



Sponsor: Director, Electronic Warfare,
OUSD(A&S)/A/Platform & Weapon
Portfolio Management (P&WPM)
Dept. No.: T841
Contract No.: W56KGU-18-D-0004-0001
Project No.: 0718D390-A2, 0719D390-A2

The views, opinions and/or findings
contained in this report are those of The
MITRE Corporation and should not be
construed as an official government position,
policy, or decision, unless designated by
other documentation.

**Approved for Public
Release; Distribution
Unlimited. Public Release
Case Number 19-0991**

©2020 The MITRE Corporation.
All rights reserved.

McLean, VA

Spectrum Highways

Rules of the Road for Collaborative Radio Frequency Spectrum Sharing among Cooperating Spectrum-Dependent Systems

John A. Stine

December 2019

Abstract

This report defines a novel dynamic spectrum access technique that uses the highway metaphor as the basis for design and applications. As described in this report, a spectrum highway consists of a band of spectrum over a geographical area that is subdivided into smaller bands called lanes and those lanes support access to any number of spectrum-dependent systems (SDSs). The devices that comprise the SDS follow rules that govern their selection of spectrum highways and lanes and the way devices contend for access using signaling. With this construct, signaling is very effective at arbitrating access. This signaling differentiates access based on the precedence of the contenders and resolves contention. These highways support spectrum use in very congested and dynamic spectrum environments.

This page intentionally left blank.

Executive Summary

Access to electromagnetic spectrum (EMS) is a key component of most technical innovation whether commercial, civil, or military. With today's processes, each of these uses requires a specific assignment in an already congested and shrinking space. In the military context access is further complicated by the contest for spectrum against adversaries. The traditional method of finding a unique and fixed assignment for all spectrum uses is no longer viable and represents a vulnerability in warfighting where key system spectrum assignments can be identified and denied.

Dynamic spectrum access (DSA) is heralded as a technical solution to the shortage of spectrum, but are current DSA systems up to the challenge? The core idea of DSA is to enable spectrum dependent systems (SDSs) to dynamically access spectrum that is not in use, aka whitespace, assuming incumbent uses are not persistent and will have downtime. The key technology is the means to identify and move SDSs to the whitespace.

Commercial vendors and national spectrum managers currently favor two approaches. In the first, a database attempts to monitor all use and then grants access to SDSs in whitespace it identifies. In the second, the SDS devices sense their environment and autonomously move to the whitespace they detect. Both approaches have severe deficiencies associated with truly knowing if what they track or detect provides sufficient information to make decisions that still protect incumbent users. Sensing approaches do not detect passive use by sensors and highly sensitive receivers. Only the database approach can distinguish among multiple competitors and determine precedence, but that process is too slow to react and arbitrate those decisions. Sensing DSA systems can make faster decisions but still may not be able to react quickly enough because of the need to coordinate the rendezvous of multiple devices to the same channel. When devices have different perspectives and sense different activity, rendezvous may be intractable. These approaches collapse in the presence of congestion. Finally, they are easily defeated in a contested environment. Electronic warfare systems can track and move faster than DSA systems. In the case of sensing DSA systems, adversaries can manipulate systems with false indications of use and corral them for easy denial at critical times. With the database approach, just maintaining the network that connects devices to the database will require significant engineering and resource allocations, but with an adversary the network is an obvious and easy opportunity to deny that access and cause widespread system failures. With these approaches, DSA is not a viable answer for the military.

The spectrum highway system is a new approach to DSA that is highly efficient at managing the access and reuse of spectrum in time and in space. Additionally, it resolves problems of congestion, distinguishes spectrum use based on precedence of users and uses, coordinates the rendezvous of SDS devices, avoids features that can be easily attacked and enables autonomous maneuver in the EMS. Further, the underlying technology that enables spectrum highways will fully support heterogeneous uses in the same spectrum, making spectrum convergence feasible even among distributed autonomous devices.

This DSA technique is most easily visualized using highways (or roadways) as a metaphor. Following this metaphor, highways are created by setting aside spectrum in geographical locations and subdividing that spectrum into lanes. Like vehicles on a highway, many SDSs can operate on the same lane of the highway, with their use separated either spatially or temporally,

or on separate lanes for spectral separation. SDSs can autonomously move among the lanes and, if necessary, merge lanes for broader bandwidth access. They can also move among highways. SDSs contend among themselves to use the highway lanes by using a signaling protocol that is highly effective at resolving contention in congested conditions and orchestrating spatial reuse of the highway lanes. This same signaling is used to arbitrate precedence among the SDSs for lane access. Similar to the way an emergency vehicle can alert other vehicles to its presence and precedence using a siren and flashing lights, an SDS can use the signaling methods of the highway. The signaling can indicate which SDS a contending device belongs thus removing the challenges of rendezvous.

The behaviors of SDSs that cause them to cooperate efficiently are created by the access protocol, its rules for access behavior, and authorizations given to SDSs that govern their use of precedence signaling when using highways.

Given a collection of SDSs designed to operate on highways, the concept of operations is to provide those SDSs with information about the highways and the precedence levels they are permitted to use and then let them operate autonomously thereafter. There is no requirement for SDSs to remain connected to a database and the use of signaling removes any ambiguity about spectrum availability. In operation, each SDS autonomously decides which highways and lanes to use and cooperates with others on the same highway and lane using the access protocol and its associated rules. There is no burden on system operators. This access paradigm can achieve the same results as access using a database, (e.g., the spectrum access system (SAS) used by the Citizen's Band Radio Service in 3.5 GHz), without the need for constant connectivity between the devices and the SAS. Better, the devices can be mobile, thus avoiding the expense of professional installation, localization, and registration.

Spectrum highway DSA can achieve the same results as a sensor-based DSA approach without risk of interfering with passive users. As long as SDSs follow the rules for access and behavior on the highway, their use of the spectrum will be compatible with other SDSs on the highways. The autonomy and flexibility this approach provides, enable SDSs to have initiative in selecting the spectrum to use which gives them an advantage in a contest against an adversary.

As with any new concept, bringing it to fruition will require the Department of Defense (DoD) to first build the systems and verify they work collaboratively as intended. The report provides a comprehensive description of how to design highways and signaling to enable the outcomes described above. Next steps are to implement the concepts and refine the details.

Implementation efforts should follow a natural progression starting with simulation, then demonstration in a closed environment (e.g. such as the Defense Advanced Research Projects Agency Colosseum which was built in support of its Spectrum Collaboration Challenge), and finally open-air demonstration. DoD could sponsor research on components of the highways. An example of relevant research would be designing the physical layer of signals to support efficient and resilient signaling. A second example would be to design spectrum highway access devices (SHADs) and to standardize protocols between SHADs and SDSs in order to separate highway access from SDS function and so simplify the requirements for SDSs to qualify for operating on highways. Building and demonstrating the operation of SHADs would represent an appropriate surrogate for demonstrating the efficacy of spectrum highways.

Spectrum highways provide an exciting new approach to spectrum access giving devices autonomy to operate while ensuring that the outcome of multiple devices operating on the same highway does not result in deadlock or the collapse of capacity as a result of congestion. It provides means for devices to arbitrate who and what use should be supported. These features are useful in pure commercial systems, enabling operations in bands where government and commercial users share the spectrum, and are mission critical in the military contexts where SDSs autonomously maneuver in spectrum.

The use of spectrum highways with their rules and contention mechanisms will result in unmatched optimality and efficiency of DSA by heterogenous SDSs operating in a dynamic environment. The need for speed and agility of EMS maneuver within the DoD enterprise dictates the pursuit of such solutions.

Table of Contents

1	Introduction	1-1
1.1	Background.....	1-1
1.2	Advanced DSA Alternatives.....	1-3
1.3	Facts, Assumptions, and Requirements	1-4
1.3.1	Facts	1-4
1.3.2	Assumptions.....	1-4
1.3.3	Requirements	1-5
1.4	Highway Design Overview.....	1-6
2	Defining RF Spectrum Highways	2-1
2.1	Spectrum Consumption.....	2-1
2.2	Spectrum Availability	2-1
2.3	Spectrum Highway Boundaries	2-2
2.3.1	Highway Location and Spectrum.....	2-2
2.3.2	Spectral Lanes.....	2-3
2.3.3	Highway Schedule and Transmission Opportunities.....	2-6
2.4	Defining Spectrum Highway Boundaries	2-7
2.4.1	Highway Exemplar	2-9
2.4.2	Protocol or Policy Construct Overview	2-18
3	Channel Access Options.....	3-1
3.1	Centralized versus Distributed Control.....	3-1
3.2	Wireless Limitations	3-1
3.2.1	Singular Task	3-2
3.2.2	Channel Symmetry.....	3-2
3.2.3	Channel Sensing.....	3-2
3.2.3.1	Hidden Nodes	3-2
3.2.3.2	Exposed Nodes	3-3
3.3	Wireless Access Mechanisms	3-4
3.3.1	Scheduled Protocols.....	3-4
3.3.1.1	Time Division Multiple Access.....	3-4
3.3.1.2	Hybrid TDMA.....	3-5
3.3.2	Contention.....	3-6
3.3.2.1	Aloha	3-6

3.3.2.2	Slotted Aloha	3-7
3.3.2.3	Carrier Sense Multiple Access (CSMA)	3-8
3.3.2.4	Packet Sensing Multiple Access (PSMA)	3-10
3.3.2.5	Deficiencies of Carrier and Packet Sensing Protocols	3-11
3.3.3	Synchronous Collision Resolution (SCR)	3-13
3.3.3.1	Collision Resolution Signaling.....	3-13
3.3.3.2	Design of CRS Contention	3-15
3.3.3.3	Spatial Demonstration of CRS	3-22
3.3.3.4	Designing Priority Phases.....	3-23
3.3.3.5	SCR Challenges.....	3-25
3.4	Conclusions.....	3-26
4	Multi-System Contention Design.....	4-1
4.1	Assessment of MAC Approaches for Spectrum Highway Use	4-1
4.2	Access Design Alternatives to Achieve Requirements.....	4-2
4.2.1	Differentiation of Service and Precedence of Access.....	4-4
4.2.1.1	Lane Precedence	4-4
4.2.1.2	Lane Selection	4-4
4.2.1.3	Use Precedence.....	4-6
4.2.1.4	Quality of Service (QoS) Precedence.....	4-6
4.2.1.5	Reservations.....	4-6
4.2.1.6	Preemption of Reservations.....	4-6
4.2.2	Contention Channels.....	4-7
4.2.2.1	In-Band Channels	4-7
4.2.2.2	Out-of-Band Channels.....	4-8
4.2.2.3	Consolidated Signaling.....	4-11
4.2.2.4	Resilient Channels	4-12
4.2.3	Contending for Alternative Duration Access.....	4-13
4.2.4	Spatial Reuse.....	4-14
4.2.4.1	Choosing Spatial Separation.....	4-14
4.2.4.2	Supporting Asymmetric Access Because of Directional Antennas.....	4-15
4.2.4.3	Support Asymmetric Access Because of Different Signaling Powers	4-15
4.2.5	Changing Lanes	4-15
4.2.6	Effectiveness in Mobile Environments.....	4-17

4.2.6.1	Mobility and Slot Reservations	4-17
4.2.6.2	Mobility with Frame and Epoch Access.....	4-17
4.3	Signaling Physical Layer Design Alternatives	4-18
4.3.1	Signal Requirements	4-18
4.3.2	Signal Bandwidth Shape, Duration, and Range.....	4-18
4.3.3	Minislot Sizing.....	4-18
4.3.4	Frequency-Hopping Alternatives.....	4-20
4.4	Performance Measures.....	4-21
4.5	Design Trades	4-22
4.6	Conclusions.....	4-24
5	Specifying Highway Access Rules.....	5-1
5.1	The Protocol or Policy Construct.....	5-1
5.1.1	The IEEE 1900.5.2 Description of the Data Elements and their Meaning	5-1
5.1.2	Highway Application.....	5-3
5.2	Identifying Highway Access Feature and Rule Options.....	5-3
5.2.1	PorPName	5-3
5.2.2	PorPParameters.....	5-3
5.2.2.1	Timeslot Duration and Hierarchy	5-3
5.2.2.2	Contention Method.....	5-5
5.2.2.3	Lane Use Precedence Phase Set	5-5
5.2.2.4	Lane Selection Phase Set.....	5-7
5.2.2.5	Priority Phase Set.....	5-10
5.2.2.6	Contention Phase Set.....	5-13
5.2.2.7	Contender ID Phase Set.....	5-15
5.2.2.8	Contention Channels.....	5-17
5.2.2.9	Range Symmetry	5-18
5.3	Distinguishing SDS Precedence Authorizations.....	5-20
5.3.1	Rules for SDS Use of a Spectrum Highway	5-23
5.3.1.1	Considerations for Highway Selection.....	5-23
5.3.1.2	Rules for Access to Spectrum Highways	5-24
5.3.1.3	Rules for Precedence Behavior on Spectrum Highways.....	5-24
5.4	Conclusion	5-27
6	Device Access Authorizations.....	6-1

6.1	Routine Mission SDSs	6-1
6.1.1	Communications SDS.....	6-1
6.1.2	Sensing SDS.....	6-4
6.2	Mission-Critical SDSs	6-5
6.2.1	Radar SDS.....	6-5
6.3	Reactive Mission SDSs.....	6-7
6.3.1	Counter Radio-Controlled IED Electronic Warfare (CREW) SDS.....	6-7
6.4	Conclusion	6-8
7	Summary and Next Steps	7-1
7.1	Applications	7-2
7.1.1	Three-Tier Spectrum Sharing	7-2
7.1.2	Intermediate-Range Spectrum Commons.....	7-3
7.1.3	Spectrum Maneuver Spaces.....	7-3
7.1.4	Efficient Spectrum Management	7-3
7.2	Further Design Work	7-4
7.2.1	Highway SDSs.....	7-4
7.2.2	Signal Designs	7-4
7.2.3	SHAD Design	7-5
7.2.4	SHAD/SDS Communication	7-5
7.2.5	SDS Precedence Guidance.....	7-6
7.2.6	Standardization	7-7
7.2.7	Regulation.....	7-7
7.3	Approaches to Validate Spectrum Highway Efficacy	7-7
7.3.1	Modeling and Simulation.....	7-7
7.3.2	Small Scale Demonstration.....	7-7
7.3.3	Spectrum Colosseum	7-8
7.3.4	Live Experimentation.....	7-8
7.4	Conclusion	7-8
8	References/Bibliography	8-1
Appendix A	Highway Design Performance Measures.....	A-1
A.1	Temporal Overhead	A-1
A.2	Spectral Overhead.....	A-1
A.3	Use Efficiency.....	A-1

A.4	Temporal Effectiveness	A-1
A.5	Congestion	A-2
A.6	Responsiveness	A-2
A.6.1	Access Delay.....	A-2
A.6.2	Slot Delay.....	A-2
A.7	Spatial Capacity	A-2
A.8	Spectral Agility	A-3
A.9	Asymmetry Risk	A-3
A.9.1	Power Asymmetry	A-3
A.9.2	Weighted Power Asymmetry.....	A-3
A.9.3	Directional Power Asymmetry	A-4
A.9.4	Adaptive Directional Asymmetry.....	A-4
A.10	Contention versus Device Asymmetry	A-4
A.11	Multifunction Count.....	A-5
A.12	Collision Resolution.....	A-5
A.13	Precedence Count.....	A-5
A.13.1	Reservation Count.....	A-6
A.14	Resilience.....	A-6
A.15	Mobility.....	A-6
Appendix B	Examples of In-Band Signaling Designs	B-1
B.1	Access Rule Policy	B-1
B.2	Timing Hierarchy and Contention Method.....	B-2
B.3	Lane Use Precedence Phase Set.....	B-5
B.4	Lane Selection Phase Set	B-6
B.5	Priority Phase Set.....	B-7
B.6	Contention Phase Set	B-12
B.7	Contention Channels.....	B-13
B.8	Range Symmetry.....	B-19
B.9	Design Metrics	B-20
B.9.1	Temporal Overhead	B-20
B.9.2	Spectral Overhead.....	B-21
B.9.3	Use Efficiency.....	B-21
B.9.4	Access Delay.....	B-22

