punctured bits are parsed to multiple spatial streams by a round-robin parser. Define:

$$s = \max\{N_{BPSC}/2, 1\} \quad (36)$$

which is the number of QAM bit order values, as in [1] equation (17). Then the parser sends consecutive blocks of s bits to different spatial streams in a round-robin fashion starting with the first spatial stream. The reason for parsing in blocks of s bits is to prevent clustering of QAM orders.

11.2.1.5.7.2 Frequency interleaver

All encoded bits shall be interleaved by a separate block interleaver for each spatial stream, with a block size corresponding to the number of coded bits in a single OFDM symbol, N_{CBPS}. The block interleavers are based on the 802.11a interleaver, with certain modifications to allow for multiple spatial streams and 40MHz transmissions as shown in Figure 65.

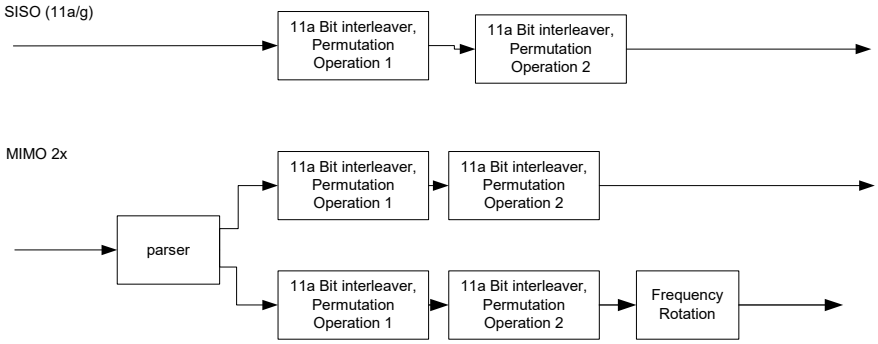


Figure 65: Structure of interleaver

The interleaver is defined by a three-step permutation. The first permutation ensures that adjacent coded bits are mapped onto nonadjacent subcarriers. The second permutation ensures that coded bits are mapped alternately onto less and more significant bits of the constellation and thereby long runs of low reliability (LSB) bits are avoided. The third permutation ensures that coded bits are mapped to achieve better frequency-spatial diversity among all spatial streams via frequency rotation. The operations of the first two permutations are the same as in 802.11a.

The basic interleaver array has N_{row} rows and N_{column} columns and the base of frequency rotation is defined by N_{rot}. These parameters are specified in

Table 27 where N_{BPSC} is the number of coded bits per subcarrier (e.g., N_{BPSC} = 1 for BPSK, 2 for QPSK, 4 for 16 QAM, etc).

Table 27: Interleaver Parameters

	N_{column}	N_{row}	N_{rot}
20 MHz channels	16	3 N_{BPSC}	11
40 MHz channels	18	6 N_{BPSC}	29

The exact amount of frequency rotation for different spatial streams is illustrated in Table 28.

Table 28: Amount of Frequency Rotation

Frequency Rotation	Total Number of Streams			
	1	2	3	4
1st stream	0	0	0	0
2nd stream		$2N_{rot}$	$2N_{rot}$	$2N_{rot}$
3rd stream			N_{rot}	N_{rot}
4th stream				$3N_{rot}$

The index of the coded bit before the first permutation shall be denoted by *k*; *i* shall be the index after the first and before the second permutation; *j* shall be the index after the second permutation and before the third permutation; *r* shall be the index after the third permutation and just prior to modulation mapping. Mathematically,

The first permutation is defined by the rule

$$i = N_{row} \times \left(k \bmod N_{column}\right) + \text{floor}\left(k / N_{column}\right) \quad k = 0,1,\ldots,N_{CBPS} - 1 \tag{37}$$

The second permutation is defined by the rule

$$j = s \times \text{floor}\left(i / s\right) + \left(i + N_{CBPS} - \text{floor}\left(N_{column} \times i / N_{CBPS}\right)\right) \bmod s \quad i = 0,1,\ldots,N_{CBPS} - 1 \tag{38}$$

where *s* is determined according to

$$s = \max\left(N_{BPSC} / 2,1\right) \tag{39}$$

The third permutation is defined by the rule

$$r = \left(j - \left(\left(2 \times i_{ss}\right) \bmod 3 + 3 \times \text{floor}(i_{ss} / 3)\right) \times N_{rot} \times N_{BPSC}\right) \bmod N_{CBPS} \quad j = 0,1,\ldots,N_{CBPS} - 1 \tag{40}$$

where $i_{SS} = 0,1,\ldots,N_{SS} - 1$ is the index of the spatial stream on which this interleaver is operating.

The deinterleaver, which performs the inverse relation, is also defined by three permutations. Here the index of the original received bit before the first permutation shall be denoted by *r*; *j* shall be the index after the first and before the second permutation; *i* shall be the index after the second and before the third permutation; and *k* shall be the index after the third permutation.

The first permutation is defined by the rule

$$j = \left(r + \left(\left(2 \times i_{ss}\right) \bmod 3 + 3 \times \text{floor}(i_{ss} / 3)\right) \times N_{rot} \times N_{BPSC}\right) \bmod N_{CBPS}, \quad r = 0,1,\ldots,N_{CBPS} - 1 \tag{41}$$

This permutation is the inverse of the permutation described in Equation (40).

The second permutation is defined by the rule

$$i = s \times \text{floor}\left(j / s \right) + \left(j + \text{floor}\left(N_{column} \times j / N_{CBPS} \right) \right) \bmod s, \quad j = 0,1,\ldots, N_{CBPS} - 1 \qquad (42)$$

where *s* is as defined in Equation (39). This permutation is the inverse of the permutation described in Equation (38).

The third permutation is defined by the rule

$$k = N_{column} \times i - \left(N_{CBPS} - 1 \right) \text{floor}\left(i / N_{row} \right), \quad i = 0,1,\ldots, N_{CBPS} - 1 \qquad (43)$$

This permutation is the inverse of the permutation described in Equation (37).

11.2.1.5.8 Subcarrier Modulation Mapping

The subcarrier modulation mapping will be identical to clause 17.3.5.7 in [1].

There is one exception as described in clause 11.2.1.4. The mapping of the HTSIG field shall be rotated by 90 degrees.

11.2.1.5.9 MIMO-OFDM symbol iFFT mapping

For the MIMO-OFDM operation in 20 MHz bandwidth, the modulation coding scheme (MCS) for each spatial stream will be identical to that in IEEE 802.11a (clause 17.3.5.9 in [1]).

Construction of the OFDM data symbol in 40MHz begins with bonding two adjacent 20MHz OFDM symbols, as shown in Figure 66. There are a total of 11 guard tones between the two 20MHz channels, ranging from -5 to +5. Note that the legacy channels occupy tones from -58 to -6 in the lower sub-channel, and from +6 to +58 in the upper sub-channel. Of the 11 guard tones, 3 are nulled out: -1, 0, +1. This provides 8 additional data tones. Two DC tones are also inherited from the legacy sub-channels. These occur at -32 and +32 in the lower and upper sub-channels respectively. There are 4 pilot tones in each legacy sub-channel, and when mapped to the 40MHz iFFT indices, they occur at ±53, ±39, ±25, ±11. Of these 8 tones, the two tones at -39 and +39 are swapped with data, leading to a total of 6 pilot tones in the 40MHz channel. Tone assigned during 40MHz for DATA symbols can be summarized as:

$$\begin{aligned} &Nulled\ Tones = \{-64\ldots-59, -1, 0, +1, +59\ldots+63\} \\ &Populated\ Tones = \{-58\ldots-2, +2\ldots+58\} \\ &Pilot\ Tones = \{-53, -25, -11, +11, +25, +53\} \\ &Data\ Tones = \{Populated\ Tones\} - \{Pilot\ Tones\} \end{aligned} \qquad (40)$$

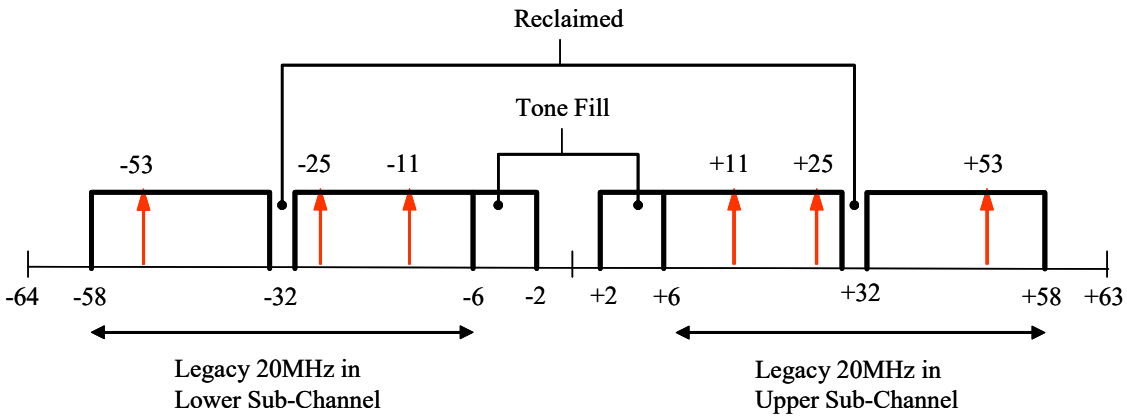


Figure 66: Tone Format for 40MHz channelization

The mapping of the input of the 40MHz channel to a 128-point IFFT is illustrated in Figure 67.