B.9.5	Slot Delay.....	B-22
B.9.6	Responsiveness	B-22
B.9.7	Spatial Capacity	B-22
B.9.8	Spectral Agility	B-23
B.9.9	Asymmetry Risk	B-23
B.9.10	Contention versus Device Asymmetry	B-23
B.9.11	Multifunction Count.....	B-23
B.9.12	Collision Resolution.....	B-24
B.9.13	Precedence Count.....	B-24
B.9.14	Reservation Count.....	B-25
B.9.15	Resilience.....	B-26
B.9.16	Mobility.....	B-26
B.10	Conclusion	B-26
Appendix C	Examples of Out-of-Band Signaling Designs.....	C-1
C.1	Access Rule Policy	C-1
C.2	Timing Hierarchy and Contention Method.....	C-2
C.3	Lane Use Precedence Phase Set.....	C-4
C.4	Lane Selection Phase Set	C-4
C.5	Priority Phase Set.....	C-6
C.6	Contention Phase Set	C-10
C.7	Contender ID Phase Set	C-11
C.8	Contention Channels.....	C-14
C.9	Range Symmetry.....	C-15
C.10	Design Metrics	C-16
C.10.1	Temporal Overhead	C-16
C.10.2	Spectral Overhead	C-17
C.10.3	Use Efficiency.....	C-17
C.10.4	Access Delay.....	C-17
C.10.5	Slot Delay.....	C-18
C.10.6	Responsiveness	C-18
C.10.7	Spatial Capacity	C-18
C.10.8	Spectral Agility	C-18
C.10.9	Asymmetry Risk	C-18

C.10.10	Contention versus Device Asymmetry	C-18
C.10.11	Multifunction Count.....	C-19
C.10.12	Collision Resolution.....	C-19
C.10.13	Precedence Count.....	C-20
C.10.14	Reservation Count.....	C-20
C.10.15	Resilience.....	C-21
C.10.16	Mobility.....	C-21
C.11	Conclusion	C-21
Appendix D	Glossary	D-1
Appendix E	Abbreviations and Acronyms	E-1

List of Figures

Figure 2-1. Geographical Spectrum Highways.....	2-3
Figure 2-2. Spectrum Lanes.....	2-4
Figure 2-3. Spectrum Lanes with Skirts	2-5
Figure 2-4. Criteria for Merging Lanes.....	2-5
Figure 2-5. Example of Three Levels of Transmission Opportunities	2-6
Figure 2-6. Example of Different Duration Transmission Opportunities at a Level	2-6
Figure 2-7. SCM Model Hierarchy.....	2-8
Figure 2-8. Exemplar of a Highway Location	2-9
Figure 2-9. Exemplar Highway Lanes	2-10
Figure 3-1. Terrestrial Exposed Node.....	3-3
Figure 3-2. Additional Exposed Node Scenarios.....	3-4
Figure 3-3. Typical Performance of Access Mechanisms Vulnerable to Congestion Collapse ..	3-7
Figure 3-4. Wireless Contention Access Scenario Showing Challenges in Using CA, Handshakes, and Virtual Sensing	3-12
Figure 3-5. Generic Implementation of SCR.....	3-13
Figure 3-6. CRS with One-Minislots Phases.....	3-14
Figure 3-7. CRS with Two-Minislots Phases, CRS with Echoing.....	3-14
Figure 3-8. CRS with Two-Minislots Phases, with both Priority and Contention Phases	3-15
Figure 3-9. CRS optimization alternatives for a k_1 contender density	3-16
Figure 3-10. Performance of the Six-Contention-Phase CRS Designs.....	3-18
Figure 3-11. Contention Collision Resolution Success Ratio of the Six-Contention-Phase CRS Designs.....	3-18
Figure 3-12. Performance of the Seven-Contention-Phase CRS Designs	3-19
Figure 3-13. Contention Collision Resolution Success Ratio of the Seven-Contention-Phase CRS Designs.....	3-19
Figure 3-14. Performance of the Eight-Contention-Phase CRS Designs	3-20
Figure 3-15. Contention Collision Resolution Success Ratio of the Eight-Contention-Phase CRS Design	3-20
Figure 3-16. Performance of the Nine-Contention-Phase CRS Designs	3-21
Figure 3-17. Contention Collision Resolution Success Ratio of the Nine-Contention-Phase CRS Design	3-21
Figure 3-18. Reduction of 500 Contending Devices to a Set of 13 Devices out of Range of Each Other	3-22
Figure 3-19. Simulation of an eight phase contention with echoing.	3-23
Figure 3-20. CRS and Frame Design for Periodic Slot Reservations.....	3-25
Figure 4-1. CRS Using Four Sets of Phases for Lane Precedence, Lane Selection, Use Precedence, and Contention Resolution	4-3
Figure 4-2. Example of a Channel Precedence Design.....	4-5
Figure 4-3. Example of In-Band Signaling.....	4-8
Figure 4-4. Example of Serial Out-of-Band Signaling.....	4-9
Figure 4-5. Demonstration of a Time Shift of the Serial Out-of-Band Signaling	4-9
Figure 4-6. Example of Serial Contention for Spectrum Highway Lanes.....	4-9
Figure 4-7. Example of Concurrent Out-of-Band Signaling	4-10

Figure 4-8. Coupling of an SDS with a SHAD to Arbitrate Spectrum Access	4-10
Figure 4-9. Example of in-band signaling using consolidated contentions	4-11
Figure 4-10. Alternative Schedules for Consolidated Contention in 10 Timeslot Frames	4-12
Figure 4-11. Identification of Slots Allowing Contentions for Timeslots, Frames, and Epochs... 4-14	
Figure 4-12. Using Additional Signaling Phases to Support Indicating an ID Number	4-16
Figure 4-13. Using a Modulated Signal to Provide an ID	4-16
Figure 4-14. Using a Single Signal on a Unique Frequency to Indicate the SDS ID	4-17
Figure 4-15. Signal Transmission with Maximum Lead Offset and Minimum Propagation	4-19
Figure 4-16. Signal Transmission with Maximum Lag Offset and Maximum Propagation	4-20
Figure 5-1. Example Priority Phase Set.....	5-11

List of Tables

Table 2-1. Spectrum highway description using a SCM Set	2-10
Table 2-2. Set Boundary for an SCM Set	2-11
Table 2-3. Band List That Provides the Spectral Boundary of the Highway	2-11
Table 2-4. Highway Schedule Specifying the Temporal Start and End of the Highway	2-12
Table 2-5. Boundary of the Highway Location	2-12
Table 2-6. Details of the Polygon Surface Used for the Location of the Spectrum Highway ...	2-12
Table 2-7. Maximum Antenna Height of Devices within the Spectrum Location	2-14
Table 2-8. SCM Transmitter Model use for Lane 1 of the Spectrum Highway	2-14
Table 2-9. Reference Power for Lane 1	2-14
Table 2-10. Spectrum Mask of Lane 1	2-15
Table 2-11. Power Map of Lane 1	2-16
Table 2-12. Omnidirectional Gain Map Used by Lane 1	2-17
Table 2-13. Spectrum Propagation Map for Lane 1	2-17
Table 2-14. Propagation Map Values for Lane 1	2-18
Table 3-1. CRS Design for Six Contention Phases.....	3-18
Table 3-2. CRS Design for Seven Contention Phases	3-19
Table 3-3. CRS Design for Eight Contention Phases	3-20
Table 3-4. CRS Design for Nine Contention Phases	3-21
Table 4-1. Purpose, Requirements, and Design Trades for the Four Possible Phase Sets.....	4-3
Table 4-2. Measures for Comparing Highway Designs.....	4-21
Table 4-3. Alternative highway design options and their performance trades	4-22
Table 5-1. SCM Policy or Protocol Construct.....	5-2
Table 5-2. Parameters for the Timing Hierarchy	5-4
Table 5-3. Parameters for the Contention Method.....	5-5
Table 5-4. Parameters for the Lane Use Precedence	5-6
Table 5-5. Parameters for the Lane Selection Phase Set	5-7
Table 5-6. Parameters for the Priority Phase Set	5-11
Table 5-7. Parameters for the Contention Phase Set.....	5-13
Table 5-8. Contention Phase Set Signaling Probabilities	5-14
Table 5-9. Parameters for the Contender ID Phase Set	5-15
Table 5-10. Parameters for Contention Channel Definition	5-18
Table 5-11. Parameters for Power Asymmetry.....	5-18
Table 5-12. Parameters Used to Govern SDS Device Access to Highways.....	5-20
Table 6-1. Example Access Authorizations for a Communications SDS.....	6-1
Table 6-2. Example Set of User Precedence Guidance for a Communications SDS	6-2
Table 6-3. Example Set of Use Precedence Guidance for a Communications SDS.....	6-2
Table 6-4. Example Set of QoS Precedence Guidance for a Communications SDS.....	6-3
Table 6-5. Example Access Authorizations for a Sensing SDS	6-4
Table 6-6. Example Set of User Precedence Guidance for a Sensing SDS	6-4
Table 6-7. Example Set of Use Precedence Guidance for a Sensing SDS	6-4
Table 6-8. Example Set of QoS Precedence Guidance for a Sensing SDS	6-5
Table 6-9. Example Access Authorizations for a Radar SDS	6-6
Table 6-10. Example Set of User Precedence Guidance for a Radar SDS	6-6

Table 6-11. Example Set of Use Precedence Guidance for a Radar SDS	6-6
Table 6-12. Example Set of QoS Precedence Guidance for a Radar SDS	6-6
Table 6-13. Example Access Authorizations for a CREW SDS	6-7
Table 6-14. Example Set of User Precedence Guidance for a CREW SDS	6-7
Table 6-15. Example Set of use Precedence Guidance for a CREW SDS	6-8
Table 6-16. Example Set of QoS Precedence Guidance for a CREW SDS	6-8
Table 7-1. Strategies for Splitting Awareness and Decision Making between SDSs and SHADs	7-6

This page intentionally left blank.

1 Introduction

This report describes a new approach to radio frequency (RF) spectrum sharing called the Spectrum Highway System. Under this approach multiple spectrum-dependent systems (SDSs)¹ of different kinds contend to use spectrum and through rules and signaling adjust their use collaboratively to a near-optimum state. In this state, the contending SDSs are packed to use all available RF bandwidth and to reuse that spectrum spatially while deferring to the most important users and uses.

This report begins in this section with background on the variety of methods for spectrum sharing with the intent of differentiating the spectrum sharing approach advocated here from others. It further lists the facts and assumptions that constrain the problem and delineates what the objective system must achieve. It concludes with a concise overview of what this report contains. Sections 2 through 4 provide technical background and analysis of the problems associated with building a rule set. Section 2 explores approaches to define the highways, the quanta of spectrum that the rules and signaling apply. Section 3 gives an overview of wireless access challenges and the various access mechanisms that exist in literature and in practice. Section 4 presents an analysis of the problem that culls through the design choices to find the most practical alternatives.

Section 5 is the most important: it explains the approach to specifying the design of the time slotting and the signaling, the rules that accompany those design elements, and the approach to authorizing SDSs and devices to use various precedence levels to activate lanes, gain access, and achieve quality of service (QoS). Appendix B provides examples using the methods of Section 5 to define a highway that uses in-band signaling. It also gives an example of a physical layer design of a signaling approach. Appendix C presents examples of defining a highway that uses out-of-band signaling. Both Appendices B and C conclude with a listing of metrics to support the comparison of the two designs.

Section 6 contains several examples of giving guidance to SDSs that govern their selection of and operation on highways and the precedence levels they can use to establish lanes and to get access ahead of other types of SDSs. Section 7 concludes the report by identifying the areas of highway design that are ripe for research and suggest some next steps to demonstrate the concept.

1.1 Background

Demand for RF spectrum is continuously increasing as existing services are more broadly used, as new services and uses are discovered, and as new SDSs are developed. In addition to satiating spectrum demand for new services and new SDSs, there is additional demand in military operations that is born out of the contest for electromagnetic spectrum (EMS) superiority. It involves identifying an adversary's use of spectrum and then denying, attacking, or exploiting that use while the adversary tries to do the same to friendly uses of spectrum. The activities

¹ The term SDS refers to one or multiple devices that operated as a system. A radar SDS may be one transmitter and one receiver. A communications network SDS will consists of multiple transceivers. Each transceiver is a device in the SDS.

within this contest can be rapid and the speed of their execution is likely to determine the victor. SDSs must be able to move fluidly through the spectrum to avoid detection and attack.

Traditionally, meeting the growing demand for spectrum has involved identifying bands that can be allotted exclusively to an RF service operated by a defined user community. This allotment is preceded by clearing and relocating legacy users in the band. The exclusive use model remains the strategy preferred by commercial service providers such as telecommunications companies, wireless Internet service providers (WISPs), and broadcasters. However, across a vast set of services that are intermittent rather than persistent, sharing is necessary because it is not feasible to move all users into exclusive spectrum. Further, in warfighting, where adversaries contend for spectrum, dynamic access, sharing, and support for maneuver in spectrum is likely to be key to achieving EMS superiority.

Dynamic spectrum access (DSA) is not a new concept, and to date has mostly used two technical approaches: database control and sense and avoid. The effectiveness of both approaches is undermined by deficiencies in tracking and detecting active spectrum use to make decisions that protect incumbent and higher priority uses. Neither approach competes well in contested environments.

In the database control approach, SDSs are connected via a reliable communications link, usually wireline, through which the SDSs receive commands from a database administrator. The commands consist of grants for spectrum use and directives to clear spectrum when it must be available for another higher priority system. Examples include the TV Whitespace database and the 3.5 GHz Spectrum Access System (SAS). The database control approach has not yet been employed across military systems and its reliance on a reliable control channel makes it impractical for the military in most cases.

In the sense-and-avoid approach, SDSs sense their environment and based on their observations choose a free band in which to operate. The alternative choices are constrained by the capabilities of the particular radio and the policy applied to the radio prior to operation. Several military systems have been developed with policy control, but they were largely experimental and have not been fielded. Challenges remain in certifying this type of equipment. Regardless of the constraint, the radio can only detect uses that involve transmission, so systems that receive faint signals (e.g. satellite ground terminals) and systems that assess the electromagnetic environment (EME) and want silence would not be protected without other coordination means. Further, in the case of communications, multiple radios must make the same decision as to which spectrum to use or else they will not be able to communicate to each other. Causing the radios, especially the multiple radios comprising a network, to rendezvous at the same channel is challenging and represents a major hurdle except in the most benign environments. These systems do not work well in congested environments because there is no whitespace to move to or the infrequent whitespace triggers many access attempts that collide and are harmful to each other. In a contested environment, adversaries can create an artificial congested environment. More dastardly, they could corral SDSs into a subset of spectrum of their choosing and then deny that spectrum when it is advantageous.

1.2 Advanced DSA Alternatives

An alternative to these approaches to DSA is to create the means for SDSs to collaborate in sharing spectrum. Under this paradigm, SDSs communicate with each other and negotiate in some way to arbitrate the access of each independent SDS to spectrum. An objective of collaboration is to achieve more effective sharing in congested conditions and to ensure the most important uses get precedence. There are two approaches, direct communication and negotiation or using a combination of rules and contention. Spectrum highways, the subject of this report, are an example of the second.

The Defense Advanced Research Projects Agency (DARPA) Spectrum Collaboration Challenge (SC2) project offers an example of the direct negotiation approach. In SC2, radio networks maneuver in the same operational space and uses machine learning (ML) to understand the external activities of other systems and their use of spectrum and to collaborate with other systems in sharing spectrum. The sharing is expected to balance the needs of each system. Systems are scored based on whether the set of competing networks perform well in aggregate. Facilitating this collaboration is a dedicated channel that reliably connects the different systems. SDSs use this channel to communicate their needs to each other and through what is anticipated to be a game theoretic algorithm informed by ML choose a channel and bandwidth that best balance the capacity among the competing systems based on their collective needs. Resolving an operating point that best balances competing demands may require several iterations of exchanging information. As in database management, dependence on a reliable communications network for collaboration can make the approach less viable in a military operations environment.

The second alternative to collaboration is similar to the way vehicles share a highway. Highways are designed for such sharing. Like vehicular highways in which the area of the road itself is subdivided into lanes, spectrum highways set aside spectrum for a geographical area and that spectrum is subdivided into sub bands that constitute the spectrum lanes. Just as multiple vehicles can share a highway by operating in the lanes and following rules, so too can a variety of SDSs operate on a spectrum highway. Just as the operator of each vehicle acts independently making choices that optimize his or her use of the highway so too do the SDSs operating on the spectrum highway. Just as some vehicles (e.g., ambulances) can be given precedence on a highway, so too can precedence be given to an SDS. The actions of each vehicle operator are quick; and, through a distributed set of decisions made by the operators, the highway supports the transport needs of the many vehicles. Similarly, a spectrum highway can support the spectrum needs of many types of SDSs as long as their use of spectrum stays within the boundary of the highway. By applying the spectrum highway collaboration approach, SDSs can achieve agility in their use of spectrum and can maneuver among spectrum highways and their lanes to find the best spectrum in which to operate. SDSs may decide to change lanes and then do so extremely rapidly. Precisely this autonomous decision making and rapid execution could make the spectrum highway approach a key enabler to achieving spectrum superiority.

The spectrum highway approach to collaborative spectrum sharing is new and MITRE is not aware of any previous attempt to build such a rule set for spectrum access among a heterogeneous set of SDSs.

1.3 Facts, Assumptions, and Requirements

A heterogeneous set of SDSs will have a variety of overlapping and unique characteristics. Most RF systems are designed without means to sense or negotiate with other systems. The following list of facts, assumptions, and requirements are the starting point of designing the ruleset and signaling. To participate, SDSs will have to be upgraded with the ability to follow the rules and with the means to send and detect signals.

1.3.1 Facts

SDS transceiver devices that operate on a single channel cannot listen and transmit at the same time on that channel and so cannot observe another's use of spectrum while transmitting, cannot detect collisions that occur during transmissions, and cannot receive warnings from peer SDSs.

Some SDSs and their activities, (e.g., sensing) are not detectable by peer systems.

SDSs have different characteristics and functions and therefore different performance. Their consumption of spectrum varies because of different transmission ranges and different sensitivities to interference.

SDSs have different spectral bandwidth requirements.

1.3.2 Assumptions

SDSs may be paired with separate devices to support signaling among the participating SDSs.

Some SDSs will use separate devices to support collaboration; those devices may not be co-located with the primary SDS in order to protect the location and presence of the primary SDS.

All sensing and signaling will be done over-the-air in RF spectrum with propagation characteristics that are comparable to or better than those of the primary SDS spectrum.

The total range of signaling will be equal to or greater than the range of the primary SDS transmissions.

Adversaries may try to observe or to disrupt the functioning of the spectrum highway.

Highways will be designed to avoid interfering with legacy SDSs.

Only SDSs that have been modified to participate in the sharing protocol will be allowed on the spectrum highway. All others will be considered unauthorized users.

Highways may be designed to accommodate particular types of SDSs.

SDSs may be designed or configured to best exploit established highway designs.

In execution, the highway boundaries may partially cover the range of the bands in which a device is capable of operating. To use the highway the device may only use the bands within the highway.

SDSs can meet the timing and location awareness requirements of the highway design they participate, generally they have time of day accuracy specified for the highway and contention and they can confirm they are within the geographical boundary of the highway.

1.3.3 Requirements

The requirements listed below guided the creation of MITRE's approach to designing highways. Experience and additional operational requirements may result in future revisions.

1. The "spectrum highway" shall accommodate a variety of SDSs including electronic attack (EA) and electronic support (ES) SDSs.
2. Specific rules shall cover cases in which all sensing and signaling is done in the bands used by the SDSs. This case must accommodate SDSs using their own transceiver for signaling, so signaling must occur between payload transmissions.
3. Specific rules shall cover cases where all SDSs use a separate, homogeneous set of transceivers for the purposes of collaboration in resolving access.
4. Rules shall support SDSs that use different spectrum bandwidth and shall pack uses so that all bandwidth of the spectrum highway is used when there is demand and it is feasible to use that bandwidth without harmful interference.
5. The spectral bands of the "spectrum highway" shall be predefined and discrete. Spectrum highways may be subdivided into non-overlapping bands (i.e., channels) that may be accessed collectively to create larger bandwidth channels. Each SDS shall contend for a band or set of bands that fully covers its spectral use.
6. Rules shall support varying durations of transmissions.
7. Rules shall support reservation of bandwidth for quality of service.
8. Rules shall provide means to preempt reservations.
9. Rules shall allow certain SDSs to have precedence over others based on their user(s) or use and shall provide a means to instruct an SDS regarding the precedence it has.
10. Devices shall only transmit on highways where it participates in signaling. In most cases, this limits devices to transmitting on one highway at a time and prevents combining spectrum lanes of multiple highways to support a single transmission.
11. The highway system shall exploit the use of directional antennas for greater spatial reuse when they are available on an SDS.
12. There shall be a standard way to define the boundaries of the spectrum highway and its lanes in terms of location, time, and frequency.
13. Signaling channels for a highway shall allow SDSs to constrain their participation to gaining access to a subset of the channels of the highway.
14. The method for highway definition shall allow definition of lanes with different bandwidth and duration on that highway.
15. SDSs shall receive definitions of the "roads and lanes" and shall only operate within their boundaries. If this is not feasible, they shall not operate on the highway.
16. SDSs will not have to be configured to participate on a particular highway (e.g., assigned a time slot on an orderwire). The highway definition shall contain sufficient information for devices to participate.

17. SDSs shall have a means to understand their location and the time of day.
18. SDSs shall understand the spectral boundaries of their transmissions in terms of a reference power spectral flux density per hertz at one meter from their antenna for far field propagation using a log-distance pathloss model.
19. SDSs that gain access to a transmission opportunity (i.e., a channel and timeslot) shall be able to use it for either transmitting or receiving with no effect on their retaining access.
20. SDSs may be paired with separate devices to perform any signaling required for access and shall meet any end-to-end delays with that device necessary to correctly use the highway

1.4 Highway Design Overview

The dynamic spectrum access technique described in this report is designed using highways as a metaphor. Following that metaphor, the objective is to set aside spectrum in geographical locations to support a highway and then to subdivide that spectrum into lanes and time periods called transmission opportunities. Both the lanes and the transmission opportunities can be combined to create wider bandwidth lanes and longer duration transmission opportunities. Transmission opportunities are defined in a hierarchy wherein multiple timeslots make a frame and multiple frames make an epoch. Transmission opportunities on lanes, i.e. timeslots and frames, can be reserved on a periodic basis to support streaming or other types of periodic access requirements.

SDSs of various types are authorized to operate on the highways and their component devices gain access to transmission opportunities on spectrum lanes by following the rules of the highway and using signaling to differentiate which device has precedence and to resolve access in congested conditions. There is no restriction on how SDSs use the spectrum as long as that use falls within the bounds of the model that defines the lane they use. Uses can be entirely passive. There is no requirement for an SDS with highway access to actually transmit: it may passively listen.

Spectrum highways combine two technologies. The first technology is defined in the IEEE 1900.5.2a Standard Method for Modeling Spectrum Consumption. All spatial and spectral aspects of highways are specified using spectrum consumption models (SCMs) built in accordance with the IEEE 1900.5.2a standard. Each spectrum highway is defined using a set of SCMs called an authorization set wherein each SCM provides the definition of a lane of the highway. Internal to an SCM is a construct called the Protocol or Policy construct. This construct contains the definition of the signaling design used for the lane.

The second technology is the media access control (MAC) paradigm, Synchronous Collision Resolution (SCR). Multiple MAC paradigms have been used for wireless access but only the SCR paradigm is capable of simultaneously differentiating access based on the precedence of users or uses of the spectrum, differentiating access based on the QoS required in the access, resolving contentions among multiple devices that have equivalent precedence, and orchestrating the spatial reuse of the lanes. Thus, SCR is used as the access mechanism of spectrum highways.

Spectrum highways use a signaling scheme that first decides the lanes to use, next determines which of the contending devices has precedence, then determines which devices have a use with

the highest QoS requirement, and finally resolves which devices that are still contending get access to the active lanes. The highway SCR scheme is designed as a generic approach that highway designers can configure using a set of parameters they choose. Configuration alternatives trade access features such as precedence in lane selection, the levels of user/user precedence, the levels of QoS differentiation, and contention resolution effectiveness for the time it requires to execute the corresponding signaling.

Rules govern the use of signaling and the precedence that SDSs and individual devices of the SDSs are allowed to use and how they must respond to the signaling of others and their observation of the use of the highway. SDSs and devices are assigned precedence levels for their operation on highways. Just as an ambulance would have the right-of-way on a highway in the case of an emergency, certain conditions can change the precedence an SDS has in gaining access. Further, SDSs are given the authorization sets of the highways on which they are allowed to operate.

There are four approaches to signaling. Two of the approaches use in-band signaling: per slot and consolidated. The other two use out-of-band signaling: serial and concurrent.

In-band signaling uses the spectrum that makes up the lanes of the highway. As a result, signaling consumes the resources of the lanes for the time it takes to execute the signaling. The fraction of time consumed by the signaling is a function of the duration of the transmission opportunities chosen for the design and the amount of signaling used in the contention design. This can be a large proportion of the total time. The advantage of in-band signaling is that the signaling will have similar propagation and directionality properties as the SDS spectrum use.

The difference between per-slot and consolidated signaling is that in per-slot signaling, as the name implies, a contention occurs prior to each slot and in consolidated signaling all signaling for the timeslots of a frame is executed consecutively prior to the start of the timeslots of a frame. The advantage of per-slot signaling is the turnaround from successful contention to the spectrum use. The advantage of consolidated signaling is that it gives devices more time to transition from signaling to SDS use and allows the continuous use of consecutive timeslots that were won in contention.

Out-of-band signaling uses spectrum separate from the lanes for signaling so that all the resources of the highway can be used efficiently. The disadvantage is that additional spectrum is required and that the spectrum used for signaling may have different propagation properties.

In highways that use serial signaling, the contention for each of the lanes of the highway occurs sequentially and uses a single signaling channel. The smallest transmission opportunity cannot be any shorter than the time required to serially execute the set of contentions for all lanes of the highway. In concurrent signaling, the contention for lanes occur on separate channels concurrently. The duration of timeslots can be reduced to the duration of the longest lane contention. The advantage of serial signaling is that the contention can be complemented with a means to identify the SDS that wins. This enables rapid changing of lanes from timeslot to timeslot and eliminates the need for devices to provide their own methods to coordinate their changing of lanes. The advantage of concurrent signaling is that the timeslots can be made comparatively shorter than those of the other signaling approaches.

A technology that can simplify the creation and exploitation of spectrum highways is a separate device designed to execute the signaling called a Spectrum Highway Access Device (SHAD). SDSs would be designed to operate with SHADs. They would provide the SHAD with information on their access needs and the SHAD would perform the signaling work to gain access. SHADs would be designed to efficiently implement the signaling for highway access. Using SHADs would make SDS development easier and eliminate the need to build the signaling capability into the SDSs. Assuming highways become a common method to dynamically share spectrum based on precedence and demand, the demand to use spectrum highways could create a scale of production that could make the SHADs the most economical signaling means to access highways.

Given this proposal, the next steps of development would be to incrementally experiment with these concepts in simulation, in a small test bed, in a large test bed, and ultimately a live experiment. Supporting activities for highways would include designing physical layer of signals for different contexts that trade characteristics of spectral and temporal efficiency with resilience, creating standards for different signaling approaches and the communications between SDSs and SHADs, developing SHADs and SDSs that enable the highway, and building the infrastructure to govern spectrum highways.

Spectrum highways provide an exciting new approach to spectrum access giving autonomy to devices to choose spectrum in which to operate while ensuring that spectrum capacity does not collapse as a result of multiple devices choosing the same spectrum. Highways provide means for those devices to arbitrate who and what uses should be supported. These features are valuable in pure commercial systems, especially at ranges that cannot be supported by Industrial, Scientific, and Medical (ISM) bands; in bands where government and commercial users share the spectrum; and in military contexts where SDSs autonomously maneuver in spectrum.

2 Defining RF Spectrum Highways

Like a highway for vehicular traffic, an RF spectrum highway system will have their equivalent of roads and lanes. A road defines the full set of alternative spectrum in which the rules apply and the lanes define how that set is subdivided and can be shared by different users. The following sections describe what constitutes these boundaries and how a highway is defined in space, time, and spectrum.

2.1 Spectrum Consumption

The highway metaphor must be suspended to understand how spectrum is occupied and consumed. Consumption is not a measure of presence but of activity. Spectrum highways do not support movement but temporary consumption. The reach of consumption is contingent on the environment, the ability of transmitters to project signal power and the sensitivity of receivers. Total consumption is also a function of signal bandwidth and the characteristics of the receiver and its sensitivity in adjacent bands. Finally, consumption is a function of time. It is rare that a transmitter or a receiver will have to operate continuously.

Further, the boundaries of spectrum occupancy by SDSs are quite irregular and variable. They can be very different among the variety of SDSs that collaborate to share the spectrum. SDSs can be wideband or narrowband. SDSs can use different power levels and different antenna gains so project their transmissions to different ranges or have receivers with different sensitivity and so require different clearance from other users. SDSs can have directional antennas and thus constrain their consumption to a direction.

SDS consumption of spectrum is a function of location of the SDS and the geographical boundary of the consumption is therefore also very irregular and unpredictable.

An SDS's consumption of spectrum starts and ends nearly instantaneously.

Spectrum consumption of SDSs is relative. The ability of one SDS to coexist with another is a function of each SDS's individual traits.

An SDS's ability to operate on a spectrum highway is a function of the capabilities of the SDS. Every SDS will have some finite amount of spectrum bandwidth over which it can operate which may or may not be within the spectrum highway. Further, SDSs may not be able to tune to and to operate within the boundaries of the lanes specified for a highway.

2.2 Spectrum Availability

It is impractical to expect any particular swath of spectrum to be universally available for system operation. Host nation rules and competing systems with higher precedence that must be protected will encroach on these boundaries. Thus, the availability of the spectrum in a given highway will vary by location and by time and the spectrum highway system will consist of geographical and temporal patches of spectrum authorizations. SDSs that are stationary may only see temporal variation in the highway definition but those that are mobile will likely have to adapt to temporal and spatial differences.

Spectrum availability for communication, tele-control, and radio location SDSs may be different than the spectrum made available for the electronic warfare functions of electronic support (ES) and electronic attack (EA). The rules of the highway will prevent friendly disruption of ES and prevent EA from harmfully interfering with other highway users. However, EA and ES devices would probably not be limited to operating on highways.

2.3 Spectrum Highway Boundaries

Unlike vehicular roads and highways, spectrum highways are not infrastructure that are permanently placed; they are authorizations, and their persistence is a function of policy. Highways could persist for long periods because spectrum is set aside for their long-term operation or they could be defined for a brief period because of an operational need or simply in response to the opportunity to use a typically occupied band of spectrum that becomes available. Further, a comprehensive redesign of the spectrum highway system can be done as quickly as it takes to load the new designs onto the SDSs that use the highways.