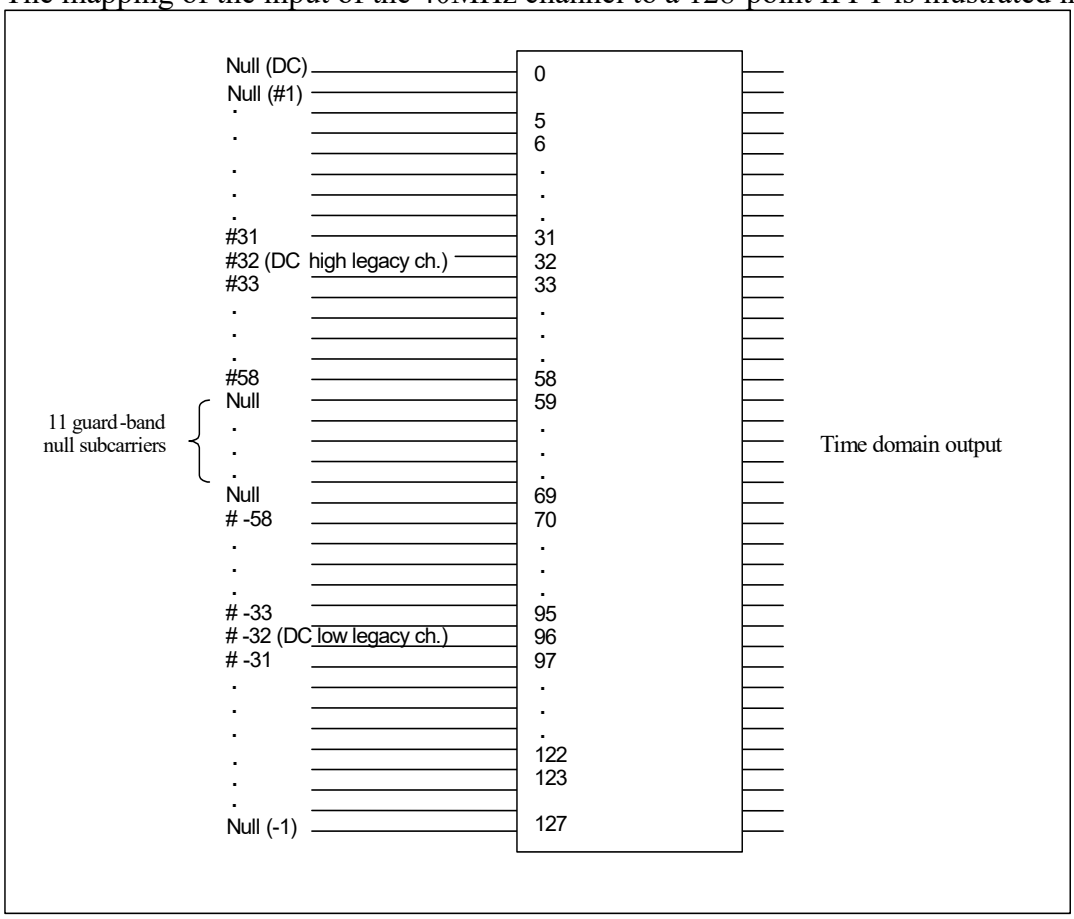


Figure 67 Inputs and Outputs of IFFT for a 40 MHz channel

The complex numbers after the mapping are divided into groups of N_{SD}, where $N_{SD} = 48$ in the 20MHz mode and $N_{SD} = 108$ in the 40MHz mode. Following the convention in clause 17.3.5.9 in [1], the complex number is denoted as $d_{k,n}^{(i_{Tx})}$, which corresponds to the k th subcarrier in the n th OFDM symbol that is transmitted from the i_{Tx} th antenna.

An OFDM data symbol, $r_{DATA,n}^{(i_{Tx})}(t)$, is defined as

$$r_{DATA,n}^{iTX}(t) = w_{SYM}(t)\left(\sum_{k=0}^{N_{SD}-1} d_{k,n}^{(i_{TX})} \exp\left(j2\pi M(k)\Delta_F\left(t-T_{GI}\right)\right) + p_{n+3}\sum_{k=-N_{SR}}^{N_{SR}} P_k^{(i_{TX})} \exp\left(j2\pi k\Delta_F\left(t-T_{GI}\right)\right)\right) \quad (41)$$

where function $M(k)$, defines the mapping from the logical subcarrier number k to a value between $-N_{SR}$ and $+N_{SR}$, while skipping the pilot subcarriers and DC and its adjacent subcarriers. $M(k)$ in the 20MHz mode is same as the one defined in clause 17.3.5.9 in [1]. $M(k)$ in the 40MHz mode is defined as

$$M(k) = \begin{cases} k-58 & 0 \le k \le 4 \\ k-57 & 5 \le k \le 31 \\ k-56 & 32 \le k \le 44 \\ k-55 & 45 \le k \le 53 \\ k-52 & 54 \le k \le 62 \\ k-51 & 63 \le k \le 75 \\ k-50 & 76 \le k \le 102 \\ k-49 & 103 \le k \le 107 \end{cases} \quad (42)$$

The contribution of pilot subcarriers in the 20MHz mode is given by (24) in clause 17.3.5.9 in ref [1]. The subcarrier frequency locations in the 20MHz mode is shown in Figure 117 in clause 17.3.5.9 in ref [1]. The subcarrier frequency locations in the 40MHz mode is shown in Figure 66.

The iFFT mapping for the 6Mbps duplicate mode in 40MHz differs slightly from the mapping shown in (41):

$$r_{DATA,n}^{(1)}(t) = w_{SYM}(t)\left(\sum_{k=0}^{N_{SD}-1} d_{k,n}\exp\left(j2\pi\left(M(k)-32\right)\Delta_F\left(t-T_{GI}\right)\right) + j\sum_{k=0}^{N_{SD}-1} d_{k,n}\exp\left(j2\pi\left(M(k)+32\right)\Delta_F\left(t-T_{GI}\right)\right) + p_{n+3}\sum_{k=-N_{SR}}^{N_{SR}} P_k\exp\left(j2\pi\left(k-32\right)\Delta_F\left(t-T_{GI}\right)\right) + jp_{n+3}\sum_{k=-N_{SR}}^{N_{SR}} P_k\exp\left(j2\pi\left(k+32\right)\Delta_F\left(t-T_{GI}\right)\right)\right) \quad (43)$$

where, $N_{SD}=48$ and $N_{SR}=26$ are values that correspond to a 20MHz transmission. $M(k)$, P_k and p_n follow the definitions in (23) to (25) in clause 17.3.5.9 in [ref 2]. P_k for the HT 40MHz mode is defined in 11.2.1.5.10.2.

The concatenation of N_{SYM} OFDM symbols can now be written as

$$r_{DATA}^{(iTx)}(t) = \sum_{n=0}^{N_{SYM}-1} r_{DATA,n}^{(iTx)}\left(t-nT_{SYM}\right) \quad (44)$$

11.2.1.5.10 Pilot subcarrier

In each OFDM symbol, four of the subcarriers in 20 MHz mode and six of the subcarriers in 40 MHz mode are dedicated to pilot signals in order to make the coherent detection robust against frequency offsets and phase noise. These pilot signals shall be put in subcarriers –21, -7, 7, and 21 in 20MHz mode, and in subcarriers -53, -25, -11, 11, 25 and 53 in 40MHz mode.

11.2.1.5.10.1 Legacy pilot subcarriers

Legacy pilots are defined as in 802.11a; refer to clause 17.3.5.8 in [1].