Spectrum highways are defined in five dimensions: the three dimensions of geographical space, and the dimensions of frequency, and time. Frequency and time are further subdivided into spectrum lanes and transmission opportunities, respectively. At any given instant, an SDS may use part or all of the geographical space of a spectrum highway but only in the discrete parts that constitute its allowable spectrum lanes and its transmission opportunities. Highways may be designed with options that merge lanes to support different bandwidth systems and to merge transmission opportunities to support variable duration temporal usage. Spatial usage of the highway is a function of the location of the device and where feasible, it is desirable that spectrum be spatially reused on the same highway.

2.3.1 Highway Location and Spectrum

Spectrum highways are inherently geographical.² In any geographical location multiple spectrum bands may be available for highway use. There is no requirement for the spectrum band of a spectrum highway to be contiguous but the same spectrum must be defined for use by the highway at all locations of the highway. Similarly it is possible for multiple spectrum highways to spatially overlap. Highways are distinguished from each other in that they are organized to enable SDSs on that highway to use all or part of the spectral bands of the highway. SDS devices are expected to operate on one highway or another when multiple highways apply to the same location. A device shall not simultaneously use a combination of bands that is only possible by accessing spectrum from multiple spectrum highways. However, a device may move among highways. It also is possible that highways may share spectrum and so a device could operate in spectrum that is part of multiple highways. In this case, the lanes, transmission opportunities, and rules for access for overlapping spectrum bands must be the same. In most cases this is impractical

² Spectrum management is tagged to locations on earth but our meaning here is more, it is spatial and that space can be three dimensional. Later the terms space and geographical location are used as synonyms.

Figure 2-1 illustrates boundaries in space and in frequency for two highways. In this scenario, two highways overlap in space and in frequency. The frequencies available on each highway are available everywhere on that highway. There is no requirement for the spectrum to be contiguous. The frequencies shared by two highways that overlap are required to be channelized into identical lanes and to be managed by identical rules.

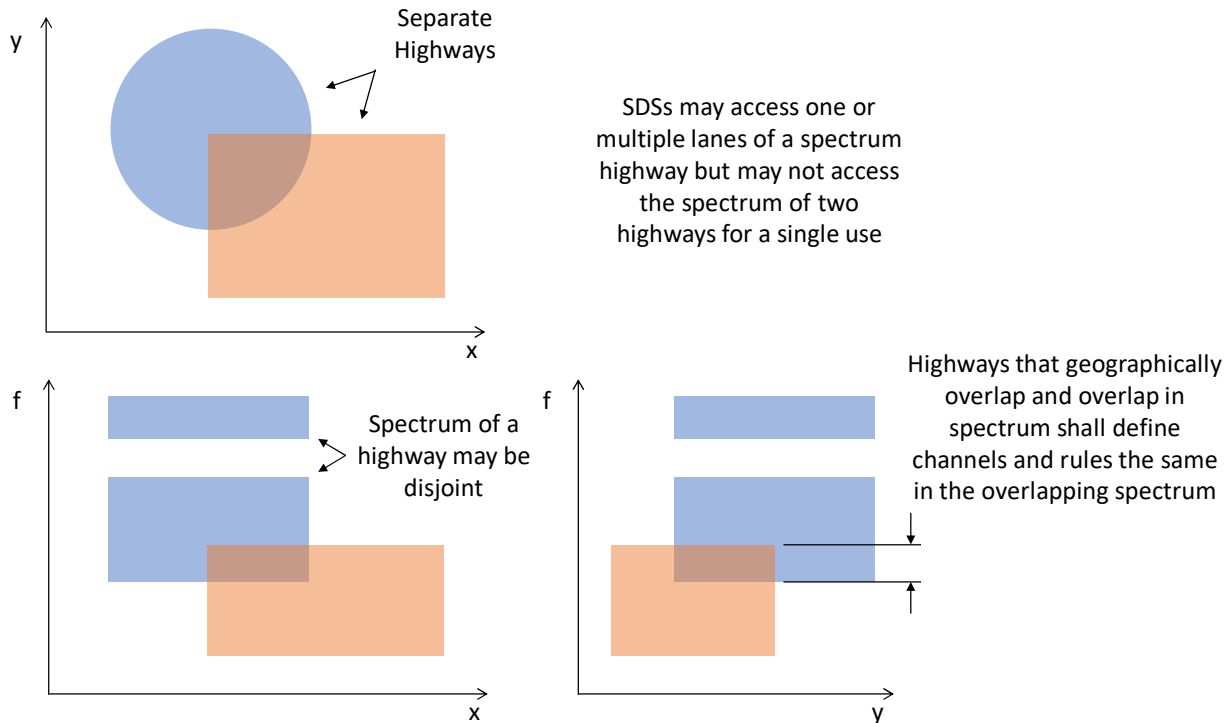


Figure 2-1. Geographical Spectrum Highways.

2.3.2 Spectral Lanes

The lanes of a spectrum highway are the quanta of spectrum that SDSs may use. They are defined with a maximum power spectral flux density³ at each frequency and in time. The bandwidth of each lane is referred to as a channel and the discrete time periods or time slots are referred to as transmission opportunities. SDSs may access and use more than one of these lanes at the same time. An SDS that gains access to a transmission opportunity may organize the opportunity between transmissions and receptions in whatever way is most useful to itself. The emitted radiated power of transmissions must be below the defined power spectral flux density.

³ Power spectral flux density is used because it includes the effects of antenna gain. Highways regulate the strength of emissions.

Figure 2-2 illustrates the ideal division of a highway band into lanes and further illustrates that lanes may be merged. Figure 2-3 illustrates a more practical definition of lanes where skirts that require a decrease of power spectral flux density at the edges in order to protect receivers operating in the adjacent lanes. In this case, merged lanes operate at the higher power spectral flux density across the skirts of the merged lanes. Figure 2-4 further illustrates that if the lanes do not spectrally touch or intersect then they cannot be merged; further, when the adjacent lanes are defined with a different maximum spectral flux density then the lower maximum power spectral flux density of the adjacent lanes is the power spectral flux density level across the skirts if the lanes are merged.⁴

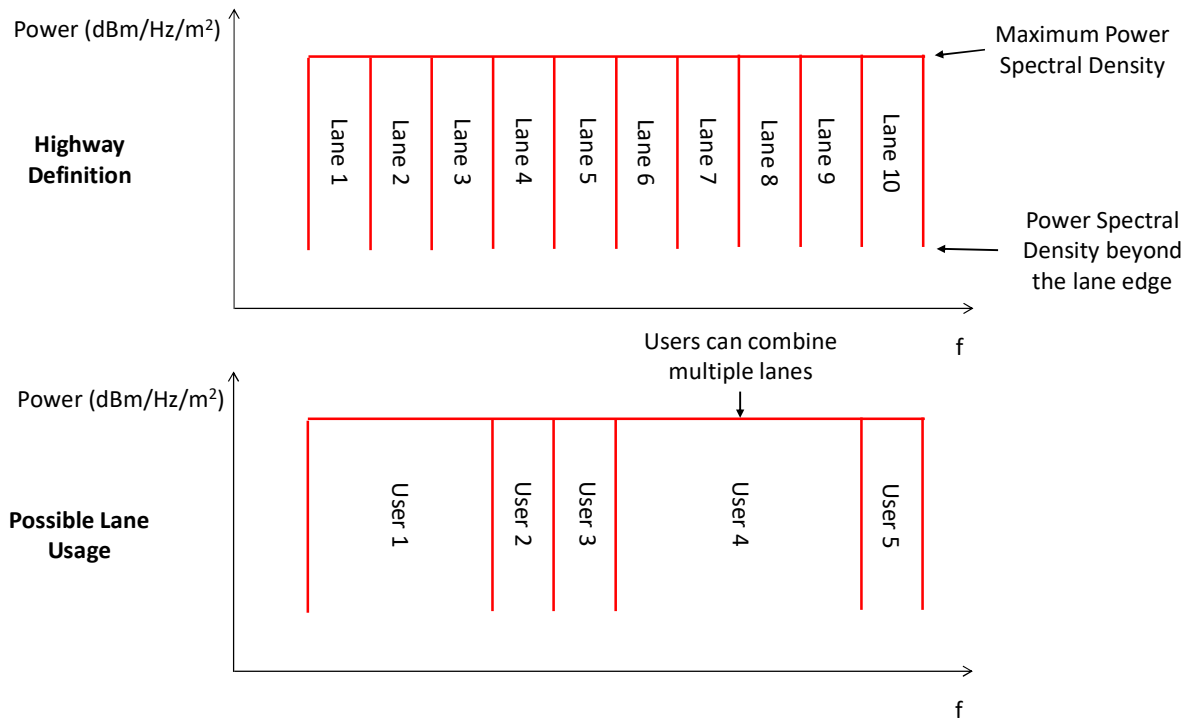


Figure 2-2. Spectrum Lanes

⁴ The method for combining lanes depends on the signaling channel that is used. In-band-signaling requires merged lanes to be explicitly defined and so it may be difficult to share spectrum across highways. Serial out-of-band signaling will allow SDSs to attempt to merge lanes ad hoc with their signaling. Signaling channels are described in Section 4.2.2





Figure 2-3. Spectrum Lanes with Skirts





Figure 2-4. Criteria for Merging Lanes

2.3.3 Highway Schedule and Transmission Opportunities

The temporal aspect of spectrum highways is described in two ways. The first is by the schedule of highway existence. It could be permanent, periodic (e.g., from 8:00 PM to 4:00 AM every day), or episodic (i.e., identified for some duration based on opportunity). The second temporal description is to subdivide highway existence into discrete transmission opportunities. The transmission opportunities of lanes are then divided into a hierarchy of time periods. Using a hierarchy has two benefits; it enables users to contend for different durations of access, and it supports contentions for periodic access. Figure 2-5 illustrates a division of transmission opportunities into three levels of hierarchy. In this example, the three levels are defined as timeslots, frames, and epochs where there are 10 timeslots per frame and 4 frames per epoch. Access rules can allow users to contend for timeslots, whole frames or whole epochs and thus allow users to obtain access for different durations. Access rules can also allow users to contend for and gain access to the same timeslot each sequential frame or the same frame for each sequential epoch for a periodic reservation. It is undesirable for devices to contend once for an arbitrary number of timeslots, because such a mechanism could be used to counter the goal of sharing. However, as illustrated in Figure 2-6, transmission opportunities at a particular level may be defined for different durations. The figure shows two timeslots per frame that are longer than the others. Additionally, devices may contend for multiple consecutive timeslots.⁵




Figure 2-5. Example of Three Levels of Transmission Opportunities




Figure 2-6. Example of Different Duration Transmission Opportunities at a Level

Highway definitions must use the same transmission opportunity definitions in order for channels to be merged. When it is useful to have lanes with different transmission opportunity definitions, those lanes should be defined as part of separate highways. Additionally, transmission opportunities are selected in conjunction with selecting and designing the access mechanism used for contention. Section 3 describes channel access mechanisms.

⁵ Three levels of hierarchy can support most access options and provide a variety of access durations and is the maximum number of levels in the proposed method for defining highways in Chapter 5.

2.4 Defining Spectrum Highway Boundaries

The description above shows that defining the spectrum highway boundaries requires defining the location of the highway, the allowed power spectral flux density as a function of frequency of each of the lanes, and the temporal division of the highway into transmission opportunities, and identifying a mechanism used to arbitrate access.

IEEE Standard 1900.5.2 is perfectly suited for describing the boundaries of spectrum highways. The Authorization Set of the standard allows the use of multiple transmitter models to define the spectrum that a user is authorized to use. Transmitter modeling can define the power spectral flux density by frequency, the spatial extent within which a transmitter may operate, and the protocols it must use to access the spectrum. Thus, an authorization set can be used to define a highway in which each transmitter model defines a specific lane or a combination of lanes that RF devices may use.

The eleven constructs of spectrum consumption modeling are [1]:

1. Reference Power: Reference power value for the spectrum mask, underlay mask, and power map constructs. It is the only construct with a true power term. Other constructs use values in dB relative to this value.
2. Spectrum mask: Data structure that defines the relative spectral power density of emissions by frequency.
3. Underlay mask: Data structure that defines the relative spectral power density of allowed interference by frequency.
4. Power map: Data structure that defines a relative power flux density per solid angle.
5. Propagation map: Data structure that defines a path loss model per solid angle.
6. Intermodulation mask: Data structure that defines the propensity of co-located signals to generate intermodulation products in a transmitter or receiver.
7. Platform Name: Name or list of names of platforms that are attributed to a particular location/platform (i.e., naval ship, airplane, etc.). The names are useful in identifying when multiple systems are co-located and could suffer intermodulation (IM) and out-of-band interference effects.
8. Schedule: Construct that specifies the time in which the model applies (start time, end time). Periodic activity (on/off) can also be defined.
9. Location: Location where an RF device may be used. Several type of locations are supported: a point, a surface area, a volume, or a path.
10. Minimum power spectral flux density: Power spectral flux density that when used as part of a transmitter model implies the geographical extent within which receivers in the system are protected.
11. Policy or protocol: Named protocol or policy with parameters that define behaviors supported by a device or systems that allow different systems to be co-located and to co-exist in the same spectrum.

These constructs can be used to model both transmitters and receivers. Figure 2-7 illustrates the feasible models. Transmitter and receiver models can be combined into models of systems and then transmitter, receiver, and system models can be combined into sets. Models and sets can have five purposes: to identify the configuration options of systems, to request spectrum, to identify spectrum consumption, to authorize spectrum use where any use of spectrum within the

individual models or the models of the set is authorized, or to provide constraints to spectrum use (i.e. uses are not permitted to interfere with the receiver models in the set and must accept the interference defined by the transmitter models in the set). A spectrum highway, as illustrated in Figure 2-1 and Figure 2-2, would be defined as an authorization set of transmitter models. Each transmitter model would define a lane or a combination of lanes. It is feasible to complement the authorization set defining a highway with a constraint set. In this case, SDSs are allowed to operate on the highway so long as they do not interfere with any of the models of the constraint set. In this report, we do not discuss this approach further assuming that highway designers will choose boundaries in the authorization set that would protect any legacy users that need protecting.

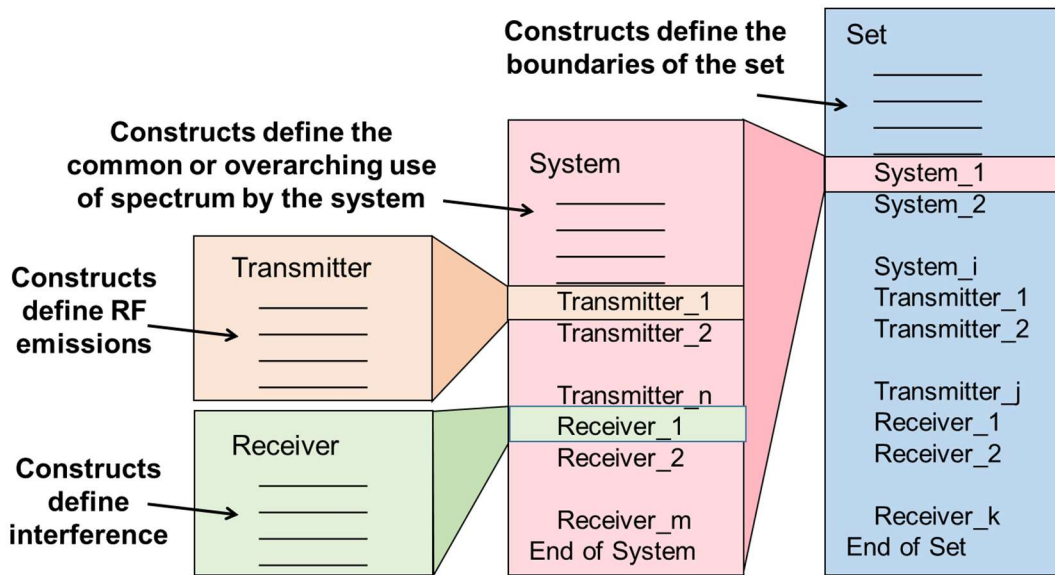


Figure 2-7. SCM Model Hierarchy

Several of the constructs listed above are not used in the definition of a spectrum lane. Spectrum lanes are defined using: reference power, spectrum mask, power map, schedule, location, and policy or protocol constructs. A propagation map, though not needed to define a lane, would also be provided to be consistent with the modeling specification. The combination of the reference power, spectrum mask, and power map defines the limits to the power spectral flux density that users can transmit. In addition, the power map allows this specification to be directional. The location defines where the transmissions may originate. The schedule defines when the spectrum highway is active. The transmission opportunities and the rules for arbitrating access for each of the lanes would be defined in the protocol or policy construct.