11.2.1.5.10.2 Pilot subcarriers in HT mode

The contribution of pilot subcarriers at the i_{Tx} th transmission stream in 20MHz mode is produced by inverse Fourier transform of sequence $P_{k,n}^{(i_{Tx})}$ given by

$$\begin{aligned} P_{-26..26,n}^{(i_{Tx})} = \{&0,0,0,0,0,P_{-21,n}^{(i_{Tx})},0,0,0,0,0,0,0,0,0,0,0,0,0,P_{-7,n}^{(i_{Tx})},0,0,0,0,0,0,0, \\ &0,0,0,0,0,0,P_{7,n}^{(i_{Tx})},0,0,0,0,0,0,0,0,0,0,0,0,0,-P_{21,n}^{(i_{Tx})},0,0,0,0,0\} \end{aligned} \quad (45)$$

where non-zero element of $P_{k,n}^{(i_{Tx})}$ is generated by the following rule

$$P_{k,n}^{(i_{Tx})} = \exp\left\{\frac{j\pi}{2}(i_{Tx}-1)m\right\}, \quad m = \left(\frac{1}{2}\left(\frac{k}{7}+3\right)+n\right) \bmod 4 \quad (46)$$

In the 40MHz mode, $P_{k,n}^{(i_{Tx})}$ is given by

$$\begin{aligned} P_{-58..58,n}^{(i_{Tx})} = \{&0,0,0,0,0,P_{-53,n}^{(i_{Tx})},0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,P_{-25,n}^{(i_{Tx})},0,0, \\ &0,0,0,0,0,0,0,0,0,0,0,-P_{-11,n}^{(i_{Tx})},0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,P_{11,n}^{(i_{Tx})},0,0,0,0, \\ &0,0,0,0,0,0,0,0,0,0,-P_{25,n}^{(i_{Tx})},0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0,0, \\ &0,0,P_{53,n}^{(i_{Tx})},0,0,0,0,0\} \end{aligned} \quad (47)$$

where non-zero element of $P_{k,n}^{(i_{Tx})}$ is generated by the following rule

$$P_{k,n}^{(i_{Tx})} = \exp\left\{\frac{j\pi}{2}(i_{Tx}-1)m\right\}, \quad m = \begin{cases} \left(\frac{(k+11)}{14}+3+n\right) \bmod 4, & k = -53,-25,-11 \\ \left(\frac{(k+3)}{14}+3+n\right) \bmod 4, & k = 11,\ 25,\ 53 \end{cases} \quad (48)$$

The polarity of the pilot subcarriers is controlled by sequence p_n, which is a cyclic extension of the 127 elements sequence as defined in clause 17.3.5.9 in ref[2].

11.2.1.5.11 GI / windowing

The windowing function in both 20MHz and 40MHz modes is the same as the one described in clause 17.3.2.4 and 17.3.2.5 in [1]. The example in clause 17.3.2.5 in [1] still holds for the 20MHz modes. In the 40MHz modes, the changes in this example is the sampling rate becomes 40MHz and $T_{TR} = 50ns$, then the discrete implementation of the windowing function is

$$w_T[n] = w_T(nT_s) = \begin{cases} 1 & 1 \le n \le 159 \\ 0.5 & 0,160 \\ 0 & otherwise \end{cases} \quad (49)$$

11.2.1.6 Clear channel assessment

When operating in the 5 GHz band the HT PHY shall follow the CCA requirements of 11a, clause 17.3.10.5 in [1]. When operating in the 2.4 GHz band the HT PHY shall follow the 11g spec according to 19.3.5 and 19.4.6 in ref [3].

11.2.1.7 PMD operating spec

In the 20 MHz bandwidth mode the HT PHY shall follow the 11a spec (clause 17.3.8 in ref [1]) and subsequent specifications when operating in the 5 GHz band.
In the 20 MHz bandwidth mode the HT PHY shall follow the 11g spec (clause 19.4) when operating in the 2.4 MHz band.

11.2.1.7.1 Transmit spectral mask

11.2.1.7.1.1 20MHz channel spacing

In the 20 MHz bandwidth mode the HT PHY shall follow 17.3.9.2 in ref [1]. The transmitted spectrum shall have a 0dBr (dB relative to the maximum spectral density of the signal) bandwidth not exceeding 18MHz, -20dBr at 11 MHz frequency offset, -28dBr at 20MHz frequency offset and -45dBr at 30MHz frequency offset and above. The transmitted spectral density of the transmitted signal shall fall within the spectral mask, as shown in Figure 1. The measurements shall be made using a 100kHz resolution bandwidth and a 30kHz video bandwidth.

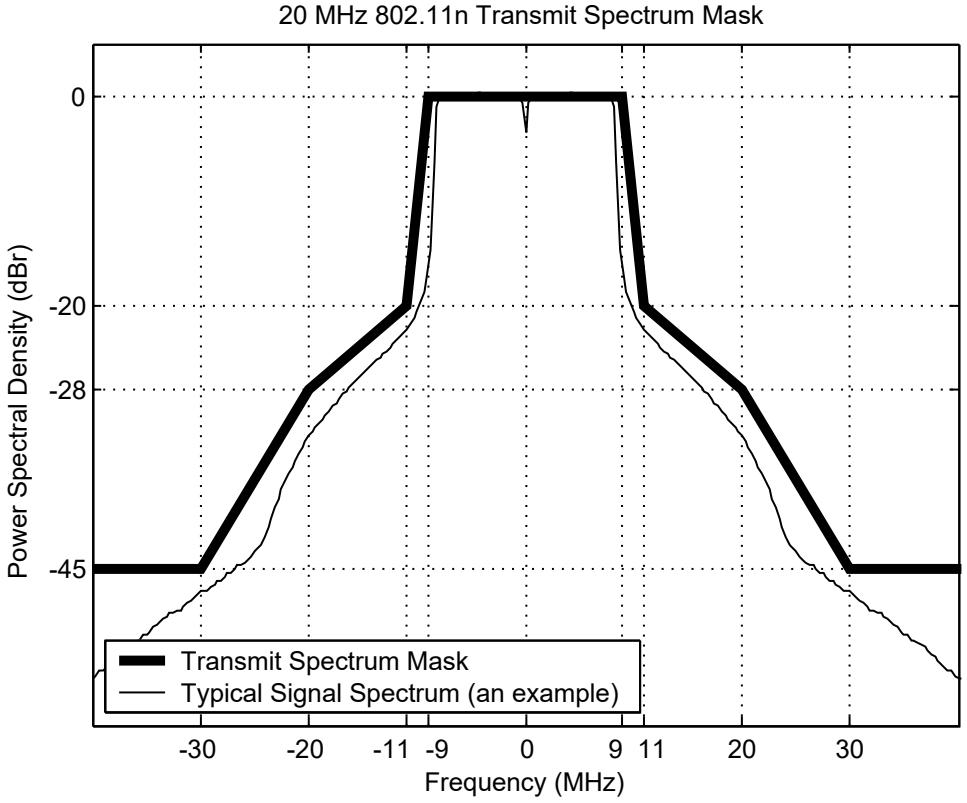


Figure 68: Transmit Spectrum Mask for 20MHz channelization

11.2.1.7.1.2 40MHz channel spacing

The transmitted spectrum shall have a 0dBr (dB relative to the maximum spectral density of the signal) bandwidth not exceeding 38MHz, -20dBr at 21 MHz frequency offset, and -45dBr at 60MHz frequency offset and above. The transmitted spectral density of the transmitted signal shall fall within the spectral mask, as shown in Figure 2. The measurements shall be made using a 100kHz resolution bandwidth and a 30kHz video bandwidth.

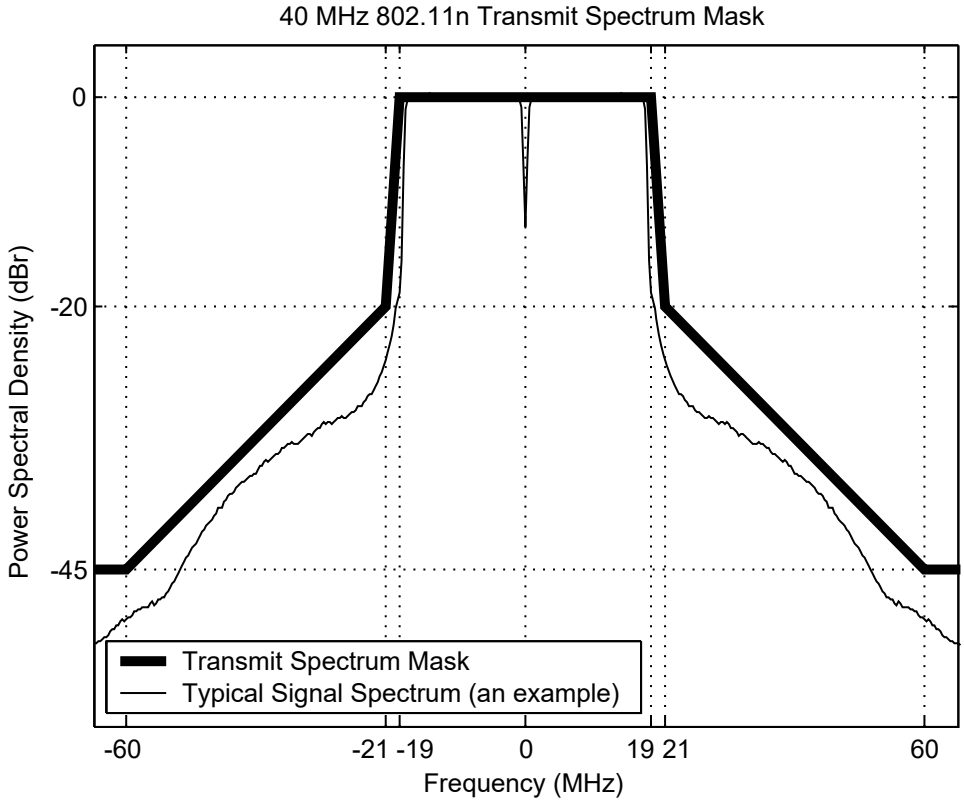


Figure 69: Transmit spectrum mask for 40MHz channelization

A HT device transmitting a 20MHz signal in a 40MHz BSS shall conform to the spectral mask shown in Figure 69.