Section 2.4.1 presents an example of defining the lanes of a highway. Section 2.4.2 describes the elements of a policy or protocol definition.

The remaining section of this report investigate and recommend protocol methods for differentiating precedence of SDSs and spectrum uses and for arbitrating access. They contain guidance on how to define these factors using the protocol or policy construct, and provide a design approach for highway access defining both the signaling and the rules for access.

2.4.1 Highway Exemplar

A highway is assigned to a location. Using the spectrum consumption modeling approach, this is accomplished by defining a polygon or a circle on a map or using a polyhedron or cylinder when the space must be allocated an aerial device. Figure 2-8 illustrates a polygon location in Northern Virginia. The highway specification would apply to devices that operate within this area. Clearly, transmissions may extend beyond the area. The creation of highways adjacent to this one in the same spectrum must consider the ramifications and mitigate interference. A device on the highway must verify it is in the assigned space and then follow the rules of the highway.

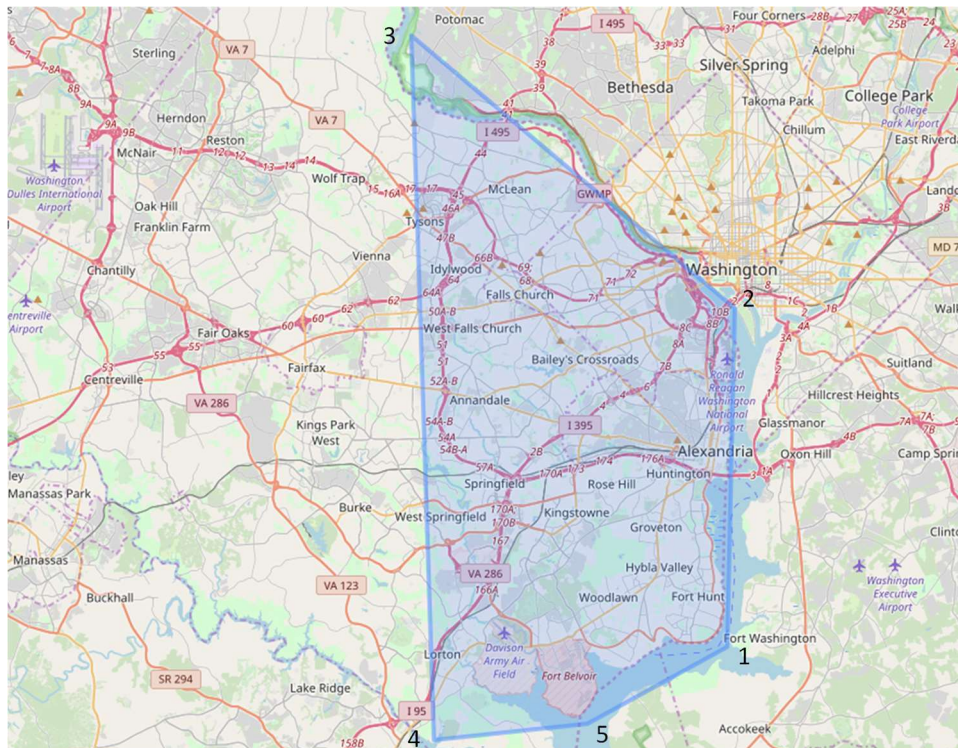


Figure 2-8. Exemplar of a Highway Location

Highways are subdivided into spectrum lanes. Figure 2-9 illustrates the spectrum masks for a five-lane highway. The scaling for these lanes captures the allowed power spectral flux density of a transmission one meter from the antenna as would be considered in a log-distance pathloss model. (This may not be measurable at one meter because it is still near-field but it can be computed based on the transmission power and antenna gain of the device.)




Figure 2-9. Exemplar Highway Lanes

Table 2-1 through Table 2-14 describe the highway using the constructs of spectrum consumption modeling. The description in the tables is specific to one of the five illustrated lanes. The lane is described using a transmitter model. The other four lanes illustrated in Figure 2-9 would be defined in almost the same way. The difference would be the RefFrequency shown in Table 2-10.

Table 2-1. Spectrum highway description using a SCM Set

Primary Data Element	Sub - elements	Value	Notes
SCMSet			
	SCMSetID	HW546	Must be a unique string value. Conventions can be developed to distinguish highways
	Name	NOVA Spectrum Highway 5	Optional
	Purpose	AUTHORIZATION	Gives permission for spectrum use
	Boundary	See Table 2-2	Optional element that identifies a boundary that encompasses the highway in frequency, time, and space. Used for quick determination of applicability.
	TxModel	See Table 2-8	Defines one lane
	⋮ TxModel	Not provided	Additional models for every lane and every combination of lanes

The Boundary data element and sub-elements assist automated systems to quickly determine if an SCM Set is relevant or not. When used as part of a set specifying a spectrum highway, it enables a system to determine if the set is irrelevant because the system cannot operate in the specified band(s), is not located within the area or volume of the highway, or must operate

outside the time period the highway is active. Table 2-2 through Table 2-7 define the boundaries.

Table 2-2. Set Boundary for an SCM Set

Primary Data Element	Sub - elements	Value	Notes
Boundary			
	BandList	See Table 2-3	Optional but when used provides one or more bands that fully cover all bands used by the highways for the lanes, e.g. the transmitter models
	SCMSchedule	See Table 2-4	Optional but when used provides a schedule within which the highway is operational.
	SCMLocation	See Table 2-5	Optional but when used provides a location that fully covers the area of the spectrum highway.

All lanes of the highway lie within the bands specified by the band list in Table 2-3. The band list may cover frequency ranges that have no lanes. A system that can operate within the bands of the band list will need to check the specific lane definitions to make a final determination that it can operate on this spectrum highway.

Table 2-3. Band List That Provides the Spectral Boundary of the Highway

Primary Data Element	Sub - elements	Value	Notes
BandList			
	Band	See below	Just one band in this example but there could be several
Band			
	StartFrequency	400 MHz	
	EndFrequency	405 MHz	

This highway schedule boundary subsumes the schedules of the lanes of the highway. The SCM Set data elements do allow the definition of a highway that has lanes with different schedules. In this case the boundary schedule (Table 2-4) would start with the start of the earliest starting lane and end with the end time of the latest end time of a lane. Final determination of whether a lane’s schedule is suitable requires the SDS to inspect the schedule definitions of the individual lanes

Table 2-4. Highway Schedule Specifying the Temporal Start and End of the Highway

Primary Data Element	Sub - elements	Value	Notes
SCMSchedule			
	StartTime	2018,07,25,00,00,0,-05,00	This time shows the highway to start \geq at the very start of 25 July 2018, local time, which is 5 hours less than UTC time.
	EndTime	2018,08,01,00,00,0,-05,00	This time shows the highway to end \leq the very start of 1 August 2018, local time, which is 5 hours less than UTC time.

The SCM standard provides several different ways to define a location. This example uses a polygon surface. In the data schema hierarchy the existence of the PolygonSurface within a location indicates just that, the location is a polygon surface. Its sub-elements define its two parts; the polygon and the antenna height. The boundary aspect of the location means that all locations of the highway fall within the geographical boundary of this location (Table 2-5). SCM modeling allows the definition of lanes that have different locations. All those locations are required to be within the boundary location, so all the lanes of this highway are within this geospatial boundary. Systems and users must check the definition of the specific highway lanes to determine if their location would permit their use of the spectrum lane.

Table 2-5. Boundary of the Highway Location

Primary Data Element	Sub - elements	Value	Notes
SCMLocation			
	PolygonSurface	See Below	
PolygonSurface			
	SCMPolygon	See Table 2-6	
	AntennaHeight	See Table 2-7	Required by the SCM standard. In this use it could mean the maximum height of the antenna.

Polygons are defined by sides that extend between points. A Side data element provides the starting point of a side and the ending point of the side is the point defined in the next Side data element. The last side of the polygon extends from the point of the last Side data element to the point of the first Side data element. The points are defined using WGS-84 coordinates. The polygon surface defined by Table 2-6 is illustrated in Figure 2-8.

Table 2-6. Details of the Polygon Surface Used for the Location of the Spectrum Highway

Primary Data Element	Sub - elements	Value	Notes
SCMPolygon			
	Side	See Point 1 sub-elements below	Each Side is defined by a point. The order of the Sides defines the polygon with the first side starting at the point
	Side	See Point 2 sub-elements below	
	Side	See Point 3 sub-elements below	

Primary Data Element	Sub - elements	Value	Notes
	Side	See Point 4 sub-elements below	defined in the Side element and extending to the point defined in the next Side element. The last side extends from the point of the last Side element to the point of the first Side element.
	Side	See Point 5 sub-elements below	
1			
Point			
	Longitude	-77.03955531	
	Latitude	38.7080177	
	Altitude	0 Meters	
2			
Point			
	Longitude	-77.03474879	
	Latitude	38.87553224	
	Altitude	7 Meters	
3			
Point			
	Longitude	-77.2400558	
	Latitude	39.01011396	
	Altitude	85 Meters	
4			
Point			
	Longitude	-77.22563624	
	Latitude	38.66192242	
	Altitude	14 Meters	
5			
Point			
	Longitude	-77.12813258	
	Latitude	38.66996443	
	Altitude	0 Meters	

The antenna height of a surface location defines the maximum antenna height allowed on surface (Table 2-7). It can be referenced to the ground level or to the average terrain height of the location.

Table 2-7. Maximum Antenna Height of Devices within the Spectrum Location

Primary Data Element	Sub -elements	Value	Notes
AntennaHeight			
	Height	30 Meters	
	Reference	AGL	References antenna height to the local ground level. AGL stands for above ground level. The alternative is height above average terrain (HAAT).

Transmitter models define lanes. The transmitter model of Table 2-8 defines Lane 1. In this highway example, the five lanes illustrated in Figure 2-9 would be defined by nearly identical models. The only differences between the lanes would be the value for RefFrequency of the SpectrumMask in Table 2-10. In this example, the location and schedule are identical to the boundary location and schedule. Table 2-8 refers to the boundary data tables but in a real description the same data would be repeated for each transmitter model.

Table 2-8. SCM Transmitter Model use for Lane 1 of the Spectrum Highway

Primary Data Element	Sub -elements	Value	Notes
TxModel			
	Purpose	AUTHORIZATION	
	ReferencePower	See Table 2-9	
	SpectrumMask	See Table 2-10	
	SCMPowerMap	See Table 2-11	
	SCMPropagationMap	See Table 2-13	Required model element but not used in regulating a spectrum highway.
	SCMLocation	See Table 2-5	Uses the same location as in the boundary but in this case the location is exact rather than subsuming.
	SCMSchedule	See Table 2-4	Uses the same schedule as in the boundary but in this case the schedule is exact rather than subsuming.
	SCMPolicyOrProtocol	See Section 2.4.2	Used to identify rule sets that support spectrum sharing.

The reference power is the only term in a transmitter model that has the units of Watts. It is the reference for the power referenced terms; the spectrum mask and the power map. The constraint that a transmitter model defines is the allowed power spectral flux density by frequency. This is conveyed by the combination of three constructs, the ReferencePower, the SpectrumMask, and the SCMPowerMap. A modeler can shift power levels among these constructs to obtain the same constraint. As a result, devices using the highway do not require the exact performance indicated by the constructs. For example, the directional gain of an antenna may be used to reduce the transmit power. The approach for this example, however, is to try to follow the physics of the components so that the derivation of the values can be described. The ReferencePower in Table 2-9 corresponds to a transmit power of 15 Watts.

Table 2-9. Reference Power for Lane 1

Primary Data Element	Sub - elements	Value	Notes
ReferencePower			
	Value	11.762 dBW	Corresponds to 15 Watts

The power scaling of the spectrum mask accounts for the spreading of energy across the breadth of the mask. The power units of the mask are dB/RBW, the resolution bandwidth. Integrating the power across the mask by frequency provides the power of transmission. When the power of transmission is defined by the ReferencePower then the total power that results by integrating across the mask is 0 dB. The scaling of the mask in Table 2-10 yields a mask with 0 dB total power.

Table 2-10. Spectrum Mask of Lane 1

Primary Data Element	Sub - elements	Value	Notes
SpectrumMask			
	ResolutionBW	10 ⁻⁶ MHz	Corresponds to 1 Hz
	SCMMask	See Below	
1			
SCMMask			
	RefFrequency	400.5 MHz	
	ControlPoint	See ControlPoint 1 Below	Mask have multiple inflection points and are assumed to be presented in the order of lowest frequency to highest frequency.
	ControlPoint	See ControlPoint 2 Below	
	ControlPoint	See ControlPoint 3 Below	
	ControlPoint	See ControlPoint 4 Below	
	ControlPoint	See ControlPoint 5 Below	
	ControlPoint	See ControlPoint 6 Below	
1			
ControlPoint			
	Frequency	-0.55 MHz	Relative to reference frequency
	RelativePower	-89.15 dB	Per the ResolutionBW
2			
ControlPoint			
	Frequency	-0.45 MHz	Relative to reference frequency
	RelativePower	-79.15 dB	Per the ResolutionBW
3			
ControlPoint			
	Frequency	-0.40 MHz	Relative to reference frequency
	RelativePower	-59.15 dB	Per the ResolutionBW
4			
ControlPoint			
	Frequency	0.40 MHz	Relative to reference frequency
	RelativePower	-59.15 dB	Per the ResolutionBW
5			
ControlPoint			
	Frequency	-0.45 MHz	Relative to reference frequency
	RelativePower	-79.15 dB	Per the ResolutionBW
6			
ControlPoint			
	Frequency	-0.55 MHz	Relative to reference frequency
	RelativePower	-89.15 dB	Per the ResolutionBW

The power map of an SCM attempts to capture two phenomena, the directionality of antennas and the power loss that results from the transition from the antenna to the atmosphere at one meter. In the case of directionality this highway does not restrict the pointing of antennas and so the pattern described is omnidirectional. The Orientation data element in Table 2-11 provides the reference for the direction components of the map. The surface orientation indicates the map is referenced to its location on the Earth; refer to The IEEE 1900.5.2 standard for more details. Table 2-12 is the actual power map.

Table 2-11. Power Map of Lane 1

Primary Data Element	Sub - elements	Value	Notes
SCMPowerMap			
	Orientation	See Below	
	GainMap	See Table 2-12	
Orientation			
	Surface	True	Indicates map directions are referenced to the location on earth.

The details of the map data structure and how it is used to assign values to directions are described in the IEEE 1900.5.2 standard. The version that appears in Table 2-12 assigns a single gain in all directions. The power scaling in the Gain data elements of the power map accounts for three factors: insertion losses at the transmitter, spatial spreading, and antenna gain. In capturing the performance of a device the Gain value may also capture insertion losses but in this definition it is not assumed to be a factor. Thus the gain value is computed as follows

$$PowerMapGain = SpatialSpreadingLoss + AntennaGain$$

$$PowerMapGain = 10 \cdot \log\left(\frac{1}{4\pi}\right) + 3$$

$$PowerMapGain = -7.992 \text{ dB/m}^2$$

In this example a 3 dBi gain is assumed for the antennas. This does not mean that the highway requires the devices to use antennas with this gain. The goal is to define through the multiple constructs of the transmitter model the allowed power spectral flux density that a device may emit.