11.2.2 Basic Transmit Beamforming (Optional Mode)

11.2.2.1 Introduction

One method of improving the SNR at the receiver, for a given distance between stations, is to apply transmit beamforming. A basic requirement for transmit beamforming is the availability of channel state information (CSI). In our proposal, CSI is obtained by assuming channel reciprocity. The initiator sends a request for a sounding packet to the intended recipient, which in turn sends a sounding packet back to the initiator. The initiator extracts the CSI from this packet, and assumes that the forward channel is identical to the reverse channel. Assumption of channel reciprocity also requires radio calibration, since the TX-RX radio pair on the reverse link is different from the TX-RX pair on the forward link. Our proposal also includes a packet exchange sequence to help perform radio calibration.

11.2.2.2 Motivation of MIMO-Beamforming (Informative)

Beamforming-MIMO can achieve higher throughput than basic MIMO at same SNR, or conversely, higher range at the same data rate. The benefit of TX beamforming is most pronounced when the number of TX antennas is larger than the RX antennas, when multiple spatial streams are being transmitted. For e.g., a typical configuration involving TX beamforming could be either 3 or 4 antennas at the AP, and 2 antennas at the client. Hence, the client is kept low-cost, while the AP shoulders a higher cost burden.

11.2.2.3 Description of BF-MIMO

11.2.2.3.1 Beamformed PPDU

The transmitter may choose to use antenna beamforming in order to increase the antenna gain in the direction of an intended receiver. The beamforming process is represented mathematically as a linear transformation, mapping the streams of data to the transmit antennas.

Let *x* be the N_{SS} dimensional vector of data, where each element of *x* is from one of the spatial streams of data. The beamforming transformation is represented by a N_{TX} x N_{SS} matrix, V. The beamforming operation, on a per subcarrier basis, is represented by,

$$y_{N_{TX}\times 1} = V_{N_{TX}\times N_{SS}} \cdot x_{N_{ss}\times 1} \quad (50)$$

The *y* vector is an N_{TX} dimensional vector whose elements get mapped to the N_{TX} antennas. The beamforming matrix V is applied to the HT-STF and HT-LTF training fields as well as the data. There are several methods for calculating V. One of them involves singular value decomposition (SVD) of the channel transfer matrix H.

It is optional for a device to be able to perform TX beamforming. However, it is mandatory that a device to receive a beamformed packet, which implies that the device should be able to send a sounding packet to the initiator.

11.2.2.3.2 Channel sounding

11.2.2.3.2.1 Channel sounding principle

In order for a STA to send a beamformed PPDU (BF-PPDU) to another STA it is necessary for the STA to have channel state information from the transmitter to the receiver. To simplify the explanation we will introduce several new terms. The STA that is to send a BF-PPDU is referred to as the beamforming STA and the STA that is to receive the BF-PPDU is referred to as the non-beamforming STA.

Using implicit feedback, the non-beamforming STA sends a sounding PPDU to the beamforming STA, which enables the beamforming STA to estimate the channel from the non-beamforming STA to the beamforming STA. This channel estimate is used to calculate the required beamforming matrix for the beamforming-STA to non-beamforming STA transmission.

11.2.2.3.2.2 Sounding PPDU format

In a non-sounding PPDU the number of HT-LTFs is equal to the number of spatial streams, which may be less than the number of transmit antennas. However, in order to calculate the beamforming matrix from one STA to another STA it is necessary to estimate the full channel matrix, which requires that the PPDU include a full set of N_{TX} LTFs. In order to facilitate estimation of the complete channel matrix, a sounding PPDU is introduced. A sounding PPDU fully excites the channel, enabling estimation of the full channel matrix. No beamforming shall be applied to either the LTFs or the data in a sounding-PPDU. In other words the beamforming matrix, V, is the identity matrix. Thus techniques such as Walsh + CDD should not be applied to sounding PPDU.

In the sounding PPDU, the "sounding bit" in the HT-SIG is set to "one". The number of HT-LTFs is equal to the number of transmit antennas, and not necessarily the number of spatial streams. Thus in a sounding PPDU the number of HT-LTFs is set by the NLTF portion of the HT-SIGNAL field, which overrides the setting of the number of HT-LTFs specified by the MCS portion of the HT-SIGNAL field. The NUMBER HT-LTF is identical to the number of TX antennas.

It is mandatory that a STA support the transmission of a sounding PPDU in response to a TRQ message in the IAC MPDU.

11.2.2.3.3 Calculation of beamforming matrix using SVD (Informative)

Once a STA has a proper estimate of the channel, then it can use that channel estimate to select a recommended beamforming matrix. There are a variety of techniques that can be used to select the beamforming matrix. One method of selecting the beamforming matrix is by calculating the singular value decomposition of the channel from this STA to the intended receiving STA. Let H be the channel matrix from the transmitter to the receiver. The SVD of the channel matrix is given by,

$$H = UDV^H \quad (51)$$

Both U and V are unitary matrices and D is a diagonal matrix of singular values. The transmitter selects the dominant N_{SS} singular values and then sends only N_{SS} spatial streams. The first N_{SS} rows of the V matrix are used as the beamforming matrix, mapping the spatial streams of data to the TX antennas. This same beamforming matrix is applied to the HT-LTFs, except for one difference. The number of HT-LTFs is equal to Nss, and each transmission of HT-LTF is shaped by one singular vector.

11.2.2.3.3.1 Channel Reciprocity

To determine the beamforming matrix, the transmitter needs to know the channel from the transmitter to the intended receiver. That channel is estimated using implicit feedback and channel reciprocity. The transmitter uses the previous sounding PPDU sent from the intended receiver to estimate the channel from the intended receiver to the transmitter. In order to estimate the channel from the receiver to the transmitter, it uses the principle of reciprocity. The principle of reciprocity is that the downlink channel (the channel from the AP to the STA) is the transpose of the uplink channel (the channel from the STA to the AP). Letting H_D be the downlink channel and H_U be the uplink channel, reciprocity is,

$$H_U = H_D^H \quad (52)$$

11.2.2.3.4 Calibration Methodology (Informative)

11.2.2.3.4.1 Calibration Overview

In order for a transmitting 802.11n STA to transmit to another 802.11n STA using beamforming, it must have a good characterization of the channel that it is transmitting over so that it can choose the correct transmit steering vectors. The STA obtains this characterization by observing the HT-LTF transmitted by the STA to which it is transmitting. This approach is based on the reciprocal property of the over-the-air (antenna-to-antenna) TDD channel, and reciprocity must be in effect for the approach to provide the best rate and range performance available.

Although the Basic MIMO and Walsh-CDD modes do not require it, calibration neither improves nor degrades performance with those modes. Thus, STAs may calibrate upon entering a BSS in which the AP is transmit beamforming capable, or at any other time, and if they are capable of transmitting in beamforming mode, the resulting calibration correction vectors may be employed for all transmission modes and on all frames. STAs that are not capable of transmitting in beamforming mode may also perform the calibration procedure with any transmit beamforming capable AP with which they associate, in order to enable that AP to transmit to them in beamforming mode.

Throughout the discussion of beamforming transmission in this document, it is assumed that communication in beamforming mode takes place between devices that have been calibrated.

The following section describes the calibration procedure for a communication link, where the two communicating STAs are referred to as STA A and STA B. Note: calibration as described here is only applicable to STAs with two or more antennas, where both transmission and reception are possible on each antenna.

11.2.2.3.4.2 Calibration procedure for Tx Corrections

In order to consider the operation of a TDD MIMO OFDM system in conjunction with beamforming techniques through transmit and receive steering, it is necessary to first consider the need to ensure that the over-the-air reciprocity of the TDD channel is preserved in the observed, baseband-to-baseband TDD channel. Reciprocity is critical in the functioning of these systems because it provides efficient means for a STA to learn the transfer function of the channel over which it will transmit.

In order to ensure that the baseband-to-baseband channel seen by the spatial processing engines is reciprocal, it is necessary to calibrate the STA radio to correct the differences in phase and amplitude response in the transmit and receive chains.