Table 2-12. Omnidirectional Gain Map Used by Lane 1

Primary Data Element	Sub - elements	Value	Notes
GainMap			
	GainMapValue		These five values define an omnidirectional gain pattern. See the IEEE 1900.5.2 standard for guidance to represent directional patterns. In this use antennas used may be directional, but this definition assumes orientation can be in any direction and limits Effective Isotropic Radiated Power (EIRP) in all directions.
	GainMapValue		
	GainMapValue		
	GainMapValue		
	GainMapValue		
1			
GainMapValue			
	Elevation	-90°	
2			
GainMapValue			
	Azimuth	0°	
3			
GainMapValue			
	Gain	-7.992 dB/m ²	This value accounts for the one meter spreading losses, and a 3 dBi antenna gain.
4			
GainMapValue			
	Azimuth	360°	
5			
GainMapValue			
	Elevation	90°	

The propagation map noted in Table 2-13 and Table 2-14 is required by the standard but does not affect the decision whether a device can use a the spectrum highway. It is used in evaluating the effects of the spectrum highway on neighboring users of the same spectrum so that adjacent uses can be managed. The data structure is very similar to that of a power map. It does not have an orientation data element since it is always surface oriented. The directional values of a propagation map are the parameters of a pathloss model. SCM modeling provides two alternatives: a linear log-distance model and a piecewise linear log-distance model. Table 2-14 uses a linear log-distance model. The propagation model in this case is omnidirectional like the power map and is equivalent to free space propagation.

Table 2-13. Spectrum Propagation Map for Lane 1

Primary Data Element	Sub - elements	Value	Notes
SCMPropagationMap			
	PropMap	See Table 2-14	

Table 2-14. Propagation Map Values for Lane 1

Primary Data Element	Sub -elements	Value	Notes
PropMap			
	PropMapValue		These five values define an omnidirectional pathloss rate. This data element is required by the standard but in this application it merely provides a worst case pathloss value to use in determining compatibility with other co-frequency adjacent spectrum uses near the highway
	PropMapValue		
	PropMapValue		
	PropMapValue		
	PropMapValue		
1			
PropMapValue			
	Elevation	-90°	
2			
PropMapValue			
	Azimuth	0°	
3			
PropMapValue			
	PropagationModel	See Below	There are two versions of a PropMapValue. We use the linear version (linear on a log distance plot).
3			
PropagationModel			
	Linear	2	Pathloss exponent for free space path loss
4			
PropMapValue			
	Azimuth	360°	
5			
PropMapValue			
	Elevation	90°	

The last essential construct of a spectrum highway is the SCMPolicyOrProtocol data element. This construct identifies the behavior that is required to use the spectrum. The remainder of this report investigates how to create access behavior in a spectrum highway and thus the definition of this construct. The following section gives an overview of this construct’s purpose and an example.

2.4.2 Protocol or Policy Construct Overview

The protocol and policy construct element was created to provide a very flexible means to specify behavioral guidance for spectrum reuse. It does not describe the behavior it merely identifies the behavior by name. Systems and devices that receive guidance with this construct are expected to know the name in order to follow the guidance.

The basic data structure of a protocol or policy is a name followed by a set of parameters. The assumption is that the name implies the general behavior of the target system and the parameters fill in the details of timing, levels, structures, and counts. Each policy and protocol name would

have an expected number of parameters associated with it that modelers must provide for the policy or protocol to be complete. Section 0 describes a possible listen before talk (LBT) protocol that is frequently identified as a protocol for spectrum sharing.

Exemplar Listen Before Talk Protocol

LBT protocols use sensing to identify reuse opportunities. This class of policies is further defined by four parameters: sensing threshold, sensing period, abandonment time, and disuse time. A possible sensing policy could require a cognitive system to sense a channel for no signal above a particular power threshold over some period of time and once it determine the channel to be clear to allow use of its spectrum. The protocol would then require the cognitive system to continue sensing the channel periodically, and if another user is detected, require the cognitive system to abandon the channel for some minimum period. For example, if the spectrum specified in the model is sensed below a power threshold of -120 dBW for 5 minutes then the cognitive system may use the spectrum within the constraints of the model so long as it senses the channel every 1 millisecond and abandons the channel if the power threshold is violated during that sensing. It must abandon the channel for 5 minutes before trying to use the spectrum again. This policy would have a name, in this case "LBT," and then four parameters, power threshold, p_{th} , free period, t_f , sensing period, t_s , and abandonment time, t_a . Since the meaning of the parameters is associated with the policy name, the complete policy can be conveyed concisely as $\langle name, p_{th}, t_f, t_s, t_a \rangle$ or specifically as $\langle \text{LBT}, -120, 300, 0.001, 300 \rangle$ where it is understood that the power term has units of dBW and all timing parameters have units of seconds.

3 Channel Access Options

Various options exist for wireless channel access and many wireless system designers have particular preferences. This section presents a survey of the multiple access options so that readers who have a preference can understand the limitations of their chosen option, and the choices made in the proposed approach to design highway access. Readers who just want to understand the access mechanism used in this highway design proposal should go directly to Section 3.3.3.

3.1 Centralized versus Distributed Control

Channel access is broadly divided into two media access control (MAC) protocol approaches: centralized control and distributed access. In the first, a master node, often referred to as a base station manages the access of several additional nodes that are sometimes referred to as slaves. This is the approach of cellular telephony and WiFi. In these cases, the base station is also the access point into the larger network. In the case of cellular, protocols work collaboratively among the base stations to manage the transfer of slave nodes' association with one base station to another as they move. This centralized control allows efficient use of spectrum and also considers the individual needs of the nodes resulting in decisions that consider both quality of service and quality of experience needed by users. However, the centralized approach is not appropriate for spectrum highways because it demands the buildout of infrastructure and is predicated on the central node understanding user needs. The goal of a spectrum highway is to allow independent and disparate spectrum users to cooperate in obtaining access with no expectation of infrastructure or an entity that understands the spectrum needs of all potential users.

Therefore, the access mechanism of a spectrum highway will necessarily be a distributed approach. Individual systems can participate because they can individually make decisions and interact with their neighbors to share the spectrum just as each driver in a vehicle independently makes decisions based on the conditions of the local environment in cooperation with other drivers and under the restrictions of the driving rules governing how to share a road. The challenging aspect of distributed spectrum access is that each user has a different view of what is happening in the spectrum based on the user's location and activity, and this can have dramatic effects on the efficacy of the chosen access mechanism. The rest of this section describes the associated issues in distributed wireless access and reviews several mechanisms used in attempts to resolve these problems.

3.2 Wireless Limitations

The wireless medium has characteristics that make it difficult to arbitrate access among radios when there is no master. These characteristics limit the effectiveness of access mechanisms that work well in other media.

3.2.1 Singular Task

Radios can either transmit or receive on a channel; they cannot do both simultaneously.⁶ Thus a device that is using a channel cannot simultaneously receive a signal on that channel that concerns its use of the channel.

The type of wireless networks that could participate in this sharing approach would be half duplex. The same channel(s) is/are used for sending and receiving transmissions, since there is no clear uplink and downlink in a collection of peer radios that may transmit or receive from any of the others. Unlike some wireline networks, a radio cannot simultaneously transmit and receive on the same channel, and unlike Ethernet, a radio cannot detect a simultaneous transmission. If two radios transmit at the same time and on the same channel, they interfere with each other and neither would be aware of it. Confirmation of a successful transmission requires an acknowledgement from the destination. Access mechanisms must be designed to mitigate the occurrence of simultaneous transmission in the same geographical area that can result in interference. The mechanisms are also typically designed to provide acknowledgements so radios can confirm their transmissions were received.

3.2.2 Channel Symmetry

The range of radio transmissions and their ability to reach other transceivers depends on a number of factors, including the environment and the specific characteristics of the transmitters and receivers involved such as the power of transmission and the gain of their antennas. It is quite possible for channels to be asymmetric; that is, a transceiver can receive the transmission of another but its transmissions cannot reach that transceiver. Thus a device may not be able to directly notify another of either its use of a channel or the degradation that it experiences from the other radio's use of a channel.

Additionally, radio range may change over time, for example as a result of changes in antenna height and orientation.

3.2.3 Channel Sensing

Many wireless access protocols are based on channel sensing: the process of simply observing the use of the channel. The limited range of sensing and the effects of geographically distributed users result in conditions that can befuddle access protocols based on sensing. The following subsections describe those conditions.

3.2.3.1 Hidden Nodes

Radios cannot detect the activity of a distant transmitter. Access mechanisms designed to operate based on first sensing the environment will detect active transmitters. A hidden node is a transmitter outside the range of a transceiver's reception that is communicating to a receiver that is in range and would be interfered with if this transceiver made the decision to transmit. Not

⁶ Some work is being done to enable pairs of radios to simultaneously transmit and receive on a channel but this type of access is static and continuous and does not support statistical multiplexing of channels which is the goal here.

detecting activity is not an indication that a channel is free. Access mechanisms must provide means for contending transceivers to know the activity of neighboring receivers.

3.2.3.2 Exposed Nodes

An exposed node is a node that sees a number of peers in a scenario that puts it at a disadvantage in gaining access to a channel. In many access protocols, the general rule for access is that a radio must wait until the channel is free. When a radio sees many other radios that wait is longer. Access becomes further complicated when the radios it sees do not see each other. This situation can result in the exposed node being starved of an opportunity to gain access. Figure 3-1 depicts an example. Here a terrestrial exposed node can see two separate networks separated by a terrain feature. For the node to gain access, both networks must be silent. These two networks operate on their own schedule and each network has silent periods, but they do not coincide, thus depriving the exposed node of the opportunity to gain access. This example is obvious but these types of situations can occur for an assortment of reasons that occur by happenstance because of movement and traffic loads.

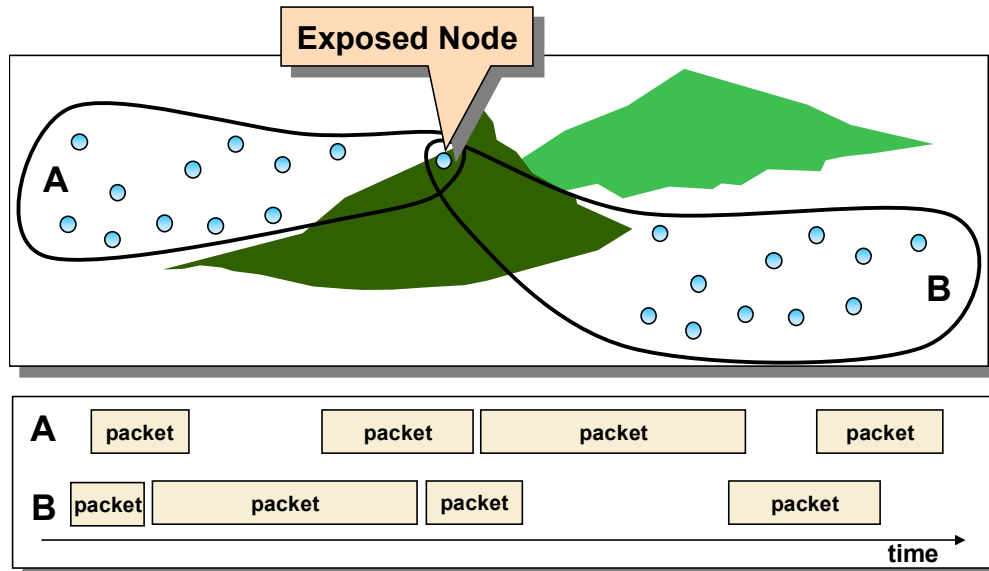


Figure 3-1. Terrestrial Exposed Node

Figure 3-2 illustrates some additional examples. In the first, an unmanned aerial vehicle is used to improve connectivity of a ground force. It is exposed to a large number of ground platforms. Those platforms may not see each other and so can gain access simultaneously and in a way that would prevent the UAV from getting access. The second scenario of this figure shows the role traffic can play. Nodes N1 and N2 are traffic hubs that many radios try to reach for access to another network or for data or services provided by those nodes. The larger circles around these nodes identify the range of their transmissions. The exposed node in this case is the node that is in range of both N1 and N2 and it must wait for both hubs to be silent to gain access.

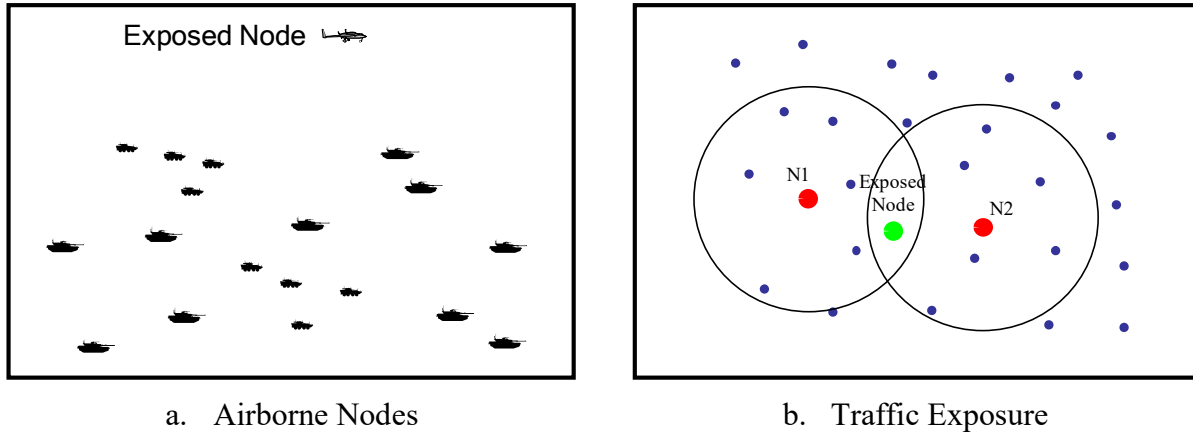


Figure 3-2. Additional Exposed Node Scenarios

Exposed nodes can also present a problem in access mechanisms that do not require silence for access. When an exposed node gains access, its use of a channel is much more costly in a geospatial sense because its access to a channel requires a larger proportion of peer nodes to defer their access.

3.3 Wireless Access Mechanisms

Many mechanisms enable access to communication channels. Most access mechanisms were first developed for wireline channels and then adapted for use with wireless channels. This section provides an overview of the basic techniques in wireless communications.

3.3.1 Scheduled Protocols

Protocols with scheduling subdivide the channel into time periods and assign some number of these time periods to each participant in the network. These resources remain assigned whether the participant uses them or not. The general benefits are that no channel time is wasted by contention, each radio is assured of getting an opportunity to use the channel, and most of the wireless limitations described above are avoided. The disadvantages are that there is no spatial reuse of the channel (a goal that results in the limitations above), there is a fixed membership in the network that cannot vary, and there is usually risk that the allocation of temporal resources to individual nodes will not match their traffic needs, and some nodes have too much resources and others have too little. When radios do not have traffic for the channel their allocations go to waste. When radios move they may lose connectivity but the resources for their links are not reallocated.

3.3.1.1 Time Division Multiple Access

Time Division Multiple Access (TDMA) is the most basic type of scheduled protocol and the term TDMA is generally used as a synonym for a scheduled access mechanism. TDMA was initially used in wireline networks and is applied in wireless environments in the same way. In the pure application of TDMA, upon configuration the channel is uniformly slotted and a proportional number of slots is allocated to each participant in the network. Network resources

cannot be reallocated while the network is in use and a change in network membership requires configuration changes to every radio in the network.

The slots of a network using TDMA may be assigned to the communications of pairs of nodes, for broadcasting from a node, or to a particular radio for use in communication to any of its peers as needed.

Because the radio channel is not as reliable as a wireline channel, wireless versions of the TDMA protocol often include signaling protocols in their packet formats that are used to confirm the reception of packets from other radios.

3.3.1.2 Hybrid TDMA

A Hybrid TDMA protocol is one that combines the use of a contention mechanism with a scheduling mechanism. The access schedule is dynamic and adapts by responding to the use of the assigned resources and to the demand by the nodes in the network for those resources. These protocols typically have three features; a set of timeslots in which communications can occur, a mechanism that advertises each nodes perception of the use of the timeslots, and then a mechanism through which nodes can contend for unreserved timeslots.

As expected and similar to TDMA, the transmission timeslots are organized on a period, for example they could be organized at 20 timeslots per frame that repeats itself. A radio that gains access to timeslot 5 would be able to use it in each frame until it releases the timeslot. Then, on some periodic basis, nodes transmit a signal indicating their perception of the use of the timeslots. As an example, the Wideband Networking Waveform (WNW), which implements this type of access mechanism, has each radio transmit several bitmaps indicating what timeslots it uses and what timeslots its neighbors use, with the implication that a radio should never contend for a timeslot in use by itself, its neighbors, and its neighbors' neighbors. A radio indicates it has freed a slot it is using by no longer advertising the reservation. Over time its neighbors will stop advertising its use as well. At the point when no neighboring radio advertises its use, another radio can gain access to the timeslot by advertising that slot as being used by itself in its bitmap. The reservation is made complete by all other neighbors indicating that the slot is not reserved by themselves but by a neighbor. As long as there are no conflicts the reservation is maintained. A reservation of this type can be made for broadcasting to neighbors or for communications with a specific peer. This approach can also support reservations across multiple channels in a common network. In the case the radios transmit a set of bitmaps for each channel on a common control channel.

This sort of mechanism allows timeslots to be allocated based on demand and to be spatially reused when the users are sufficiently far apart that they share none of their neighbors. These are both advantages over pure TDMA. Hybrid TDMA provides a level of certainty that the radio will retain capacity over time – a desirable characteristic when making routing decisions, while allowing the release of those resources when they are not needed. Finally, it can minimize the cost of contentions by controlling the rate at which they occur (e.g., the transmission of the bitmaps may occur once every 10 frames). Here, one successful contention creates a continuous access to the timeslot. However, there is an obvious tradeoff in the responsiveness of access. It may take several frames to release a timeslot and then several additional frames to reserve the slot. There is some risk that several radios hidden from each other will assert their interest in

using the slot in the same contention cycle, in which case a collision will occur. The access mechanisms must provide some means to recognize the conflict and to resolve it. In the case of WNW, the expectation is that any radio that detects a conflict with one of its reservations will cancel that reservation immediately and try to reestablish it later.

Contentions in Hybrid TDMA can use any of the contention mechanisms described in Section 3.3.2. There are various mechanisms for indicating channel activity, avoiding conflicting reservations, and indicating when a channel becomes free.

3.3.1.2.1 Challenges

Hybrid TDMA approaches, as described above are very sensitive to mobility. The act of moving can cause radios that have reserved the same slots to come into proximity to each other, where their use of the channel can interfere. This can occur because the radios move toward each other or because a radio becomes more exposed; for example, a radio on a helicopter that takes off and moves to a higher elevation or a radio on a platform that moves to high ground. This issue can become quite severe when the goal is to avoid using the same timeslot that any neighbor of a neighbor uses. The outcome can be a dramatic resetting of the reservations, requiring the network to almost start from scratch to reestablish its connectivity, thus rendering the network ineffective until this is resolved.

3.3.2 Contention

The primary goal of contention is to statistically multiplex traffic, which means the ability to self-adapt and provide capacity to individual radios based on the traffic they have and avoid the allocation of resources that cannot be used, which frequently happens with scheduled protocols. The anticipated outcome is a more efficient temporal use of the channel. In wireless domains, another objective is that contention-based mechanisms will contribute to the spatial reuse of the wireless channel.

3.3.2.1 Aloha

The Aloha protocol was first used in a wireless network across the islands of Hawaii, hence its name. It is the most basic of access mechanisms. A radio that receives traffic to send sends it immediately upon its arrival to the radio. There is no effort to develop awareness whether the channel is busy so transmissions of a peer radio can start at any time. As a result, there is risk of “collisions” where one transmission overlaps another.

A well-known phenomenon, congestion collapse, can occur in a network that uses Aloha. With the increase of traffic the capacity of the wireless channel collapses because of the increased failure of transmissions that result from the corresponding increase in collisions. Figure 3-3 illustrates this phenomenon mapping throughput and latency against load. In Aloha, the peak throughput is merely 18% of the channel capacity. In network instantiations where radios receive feedback on the success of the transmission, the failure to receive an acknowledgement from a destination radio can cause a source radio to retransmit. Since the acknowledgement transmission can also fail, the collapse illustrated can be even more severe because even successful transmissions will be retransmitted if the source radio is not made aware of the success. Figure 3-3 shows that the average latency of a successful packet (the time from arrival

to successful delivery) will explode before the maximum capacity is reached and so the preferred operating region would be somewhere below the maximum capacity where the latency is acceptable.

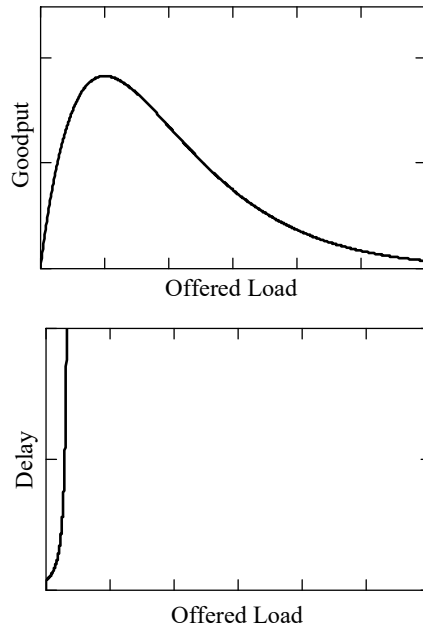


Figure 3-3. Typical Performance of Access Mechanisms Vulnerable to Congestion Collapse

This analysis may be pessimistic in cases that do not involve multiple radios contending for access in a network. If the offered load is mostly to a single radio that queues its packets and is performing most of the transmissions, the likelihood of collisions may be small.

Aloha is not used often, since other alternatives are more effective. In some cases, however, when there are few radios in a network or traffic is infrequent Aloha remains an economical choice. Nevertheless, it would not be well suited for arbitrating the access in a spectrum highway where congestion is likely and where radios must defer to each other.

3.3.2.2 Slotted Aloha

Slotted Aloha is a variant of Aloha in which the channel is time slotted and access attempts are made for the slots only. Transmissions do not extend past the slot boundaries. Slotted Aloha still suffers from congestion collapse, but the additional structure reduces the occurrence of collisions. This modification to the Aloha access mechanism doubles the theoretical maximum throughput to 36% of the channel capacity (assuming all successful transmissions fill the slots).

Slotted Aloha can be tuned to perform better by using a modification called p-Persistent Slotted Aloha when the quantity of contending nodes can be estimated. In this case, nodes with packets attempt to gain access to a slot with some probability that optimizes the chance that just one of the contenders is selected. However, when the quantity of contending nodes varies from the amount used in selecting the persistence parameter p , the result is either greater latency in delivery at the lower quantity or a lowering of the maximum throughput breakpoint as the number of contenders increases above the design quantity.

Slotted Aloha remains a frequently used contention mechanism and is often the preferred method for satellite access where propagation times and the separation distances among terrestrial contenders make other techniques impractical. Although Slotted Aloha and the p-Persistent variation are more effective than pure Aloha, these protocols have the same deficiencies when it comes to supporting a spectrum highway: they do not perform well in congested environments and do not have mechanisms to allow radios to defer to each other.

3.3.2.3 Carrier Sense Multiple Access (CSMA)

In CSMA, a contending radio first listens to a channel to determine if it is busy and then transmits a packet only if the channel is free. This access mechanism is the core of the very successful Ethernet protocol. In some cases, multiple terminals may have packets and simultaneously decide to transmit, but in Ethernet it is possible to detect collisions. Upon the detection of a collision, the terminals immediately stop transmitting and choose a random backoff time to wait before trying again under the expectation that the other radios will choose different backoff times. As described in Section 3.2.1, in wireless environments radios cannot detect a collision. Thus, sending a packet is a commitment and transmission continues until it is done. Further, when a collision occurs, the sending radio is unaware of the failure. Additionally, because the radios are dispersed and are unable to hear all peers, the protocol suffers the problems of hidden nodes (Section 3.2.3.1) and exposed nodes (Section 3.2.3.2), and performance degrades to that of Aloha. Like Aloha, CSMA can suffer congestion collapse and has no mechanism that allows radios to defer to each other.

Because of the deficiencies of CSMA in wireless, a pure version does not exist; instead, multiple other mechanisms are added on to improve its performance. The following subsections describe the additional protocol features that are frequently added to the CSMA mechanism to make it suitable for wireless media access control.

3.3.2.3.1 Collision Avoidance

Since wireless radios cannot detect collisions, packet collisions are costly in the unmodified CSMA. The channel is used but there is no goodput. Since the CSMA mechanism allows radios to start transmitting once they detect a channel is free, there is a high probability that collisions will occur immediately after the channel turns from a busy to a free state. Packet arrivals at multiple radios during a transmission will result in those radios being cued to start transmitting at the same time when the channel state transitions from busy to free. Collision Avoidance (CA) is a mechanism to prevent these types of collisions.

The standard CA mechanism is to require and randomize a backoff time after a channel returns to a free state. The goal is to cause radios to plan the start of their transmissions at different times and then require the radios to use the carrier sensing mechanism and defer to whichever radio starts first.

3.3.2.3.2 Mini-slotting

Radios cannot immediately detect the transmission of a peer radio. The signal of that peer must propagate to the sensing radio, and the sensing radio must receive some portion of the signal and process it in order to make a decision whether the channel is busy or not. After the detection of a free channel, a period of time is required to make the decision whether to transmit and to

transition to a transmitting state. These activities create a window of time during which radios employing CA could collide even though their scheduled transmission start times are different.

In the same sense that slotting improves the performance of Aloha by causing collisions to occur only in discrete time periods, mini-slotting improves the performance of the CA mechanism. Rather than allow the backoff time to be continuous, the time is made discrete with the minislots. All transmissions are to start at the beginning of a minislot. The minislot is sized to account for the longest of propagation times for the radio deployment, the time to capture and process a portion of a signal, the time to detect a signal and make a decision, and finally the time to transition to the transmit state.

3.3.2.3.3 Handshake and Virtual Sensing

Despite CA, collisions can still occur because of hidden nodes (Section 3.2.3.1). The sending of handshake packets is a solution used in many protocols. In the handshake, the source radio first sends a short Request to Send (RTS) packet announcing its intention to send a packet to a destination that is identified in the RTS packet. That destination responds with a Clear to Send (CTS) packet if it receives the RTS and perceives the channel to be free, which is the trigger for the source to send the full and larger packet. This mechanism has two purposes. First, if a collision does occur, or if the destination is either deaf (Section 3.3.2.5.1) or mute (Section 3.3.2.5.2), it only disrupts the shorter RTS transmission. Second the CTS packet can inform its neighbors of its intent to receive a packet and so mitigate the collisions that could occur because this source is hidden. These protocols use virtual sensing where the RTS and CTS packets inform neighbors of the intended duration of a packet exchange as part of the packet. This duration is frequently referred to as a Network Allocation Vector (NAV). A radio that correctly receives a RTS or CTS NAV executes virtual sensing, it assumes the channel is busy for the amount of time indicated in the NAV and acts as though it is, even if the radio does not detect any transmissions.

3.3.2.3.4 Acknowledgements

The only way for a radio to be certain that another radio has received its transmission is for the destination to confirm receipt. An acknowledgment is a packet or content within a packet sent by a destination to confirm receipt of the packet. There are two approaches to sending acknowledgments. In the first, the destination radio sends a separate acknowledgment packet (ACK) directly to the source following receipt of a payload packet, thus providing immediate confirmation that a packet is received. Not receiving this ACK indicates to the source that the transmission was unsuccessful. In the second approach, a destination radio sends a broadcasted acknowledgment indicating all the packets it has received over some period of time. These acknowledgements are frequently aggregated. Source radios have to wait until this packet is transmitted and so unacknowledged packet retransmissions will suffer greater delays than with the first approach.

3.3.2.3.5 Priority Based Windowing

CSMA relies on the randomness of channel access attempts for arbitrating access. The CA mechanism ensures that attempts remain random after a channel becomes free. In some protocols, the random backoff used for CA is computed for all packet arrivals. Radios wait for

some period after seeing the channel is free regardless of the channel state at the time of the packet arrival. An interruption during the waiting period by a packet transmission of another radio only curtails the backoff, which continues where it left off once the channel becomes free again. Priority is differentiated in these systems by allowing the radio to choose a shorter window of time for selecting the wait time for a higher priority packet thus giving it an advantage in gaining access quickly. As an example, there could be three priority levels where the radio selects a random minislot in a 10-minislot window for the highest-priority packets, a 20-minislot window for the medium-priority packets, and a 30-minislot window for the lowest-priority packets.

3.3.2.3.6 Exponential Backoff

CSMA protocols can also suffer congestion collapse. The susceptibility grows as the backoff windows become smaller. But using large windows can be very wasteful and contribute to latency when networks are not congested. Exponential backoff is a procedure to adapt and to mitigate congestion by increasing the backoff window size with every failed contention. For example, the backoff may double with every failed contention up to some maximum size. An upper limit prevents a radio from continuing to contend for an unavailable destination, and after some number of attempts, protocols typically assume the destination is unreachable and drop the packet.

3.3.2.4 Packet Sensing Multiple Access (PSMA)

A deficiency of CSMA is that it is quite vulnerable to jamming in adversarial environments. Radios using the CSMA mechanism look for energy in a channel to make decisions. Adversaries can put even small amounts of energy in the channel and radios will decide to defer communication, since radios use this detection to avoid interfering with what they assume are valid communications. CSMA does not distinguish a peer from an adversary that deliberately puts a busy signal signature on a channel to block communications. The counter to this vulnerability is to cause radios to look for specific bit sequences or to decode full packets to detect a busy channel.

Protocols based on packet sensing are very similar to wireless CSMA protocols and those who are not familiar with the distinction may call them CSMA. The typical packet sensing protocol uses the add-on features described above, and the packets that are sensed are the handshake packets: the RTS, CTS, and ACK. The NAV within these packets allows radios to perform virtual sensing. Radios constantly try to receive these control packets when they believe the channel is free and virtual sensing rather than carrier sensing causes radios to defer access. Unless an adversary can replicate one of these control packets, it will not cause a radio to defer just because the radio senses a busy channel.

A deficiency of packet sensing is that it requires mini-slotting and CA and the minislots must be large enough to accommodate the RTS, CTS, and ACK control packets. The minislots tend to be nearly 100 times larger than those for CSMA. The result is a much less efficient CA mechanism where the backoff (i.e., idle time on the channel) will take much longer.

3.3.2.5 Deficiencies of Carrier and Packet Sensing Protocols

The mechanisms used to avoid collisions and to overcome hidden nodes result in additional problems when applied in geographically dispersed networks, as described in the following subsections.

3.3.2.5.1 Deafness

Deafness is similar to the hidden node problem but is a phenomenon of virtual sensing. In this case, a node is in range of an active transmitter and so cannot receive a transmission from another network participant that might have sent an RTS or CTS and as a result is unaware of the channel state. This node is considered to be deaf. Later, if the node must transmit a packet it is unaware of the channel being in use and therefore could contend, gain access, and transmit, thus disrupting a packet reception by its neighbor.

3.3.2.5.2 Muteness


One of the mechanisms to mitigate the problem of hidden nodes is to have transmitters and receivers perform a handshake whereby the transmitter first announces its transmission with an RTS followed by the receiver acknowledging its availability with a CTS and announcing to its neighbors that it will be busy with the reception inserting a NAV in the CTS. The desired result is to mitigate the occurrence of hidden nodes so radios that receive the destination's CTS will wait to transmit until the reception is over. In this protocol approach, the failure of the receiver to respond to an RTS is taken as indication that it is busy or not available. However, there are cases where radios can receive a transmission but do not respond to an RTS because they are waiting until a neighboring receiver finishes its reception. These radios are considered to be mute. The node that sent the RTS, however, cannot distinguish this response from what would happen if the muted node were interfered with or had moved out of range.