The reciprocity of the over-the-air TDD MIMO channel can be expressed as a transpose relationship between the channel response matrices associated with forward and reverse channels between two STAs. Consider, for example, STA A and STA B, where we will refer to the channel over which STA A transmits as the forward link, and the channel over which STA B transmits as the reverse link. Thus we will define the

following two gain matrices for the over-the-air channels (the over-the-air channel is the complex baseband channel that would be realized if the transmitter and receiver responses have ideal constant gain and zero phase across the channel bandwidth):

- $\mathbf{H}_{A\to B}(\ell)$: forward link gain matrix for subcarrier ℓ
- $\mathbf{H}_{B\to A}(\ell)$: reverse link gain matrix for subcarrier ℓ,

The reciprocal relationship between these two channels is expressed as:

$$\mathbf{H}_{B\to A}(\ell) = \mathbf{H}^{t}_{A\to B}(\ell).$$

The observed channels, however, are effected by the amplitude and phase responses of the transmit and receive chains. Assuming good isolation, these can be expressed as diagonal matrices with complex valued diagonal entries, of the form $\mathbf{C}_{A,Tx}(\ell)$. The relationship between the observed channels, $\tilde{\mathbf{H}}_{A\to B}(\ell)$ and $\tilde{\mathbf{H}}_{B\to A}(\ell)$, and the over-the-air channel is

$$\tilde{\mathbf{H}}_{A\to B}(\ell) = \mathbf{C}_{B,Rx}(\ell)\mathbf{H}_{A\to B}(\ell)\mathbf{C}_{A,Tx}(\ell), \text{ and}$$

$$\tilde{\mathbf{H}}_{B\to A}(\ell) = \mathbf{C}_{A,Rx}(\ell)\mathbf{H}_{B\to A}(\ell)\mathbf{C}_{B,Tx}(\ell).$$

The resulting observed channels are not, in general, reciprocal. However, through a simple channel sounding procedure that involves an exchange of training sequences, correction vectors can be calculated that are applied on transmit to restore reciprocity. The resulting calibration is quite stable and needs to be updated infrequently, typically on association.

The object of the procedure is to compute correction vectors $\mathbf{K}_A(\ell)$ and $\mathbf{K}_B(\ell)$ that restore reciprocity such that

$$\tilde{\mathbf{H}}_{A\to B}(\ell)\mathbf{K}_A(\ell) = \left[\tilde{\mathbf{H}}_{B\to A}(\ell)\mathbf{K}_B(\ell)\right]^{t}.$$

The correction vectors are applied on transmit in such a way that each element in the correction vector scales the transmit vector element associated with a specific antenna. Given a modulation symbol vector, $\mathbf{s}_A(\ell)$, in subcarrier ℓ at STA A, and beamforming matrix $\mathbf{V}(\ell)$, then the transmitted vector in subcarrier ℓ is given by $\mathbf{x}(\ell) = \mathbf{K}_A(\ell)\mathbf{V}_A(\ell)\mathbf{s}_A(\ell)$.

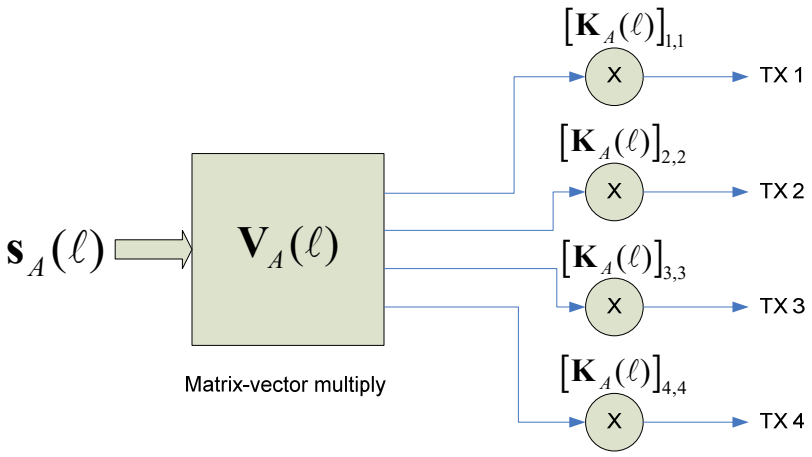


Figure 70: Application of calibration correction vector on Tx

Correction vectors that will satisfy this relationship are given by:

$$\mathbf{K}_A(\ell) = \left(\mathbf{C}_{A,Tx}(\ell)\right)^{-1} \mathbf{C}^t_{A,Rx}(\ell), \text{ and}$$

$$\mathbf{K}_B(\ell) = \left(\mathbf{C}_{B,Tx}(\ell)\right)^{-1} \mathbf{C}^t_{B,Rx}(\ell).$$

Calibration can also be applied on receive in a similar manner. Tx calibration only is addressed here.

For STAs that support only 20 MHz operation, calibration must be performed for subcarriers with indices between -26 and 26, except for the subcarrier with index 0. For STAs supporting both 20 MHz and 40 MHz operation, calibration must be performed for subcarriers with indices between -58 and 58, except for the subcarriers with indices -1, 0, and 1.

The calibration procedure is as follows:

1. STA A transmits a Calibration Training message, initiating the calibration procedure.
2. STA B transmits a Calibration Training message, which includes the HT-LTF that allows a receiving STA A to estimate the channel matrices $\mathbf{H}_{B\to A}(\ell)$.
3. STA A transmits a Calibration Training message, which includes the HT-LTF that allows STA B to estimate the channel matrices $\mathbf{H}_{A\to B}(\ell)$.
    a. The previous two steps must occur over a sufficiently short time interval that the channel does not change over the interval.
4. STA A computes an estimate of $\mathbf{H}_{B\to A}(\ell)$ based on the received HT-LTF.
5. STA B computes an estimate of $\mathbf{H}_{A\to B}(\ell)$ based on the received HT-LTF.

6. STA B sends the quantized estimates of $\mathbf{H}_{A\to B}(\ell)$ to STA A using the Calibration Measurement Message that contains the Channel Measurement Report.

7. STA A uses the estimates of $\mathbf{H}_{B\to A}(\ell)$ and $\mathbf{H}_{A\to B}(\ell)$ to compute the correction vectors $\mathbf{K}_A(\ell)$ and $\mathbf{K}_B(\ell)$.

8. STA A sends the quantized correction vectors $\mathbf{K}_B(\ell)$ to STA B using the Calibration Correction Vector Message that contains the Correction Vector element.

9. STA A and STA B now have correction vectors and are considered to be calibrated.

Once a STA has completed the calibration procedure, the correction vectors, $\mathbf{K}(\ell)$, are assumed to be available at the STA's transmitter, and are incorporated into the transmit spatial processing procedure for all subsequent communications.

This procedure needs to be performed only once for any STA that will participate in a beamforming exchange with another STA. For example, once an AP performs this procedure with one STA, it is calibrated and may participate in a beamforming exchange with any other calibrated STA. However, in many networks, the AP will be the primary agent for calibrating client STAs. Thus a typical AP will participate in many of these calibration procedures.

11.2.2.3.4.3 Calibration procedure for Beamforming Rx-only STAs

STAs that have no transmit beamforming capabilities, but that still are capable of receiving beamforming transmissions, are referred to as Rx-only STAs. Such STAs must also participate in a calibration procedure similar to that described above. However, it is not necessary for such a STA to apply the calibration correction matrix on transmit. In this case, a STA A that wishes to transmit to a STA B using beamforming, must initiate the calibration sequence described above, with the exception that it does not execute step 8. Instead, STA A makes use of $\mathbf{K}_B(\ell)$ to precondition its channel estimate for the purposes of computing transmit steering vectors. This may be achieved by right-multiplying STA A's estimate of $\mathbf{H}_{B\to A}(\ell)$ by $\mathbf{K}_B(\ell)$.

11.2.2.3.4.4 Calibration Processing

The MIMO calibration training sequence consists of $N_{\mathrm{CAL}} = N_{Tx}$ HT-LTFs. Each HT-LTF consists of an 800 ns guard interval followed by two periods of a 3.2 μs training symbol sequence, for a total duration of 7.2 μs. For details on the format of the HT-LTF, see section 11.2.1.3.4.

11.2.2.3.5 HT-Short Training Field in BF-MIMO mode

HT-STF is identical to basic MIMO case (refer to section 11.2.1.3.3). The HT-STF undergoes the same spatial steering as the DATA.

11.2.2.3.6 HT-Long Training Field in BF-MIMO mode

HT-LTFs are identical to the basic MIMO case (refer to section 11.2.1.3.4). The HT-LTF undergoes the same spatial steering as the DATA.

11.2.2.3.6.1 MCS field in BF-MIMO mode

MCS set for BF-MIMO is identical to that of basic MIMO, as described in section 11.2.1.4.2.2.

11.2.2.3.7 Antenna streams / Transmitter architecture in BF-MIMO mode

The only difference in the transmitter architecture with respect to the basic MIMO mode is the introduction of the spatial steering matrix, as shown in Figure 71:

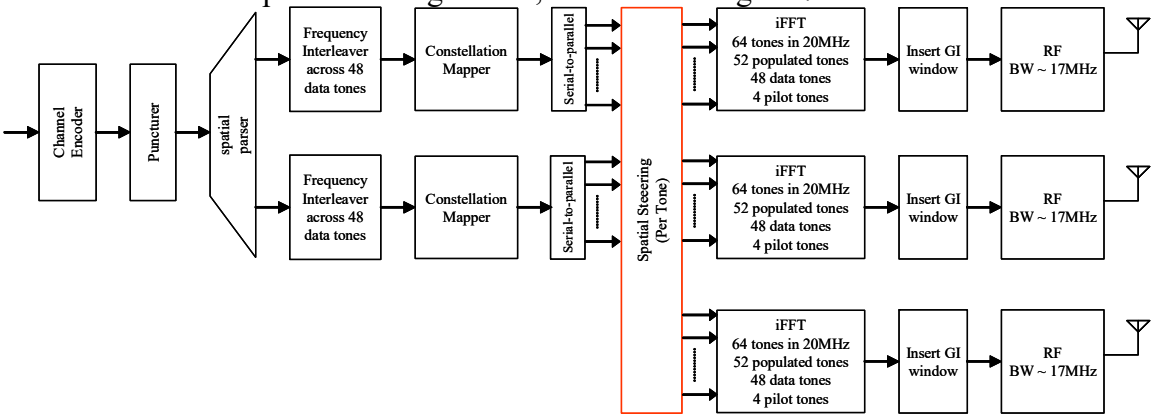


Figure 71: Transmitter architecture for TX Beamforming for N_{ss}=2,and N_{TX} = 3

11.2.2.3.8 Puncturing of Coded Data in BF-MIMO mode

TX Beamforming uses the same puncturing pattern as the basic MIMO mode, described in 11.2.1.5.6.

11.2.2.3.9 MCS set in BF-MIMO mode

The MCS is identical for all spatials streams in basic TX beamforming, as is the case in the basic MIMO mode. The MCS set is described in Table 26.

11.2.3 Advanced Transmit Beamforming (Optional Mode)

In our proposal, there are two class of BF-MIMO. The first one is basic BF-MIMO (as described in section 11.2.2). When AP uses this mode, STA would receive a beamformed frame as a basic MIMO frame. No additional functionality is required at the STA, except for the ability of the STA to send a sounding-PPDU and support feedback calibration. An extension to the basic BF-MIMO mode is the advanced BF-MIMO mode, which includes the following highlights:

- Unequal power among the spatial streams
- Independent MCS for each spatial stream

- Support for bi-directional beamforming

11.2.3.1 Required Capability for supporting of ABF-MIMO mode

In this mode, STA should support additional functionality to achieve better performance. These include:

- The receiving STA has to recognize a TRQ, and should be able to transmit a sounding PPDU as a response to it. (Common to basic BF-MIMO)
- Calibration support using MAC frame. (Common to basic BF-MIMO)
- Power level detection at each spatial stream to support unequal power loading
- Supporting unequal MCS at different spatial stream
- UD decomposition (or an equivalent method) to extract eigenvalues and eigenvectors of the channel matrix. Eigenvalues are used to determine the MCS at the receiver to support pairwise spoofing, and/or link adaptation. Eigenvectors in the U matrix allow bi-directional beamforming to be supported

11.2.3.1.1 Power Loading of spatial streams in ABF-MIMO

In ABF-MIMO, additional performance gain is possible if different power levels are applied to different spatial streams. However, to decode the PPDU with different power levels on different spatial streams, the receiver requires knowledge of these power ratios.

In the case of an ABF-PPDU, the values of the power level of each spatial stream can be varied as long as the total output power is maintained constant. Mathematically this can be viewed as multiplying the data vector by an N_{SS} x N_{SS} diagonal matrix, P, where the entries on the diagonal matrix are positive and the sum of the squares of the diagonal entries is N_{SS}. In other words,

$$p_1^2 + p_2^2 + \ldots + p_{N_{SS}}^2 = N_{SS} \quad (53)$$

After the scaling of power ratios, the beamforming matrix, V, is applied. The operation of power scaling and beamforming is represented as,

$$\vec{y} = V \cdot P \cdot \vec{x} \quad (54)$$

Transmitter should determine the power levels for each spatial stream based on the channel estimates obtained from the sounding packet.

11.2.3.1.1.1 Unequal power loading (Informative)

In order to decode the received ABF-PPDU, the receiver must estimate the power ratios between the spatial streams, which may not be equal to each other.

Once the streams of data have been separated at the receiver, the power ratios of the separated spatial streams can be estimated by estimating the power ratios of the pilot tones. The spatial streams of data can then be scaled using the estimate of the power ratios, derived from the power of the pilot tones.

11.2.3.1.1.2 Extended MCS set for ABF-MIMO mode

ABF-MIMO would use different modulation/coding scheme at each spatial stream. As a result, MCS fields in ABF-MIMO mode have been extended, as shown in Table 29.

Table 29: Extended MCS set for ABF-MIMO

ID#	Stream Count	Stream ID vs Modulation&Coding				Full GI		Half GI	
		stream 1	stream 2	Stream 3	stream 4	Rate in 20MHz	Rate in 40MHz	Rate in 20MHz	Rate in 40MHz
33	1	256QAM, 3/4	N/A	N/A	N/A	72	162	80	180
34	2	QPSK, 3/4	BPSK, 1/2	N/A	N/A	24	54	26.667	60
35	2	QPSK, 3/4	BPSK, 3/4	N/A	N/A	27	60.75	30	67.5
36	2	16QAM, 1/2	QPSK, 1/2	N/A	N/A	36	81	40	90
37	2	16QAM, 3/4	QPSK, 1/2	N/A	N/A	48	108	53.333	120
38	2	16QAM, 3/4	QPSK, 3/4	N/A	N/A	54	121.5	60	135
39	2	64QAM, 2/3	BPSK, 1/2	N/A	N/A	54	121.5	60	135
40	2	64QAM, 3/4	QPSK, 3/4	N/A	N/A	72	162	80	180
41	2	64QAM, 2/3	16QAM, 1/2	N/A	N/A	72	162	80	180
42	2	256QAM, 3/4	16QAM, 1/2	N/A	N/A	96	216	106.67	240
43	2	256QAM, 3/4	16QAM, 3/4	N/A	N/A	108	243	120	270
44	2	256QAM, 3/4	256QAM, 3/4	N/A	N/A	144	324	160	360
45	3	16QAM, 3/4	16QAM, 1/2	QPSK, 1/2	N/A	72	162	80	180
46	3	64QAM, 2/3	QPSK, 3/4	BPSK, 1/2	N/A	72	162	80	180
47	3	64QAM, 3/4	16QAM, 1/2	QPSK 3/4	N/A	96	216	106.67	240
48	3	64QAM, 3/4	16QAM, 3/4	QPSK, 3/4	N/A	108	243	120	270
49	3	64QAM, 3/4	64QAM, 2/3	BPSK, 1/2	N/A	108	243	120	270
50	3	64QAM, 3/4	64QAM, 3/4	16QAM, 3/4	N/A	144	324	160	360
51	3	256QAM, 3/4	64QAM, 2/3	16QAM, 1/2	N/A	144	324	160	360
52	3	256QAM, 3/4	64QAM, 3/4	QPSK 3/4	N/A	144	324	160	360
53	3	256QAM, 3/4	64QAM, 3/4	16QAM, 3/4	N/A	162	364.5	180	405
54	3	256QAM, 3/4	256QAM, 3/4	QPSK 3/4	N/A	162	364.5	180	405
55	3	256QAM, 3/4	256QAM, 3/4	64QAM, 2/3	N/A	192	432	213.33	480
56	4	64QAM, 2/3	16QAM, 3/4	BPSK, 1/2	BPSK, 1/2	96	216	106.67	240
57	4	64QAM, 3/4	64QAM, 3/4	16QAM, 1/2	QPSK, 1/2	144	324	160	360
58	4	256QAM, 3/4	64QAM, 3/4	64QAM, 2/3	QPSK, 3/4	192	432	213.33	480
59	4	256QAM, 3/4	64QAM, 3/4	64QAM, 3/4	QPSK, 1/2	192	432	213.33	480
60	4	256QAM, 3/4	256QAM, 3/4	16QAM, 3/4	QPSK, 1/2	192	432	213.33	480
61	4	256QAM, 3/4	256QAM, 3/4	64QAM, 2/3	16QAM, 1/2	216	486	240	540
62	4	256QAM, 3/4	256QAM, 3/4	64QAM, 3/4	QPSK, 3/4	216	486	240	540
63	reserved	reserved	reserved	reserved	reserved				

11.2.3.2 Transmitter Architecture in ABF-MIMO mode

The transmitter architecture for the ABF-MIMO mode includes the provision for independent MCS for each spatial stream. The Transmit datapath architecture is shown in Figure 72:

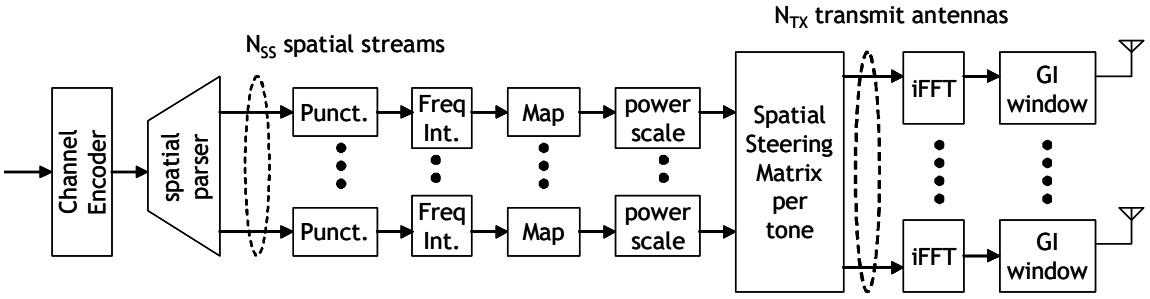


Figure 72: Transmit Datapath Architecture for ABF-MIMO mode

11.2.3.2.1 Subcarrier Modulation Mapping in ABF-MIMO mode

256QAM is introduced in the ABF-MIMO mode. The normalization factor K_{MOD} for 256QAM is $1/\sqrt{170}$. For 256-QAM, $b_0b_1b_2b_3$ determines the I value and $b_4b_5b_6b_7$ determines the Q value, as illustrated in Table 30. The constellation mapping for 256QAM is as follows.

Table 30: Constellation Mapping for 256-QAM in ABF-MIMO Mode

Input bits $b_0b_1b_2b_3$	I-out
0000	-15
0001	-13
0011	-11
0010	-9
0110	-7
0111	-5
0101	-3
0100	-1
1100	1
1101	3
1111	5
1110	7
1010	9
1011	11
1001	13
1000	15

Input bits $b_4b_5b_6b_7$	Q-out
0000	-15
0001	-13
0011	-11
0010	-9
0110	-7
0111	-5
0101	-3
0100	-1
1100	1
1101	3
1111	5
1110	7
1010	9
1011	11
1001	13
1000	15

11.2.3.2.2 Spatial Parsing with Unequal Puncturing Rates

Puncturing is performed after parsing, as shown in Figure 72 (which is different from basic BF-MIMO, as shown in Figure 71). Parsing is done in cycles. During each cycle, a block of coded bits is parsed to each spatial stream. These blocks are taken to be integer multiples of the puncture cycle of the code; that is, for a rate n/(n+1) puncturing, the puncture cycle is n trellis stages, or 2n un-coded bits (the mother code is rate ½). Parameters of the parser are selected to

yield the smallest parsing cycle length, hence ensure good mixing of spatial streams within the constraint length of the code.

11.2.4 Low-Density Parity-Check Coding (Optional Mode)

Optional Low-Density Parity-Check (LDPC) coding may be used with all MCS's from Table 26. Utilization of LDPC codes with the optional advanced transmit beamforming (clause 11.2.3) is TBD. When encoding specified in this clause is applied, convolutional encoder and associated puncturing and padding shall not be used.

11.2.4.1 Code Rates and Codeword Lengths

Code rate (R) of the LDPC code may be selected from the set of code rates defined for mandatory convolutional codes (1/2, 2/3, 3/4 and 7/8). One of two codeword sizes is used to encode packets of length specified by HT-LENGTH parameter. Codeword size in bits, N, may be 576 or 1728, depending on the value of HT-LENGTH parameter, as follows:

If HT-LENGTH > 226 bytes, N = 1728 bits, otherwise N = 576 bits.

11.2.4.2 PPDU Encoding using LDPC Codes

This subclause, which specifies the LDPC encoding, replaces corresponding procedure defined for mandatory convolutional codes.

LDPC codes are block codes that are specified by the parity check matrix of size MxN , which defines an LDPC block code (N,K), where K is the information block size in bits, N is the length of the codeword, and M is the number of parity check bits, M=N-K. Code rate is defined as R=K/N. The general characteristic of an LDPC parity check matrix is a low density of non-zero elements (1's in this case), which allows utilization of efficient decoding algorithms. LDPC parity check matrix is designed in a way that enables simple encoding as well. It can be represented in the following form:

$$\mathbf{H} = [\mathbf{H}^d \mid \mathbf{H}^p]$$

$\mathbf{H}^d$ is an MxK matrix and corresponds to the "data" bits of the codeword. The design of this matrix ensures high coding gain. $\mathbf{H}^p$ is an MxM modified "dual diagonal" matrix, which corresponds to the "parity" bits of the codeword. LDPC codes applied here are systematic and canonic block codes. Codewords have the following structure:

$$\mathbf{c} = \begin{bmatrix} \mathbf{d} \\ \mathbf{p} \end{bmatrix}$$

where $\mathbf{d} = [d_0, d_1, ..., d_{K-1}]^T$ is the block of (uncoded) data bits and $\mathbf{p} = [p_0, p_1, ..., p_{M-1}]^T$ are the parity bits. A codeword is any binary N-vector $\mathbf{c}$ that satisfies:

$$\mathbf{Hc} = \mathbf{H}^d\mathbf{d} + \mathbf{H}^p\mathbf{p} = \mathbf{0}.$$

Since the information packet size may not in general be a multiple of number of information bits per codeword, procedure for encoding is defined in this subclause. The encoding procedure uses the following parameters as input:

HT-LENGTH – the number of information octets to encode and transmit using selected MCS
MCS – modulation type and code rate set

Based on the MCS, following two parameters are derived first:

$N_{CBPS} = N_{BPSC} * N_{SD} * N_{SS}$, number of coded bits per OFDM symbol
$N_{DBPS} = N_{CBPS} * R$, number of information bits per OFDM symbol

where: N_{BPSC} is the number of coded bits per sub-carrier, N_{SD} is the number of data sub-carriers, and N_{SS} is the number of spatial streams.

Both shortening (reduction of number of bits in the information portion of the codeword) and puncturing (reduction of number of bits in the parity portion of the codeword) may be applied in order to ensure minimum overhead without degrading performance. If N_{DBPS} is above a pre defined threshold, only shortening is used. The procedure for encoding N_{INFO_BITS} = HT-LENGTH∗8 consists of the following steps:

a) Encoding parameters computation

Determine the number of codewords, N_{CWORDS}, to transmit. For each of the codewords of length $N_1 \le N$ determine amount of the information bits to encode, $K_1 \le K$, as well as the number of parity check bits to be transmitted, $M_1 \le M$ ($M_1 = N_1 - K_1$). This, in turn, determines by how many bits to shorten information portion of the codeword (number of zeros to pad the information field), N_{PAD_CW}, as well as how many bits to puncture (remove bits from the parity portion of the codeword), $N_{PUNCTURE_CW}$. In some cases it may be required that the last codeword has different number of data and parity bits, K_{1_LAST} and M_{1_LAST} , respectively, (and consequently $N_{PAD_CW_LAST}$ and $N_{PUNCTURE_CW_LAST}$) than the other codewords.

```
if NDBPS ≤ 216
        NOFDM = ceiling(NINFO_BITS /NDBPS) + 1
        NCWORDS = ceiling(NOFDM*NCBPS/N)
        N1 = floor(NOFDM*NCBPS/ NCWORDS)
        K1 = ceiling(NINFO_BITS/ NCWORDS)
        NPAD_CW = K - K1
        NPUNCTURE_CW = max(0,(N-N1- NPAD_CW))
        if NCWORDS > 1
                K1_LAST = NINFO_BITS - K1*(NCWORDS-1)
                NPAD_CW_LAST = K - K1_LAST
                N1_LAST = NOFDM*NCBPS - N1*(NCWORDS-1)
                NPUNCTURE_CW_LAST = max(0,(N-N1_LAST- NPAD_CW_LAST))
        end
else
        NCWORDS = ceiling(NINFO_BITS /K)
        K1 = ceiling(NINFO_BITS/ NCWORDS)
        NPAD_CW = K - K1
```

if $N_{CWORDS} > 1$
$K_{1_LAST} = N_{INFO_BITS} - K_1*(N_{CWORDS}-1)$
$N_{PAD_CW_LAST} = K - K_{1_LAST}$
end
end

b) Last OFDM symbol padding field parameter computation

If the total number of coded bits, $N_{CODED_BITS_TOTAL}$, doesn't fit integer number of OFDM symbols, then additional padding (scrambled zeros) bits are added to the last OFDM symbol. Number of those bits, $N_{PAD_OFDM_LAST}$, is computed as follows:

if remainder($N_{CODED_BITS_TOTAL}$,N_{CBPS}) > 0
$N_{PAD_OFDM_LAST} = N_{CBPS}$ - remainder($N_{CODED_BITS_TOTAL}$,N_{CBPS})

where $N_{CODED_BITS_TOTAL}$ is computed as follows:

if $N_{CWORDS} > 1$
$N_{CODED_BITS_TOTAL} = N_1*(N_{CWORDS}-1) + N_{1_LAST}$
else
$N_{CODED_BITS_TOTAL} = N_1*N_{CWORDS}$
end

$N_{PAD_OFDM_LAST}$ zeros are appended to N_{INFO_BITS} from the input data buffer and passed to the scrambler.

c) LDPC encoding

Generate all N_{CWORDS} codewords by applying any of the encoding methods, such that $\mathbf{Hc}_i = \mathbf{0}$, where i=0, 1, …, N_{CWORDS}-1 with shortening and puncturing as necessary. **H** is the parity check matrix for the selected codeword size and code rate. Only one **H** matrix is used for the entire packet. The information vector for the i-th codeword to be encoded is formed by concatenating K_1 scrambled information bits and $N_{PAD_CW} = K - K_1$ zeros:

$$\mathbf{d}^i = [\underbrace{d^i_0, d^i_1, ..., d^i_{K_1-1}}_{K}, \overbrace{0,0,...,0}^{N_{PAD_CW}}]^T$$

After the encoding, obtained codeword is punctured (if necessary) by removing the last $N_{PUNCTURE_CW}$ parity bits:

$$\mathbf{c}^i = [\underbrace{d^i_0, d^i_1, ..., d^i_{K_1-1}}_{K}, \overbrace{0,0,...,0}^{N_{PAD_CW}}, \underbrace{p^i_0, p^i_1, ..., p^i_{M_1-1}}_{M}, \overbrace{p^i_{M_1}, p^i_{M_1+1}, ..., p^i_{M-1}}^{N_{PUNCTURE_CW}}]^T$$

Zeros used for padding data part and punctured bits are not transmitted over the air. Rather, a concatenation of K_1 data bits and M_1 parity bits is transmitted only:

$$\mathbf{c}^i_{transmit} = [\underbrace{d^i_0, d^i_1, ..., d^i_{K_1-1}}_{K_1}, \underbrace{p^i_0, p^i_1, ..., p^i_{M_1-1}}_{M_1}]^T$$

d) Codeword and the last OFDM symbol padding field concatenation

Concatenate N_{CWORDS} codewords and $N_{PAD_OFDM_LAST}$ padding bits (denoted as q's below) to form a vector of bits to be transmitted:

$$\mathbf{c}_{transmit} = [d_0^0,...,d_{K_1-1}^0, p_0^0,..., p_{M_1-1}^0,.....,d_0^{N_{CWORDS}-1},...,d_{K_{1_LAST}-1}^{N_{CWORDS}-1}, p_0^{N_{CWORDS}-1},..., p_{M_{1_LAST}-1}^{N_{CWORDS}-1}, q_0,...,q_{N_{PAD_OFDM_LAST}-1}]^T$$

Bit d_0^0 is transmitted first.

11.2.4.3 Parity Check Matrices

Parity check matrices **H** used in Step c) of the encoding procedure, described in sub-clause 11.2.4.2, are derived from one of the "base" parity check matrices, $\mathbf{H_b}$, specified below. One base matrix is defined per code rate. Size of a base parity check matrix is denoted as M_bxN_b. N_b, the number of columns in the base matrix, is fixed for all code rates, N_b=24. M_b, the number of rows in the base matrix, depends on the code rate as follows: $M_b = N_b(1-R)$. Parity check matrix **H** of size MxN is generated by expanding the base matrix for the selected rate, $\mathbf{H_b}$, z-times: $z=N/N_b=M/M_b$. The expansion operation is defined by element values of the base matrix. Each non-negative base matrix element, s, is replaced by a zxz identity matrix, $\mathbf{I}_z$, cyclically shifted to the right $s'=s_{mod(z)}$ times. Each negative number (-1) in the base matrix is replaced by a zxz zero matrix, $\mathbf{0}_{zxz}$. For the codeword of size 576 bits, z = 24 and for the codeword of size 1728 bits, z = 72.

Base matrices specification

Rate 1/2: M_bxN_b = 12x24

```
 0   0  -1   0  -1   0  -1  -1  -1   0  -1  -1   1   0  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1
29  -1   0  26  -1  -1   0  -1   0  -1  -1  -1  -1   0   0  -1  -1  -1  -1  -1  -1  -1  -1  -1
-1  -1  -1  21   0  -1  17  -1  -1  38  -1   0  -1  -1   0   0  -1  -1  -1  -1  -1  -1  -1  -1
43  -1  -1  30  -1  -1  -1   0  -1  41   0  -1  -1  -1  -1   0   0  -1  -1  -1  -1  -1  -1  -1
 5  -1   1  -1  -1  20  35  -1  -1   2  -1  -1  -1  -1  -1   0   0  -1  -1  -1  -1  -1  -1  -1
-1  46  -1  -1  -1  -1  22  -1  40   8  -1  -1   0  -1  -1  -1  -1   0   0  -1  -1  -1  -1  -1
-1  -1  -1   9  -1  -1  18  13  -1  35  -1  27  -1  -1  -1  -1  -1  -1   0   0  -1  -1  -1  -1
 2  -1  44  -1  -1  -1  27  -1  -1  25  18  -1  -1  -1  -1  -1  -1  -1  -1   0   0  -1  -1  -1
33  35  -1  29  -1  -1  16  -1  -1  -1  -1  30  -1  -1  -1  -1  -1  -1  -1  -1   0   0  -1  -1
-1  -1  -1   4   4  -1  -1  -1  15  17  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1   0   0  -1
 5  -1  -1  19  -1  14  -1  -1  -1  -1  11  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1   0   0
10  -1  -1  -1  21  -1  18   8  -1  -1  -1  -1   1  -1  -1  -1  -1  -1  -1  -1  -1  -1  -1   0
```

Rate 2/3: M_bxN_b = 8x24

```
68  -1  -1  24  -1  -1  64  -1   2  -1  11  -1  -1  44  -1   3   1   0  -1  -1  -1  -1  -1  -1
-1  14  57  -1  -1  66  -1  -1  32  -1  13  -1  55  -1  67  -1  -1   0   0  -1  -1  -1  -1  -1
-1   4  -1  -1  20  -1  19  -1  -1  44  -1  19  -1  30  18  -1  -1  -1   0   0  -1  -1  -1  -1
-1  -1  26  -1  10  21  -1  -1  -1  -1  56  54  54  -1  -1  37  -1  -1  -1   0   0  -1  -1  -1
19  -1  -1  -1  31  -1  -1  50  27  -1  -1  40  -1   4  -1  -1   0  -1  -1  -1   0   0  -1  -1
-1  -1  37  18  -1  -1  -1  41  -1  68  -1  -1   2  -1  50   5  -1  -1  -1  -1  -1   0   0  -1
67  -1  -1  -1  -1   4  68  -1  -1  69  -1  34  -1  40  -1  61  -1  -1  -1  -1  -1  -1   0   0
-1  45  -1  26  -1  -1  -1  16  -1  35  26  -1   5  -1  68  -1   1  -1  -1  -1  -1  -1  -1   0
```

Rate 3/4: M_bxN_b = 6x24

```
 0  23  -1  -1   0  47  53   1  -1   5  -1  13  66  34  -1  62  65  -1   1   0  -1  -1  -1  -1
14  57   7  11  -1  -1  69  -1  25  34  -1  25  13  -1  19  70  -1  35  -1   0   0  -1  -1  -1
59   3  -1  18  -1  43  48  58  -1  29  -1  54  49  -1  25   0  66  -1  -1  -1   0   0  -1  -1
15   1   2  -1  -1  39  19  -1  37  64  52  -1  28   4  -1  21  20  -1   0  -1  -1   0   0  -1
22  64  -1  48  60  -1  36  -1  15  62  68  -1  46  55  -1  41  -1   7  -1  -1  -1  -1   0   0
68  36  57  -1  55  -1  41  34  -1  45  46  -1   8  -1  37  44  -1  15  -1  -1  -1  -1  -1   0
```

Rate 7/8: $M_b x N_b = 3x24$

```
 34  61  67  67  12  35  62  61  37  18  51  23  22  31  67  32   7  63  67  14  57   1   0  -1
  1  62  14  71  17  42  50  39  33  27  43  33  15  10  61  38  53  53  16  70  46   0   0   0
  8  65  34   8   9  30  63  57  24   0  45  31  49  24  42  66  40  70   0  28  69   1  -1   0
```

(END)