3.3.2.5.3 Unfairness

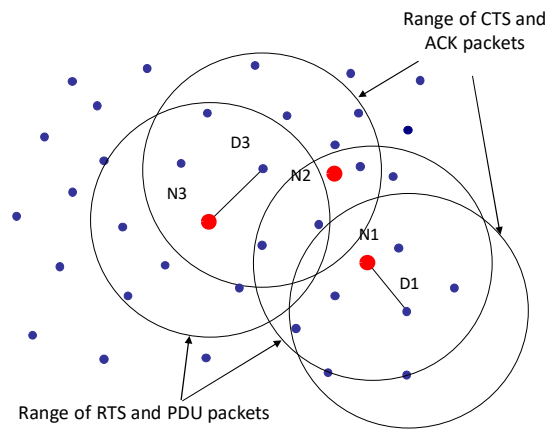
Unfair access may occur as a result of the exponential backoff mechanism. A radio that contends, sends an RTS and receives no response because the destination is deaf or mute will increase the backoff window of CA and try again. The radio that successfully sends a packet will begin its backoff for the next collision using the smallest window and so have an advantage in gaining access. A possible outcome, which partially depends on the radio that just completed the exchange having more traffic, is that the successful radio will continue to gain access while the radio with the larger backoff will continue to try to send its packets when its destinations cannot respond. Packet transmissions can occur across multiple backoff minislots, so the longer the packet transmission the more likely it is that collisions will continue. The network can become stuck in this clearly unfair condition.

3.3.2.5.4 Exemplar Scenario

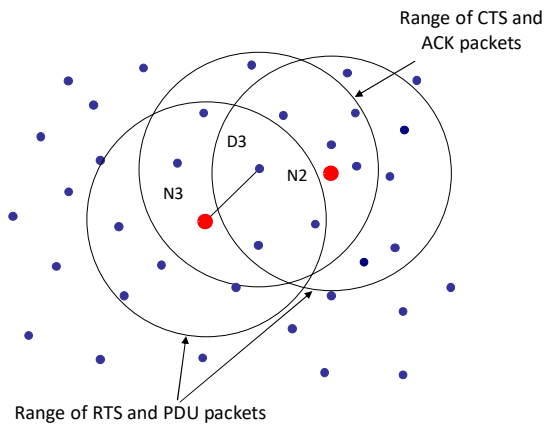
Carrier and packet sensing protocols work well when all radios are in range of each other, but can function poorly when they are not, even with the added features of collision avoidance, exponential backoff, handshakes, and virtual sensing. Figure 3-4 illustrates a scenario in which multiple radios that use these access mechanisms can interact in a compromising way.



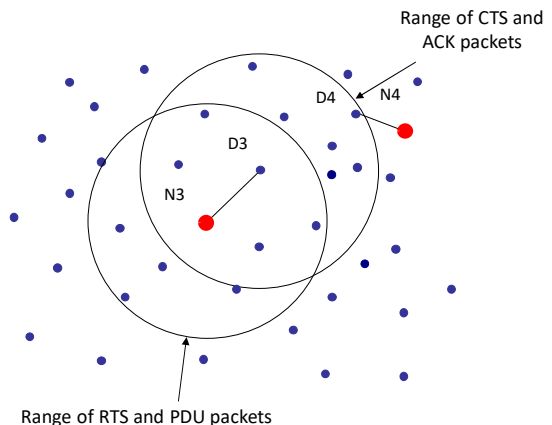
a. N1 is sending a packet to D1 when N2 has a packet to send. N2 waits for N1 to finish.



b. Meanwhile, N3 begins transmitting a packet to D3. N2 does not hear D3's CTS transmission. A.k.a. deafness problem.



c. N1 and D1 complete their exchange so N2, which is unaware N3 is sending a packet to D3, sends an RTS that interferes with the packet being received at D3. A.k.a. hidden terminal problem.



d. N4 wants to send a packet to D4 but D4 is deferring to the exchange between N3 and D3. D4 does not respond to N4's RTS's and so N4 backs-off and contends again. A.k.a. muteness problem.

Figure 3-4. Wireless Contention Access Scenario Showing Challenges in Using CA, Handshakes, and Virtual Sensing

In Figure 3-4a node N1 gains access in the typical way first sending an RTS packet, receiving a CTS packet from the destination D1, followed by a packet transmission from N1. During this transmission N2 becomes aware that it has a packet to send. Either because of hearing the carrier or because of receiving the virtual sensing guidance from N1 in the RTS or protocol data unit packet (PDU), it defers from gaining access. Then, as illustrated in Figure 3-4b, node N3, which is outside of the range of N1, contends and sends a packet to D3. Unfortunately, D3 sends its CTS during a period when N1 is transmitting and N2 does not receive it, thus breaking the expected virtual sensing mechanism. In this case N2 is suffering the deafness problem. The significance is illustrated in Figure 3-4c, which shows that Node N2 starts to transmit not realizing that D3 was in the process of receiving a packet, and so the transmission interferes and disrupts the packet exchange between N3 and D3. Nevertheless, the waiting caused by virtual

sensing still affects all radios in range of the destination, D3. So, as illustrated in Figure 3-4d, when the node N4 attempts to send a packet to D4, D4 receives the RTS but is mute because it is not supposed to transmit for the period indicated by the virtual sensing. Since N4's backoff is not curtailed by any other transmission N4 quickly goes through the backoff and potentially the multiple iterations of the exponential backoff. After the final backoff iteration it will drop the packet and might even assess that the destination D4 is no longer in range even though it is. This is an example of unfairness in action.

3.3.3 Synchronous Collision Resolution (SCR)

SCR is an alternative paradigm for wireless access where collisions are anticipated in access and signaling is used to resolve those contentions. Rather than causing radios to try access at different times, radios that have packets to send contend at the same time and allow the signaling protocol to resolve those contentions. A benefit of this access approach is that it also orchestrates the reuse of spectrum spatially [2]. SCR has four key characteristics:

1. The wireless channel is slotted.
2. All nodes with packets to transmit attempt to gain access every transmission slot.
3. Contending nodes use signaling to arbitrate their access.
4. All packet transmissions that occur during a transmission slot are sent simultaneously.

Design choices that determine the capabilities of SCR are the size and framing of transmission slots, the specific details of signaling, and physical layer characteristics that can exploit the synchronous and spatial nature of the access.

Figure 3-5 illustrates a general version of SCR. The channel is time slotted and Collision Resolution Signaling (CRS) precedes the slot during which a radio that wins the CRS may transmit a packet. In this version, the CRS is made part of the transmission slot but the only requirement is for the CRS to occur before the transmission slot.

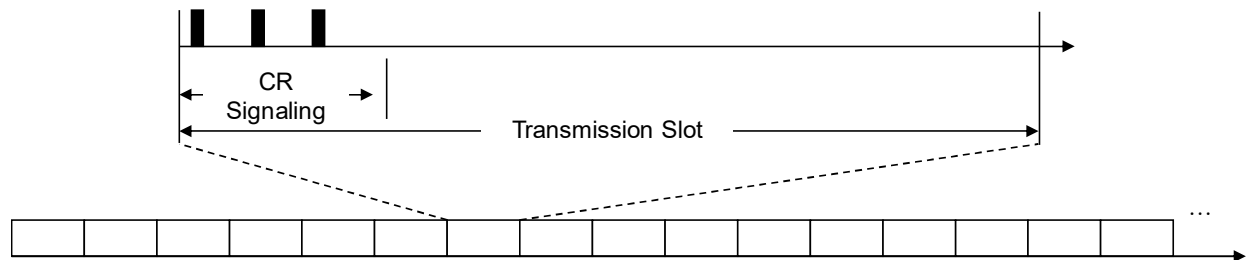


Figure 3-5. Generic Implementation of SCR

3.3.3.1 Collision Resolution Signaling

CRS consists of a series of signaling minislots organized into phases in which radios send signals to resolve contentions in the network for channel access. A minislot in SCR is the specific period of time in which a signal may be sent. A phase may consist of one or two minislots. In the one-minislot phase design, only the contending radios send signals. In the two-minislot phase design, the contending radio may send signals in the first minislot and any radio that hears

that signal then echoes the signal in the second minislot. In the two-minislot phase design, the resulting contention has a two-hop range.

Figure 3-6 illustrates a one-minislot phase CRS design. The rules for signaling and winning contention with this design are:

- At the beginning of each signaling phase a contending node determines if it will signal. (The contending node will signal with the probability assigned to that phase.)
- A contender survives a phase by signaling in a slot or by not signaling and not hearing another contender’s signal. A contender that does not signal but hears another contender’s signal loses the contention and defers from contending any further for that transmission slot.
- Contending nodes that survive all phases win the contention and may transmit in the corresponding transmission slot.




Figure 3-6. CRS with One-Minislot Phases

Figure 3-7 illustrates a two-minislot phase CRS design. The rules for signaling and winning contention with this design are:

- At the beginning of a CRS phase a contending radio determines if it will signal. A contending node shall signal in the contention minislot with the probability assigned to that phase.
- Any node that does not signal in the contention minislot but hears a signal shall send a signal in the second echo minislot.
- A contender survives the phase by signaling in the contention minislot or by not signaling and not hearing another contender’s signal in the contention minislot or a signal in the echo minislot. A contender that does not signal but hears another contender’s signal or hears an echoed signal loses the contention and shall defer from contending any further for the transmission slot.
- Contending nodes that survive all phases win the contention and may transmit in the corresponding transmission slot




Figure 3-7. CRS with Two-Minislot Phases, CRS with Echoing

CRS may be divided into two types of phases, priority and contention. Priority phases always precede contention phases. The rules above apply to contention phases. Figure 3-8 illustrates a

two-minislot phase CRS design with both priority and contention phases. The rules for priority phases are:

- At the beginning of a priority signaling phase a contending radio determines if it is the phase that matches the priority of its intended use of the transmission slot. If so, it will transmit a signal in the contention minislot of that phase. If not, it will listen for signals.
- Any node that does not signal in the contention minislot but hears a signal sends a signal in the second echo minislot.
- A contender survives the phase by signaling in the contention minislot or by not signaling and not hearing another contender's signal in the contention minislot or a signal in the echo minislot. A contender that does not signal but hears another contender's signal or hears an echoed signal loses the contention and defers from contending any further for the transmission slot.




Figure 3-8. CRS with Two-Minislot Phases, with both Priority and Contention Phases

Given n priority phases, the CRS design distinguishes up to 2^n priority levels. These priority phases can be mapped to specific transmission slot uses. Priority phases can be designed to accomplish a number of goals including distinguishing precedence, distinguishing transmission type (i.e., peer-to-peer versus broadcast), reserving capacity, and preempting reserved capacity. The use of the priority phases for these goals is further described below.

3.3.3.2 Design of CRS Contention

The design of the contention phases optimizes the probability of resolving to a single contender when all contenders are in range of each other. When radios are out of range of each other the protocol has the effect of also orchestrating the spatial reuse of the packet transmission slots.

Designing the contention phases of the CRS involves selecting the number of phases to use and assigning probabilities to those phases. The paragraphs below define a method to compute the probability that there is one survivor at the conclusion of the contention phases based on the probabilities assigned to those phases.

Let p^x be the probability that a contender will signal in phase x . Let \mathbf{P}^x be the transition matrix of phase x and \mathbf{Q}^n the transition matrix of the CRS design. The elements of \mathbf{P}^x may be defined using

$$\mathbf{P}_{k,s}^x = \begin{cases} \binom{k}{s} (p^x)^s (1-p^x)^{k-s} & 0 < s < k \\ (p^x)^k + (1-p^x)^k & 0 < s = k \\ 0 & \text{otherwise} . \end{cases}$$

where the entry $\mathbf{P}_{k,s}^x$ is the probability that s of k contenders survive the signaling phase. Note that s will never be 0 when $k > 1$ and will never be greater than k . The transition matrix of an n phase CRS design is $\mathbf{Q}^n = \prod_{x=1}^n \mathbf{P}^x$ and the probability that there will be just one surviving contender when there are k contenders at the beginning of signaling is $\mathbf{Q}_{k,1}^n$.

The objective of CRS design is to optimize the probability that just one node will survive the contention by selecting the signaling probabilities, p^x . Designing CRS to maximize the probability that just one node survives when k_1 nodes contend is relatively simple; however, a characteristic of CRS is that this maximum may result in a lower resolution probability when k nodes contend, $k < k_1$. The thick curve of Figure 3-9 illustrates the effect. The probability that there is one survivor when k_1 radios contend is greater than the probability when k radios contend in the range, $k_2 \leq k < k_1$.

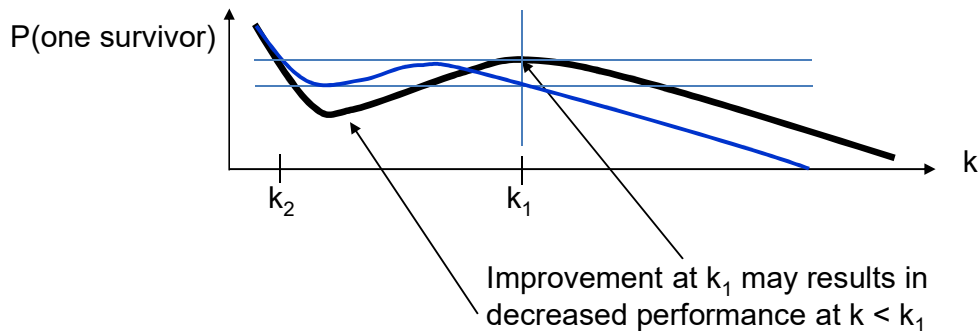


Figure 3-9. CRS optimization alternatives for a k_1 contender density

Thus the optimization used in CRS design seeks the result illustrated by the thin curve of Figure 3-9. The design maximizes the single radio survivor probability for k_1 without letting the single radio survivor probability for all k , $k < k_1$, to be less than that at k_1 . This design approach is defined more formally as follows:

Let q^n be the set of p^x for an n phase CRS design, k_t be a target density of contending radios, m be the total number of signaling slots allowed, and $S(q^n, k_t, m)$ be the probability that there will be only one surviving contender. Then the optimization problem is

$$\begin{aligned} & \max_{q^n} S(q^n, k_t, m) \\ \text{s.t. } & S(q^n, k, m) \geq S(q^n, k_t, m) \quad \forall k, 0 < k < k_t . \end{aligned}$$

The best solution for a finite set of signaling probability values can be found through an exhaustive search.

Table 3-1, Table 3-2, Table 3-3, and Table 3-4 list designs using this approach for sparse, $k_t = 50$, moderate, $k_t = 200$, and dense, $k_t = 1000$, radio densities for four different quantities of contention phases, six, seven, eight, and nine respectively. Figure 3-10, Figure 3-12, Figure 3-14, and Figure 3-16 illustrate the probability of one survivor with these designs as a function of the contender density. Figure 3-11, Figure 3-13, Figure 3-15, and Figure 3-17 illustrate the ratio of contentions without collisions to those with collisions for these same designs. Designs using more than nine contention phases are an extension of the nine-phase designs where the signaling probability of all additional phases is 0.50. The improvement in performance for each additional phase is the doubling of the contention collision resolution success ratio.

Table 3-1. CRS Design for Six Contention Phases

Phase	Design Density, k_t		
	50	200	1000
0	0.06	0.03	0.01
1	0.26	0.19	0.10
2	0.33	0.31	0.22
3	0.41	0.40	0.36
4	0.45	0.45	0.43
5	0.48	0.48	0.46
$P_{\text{One_Survivor}}(k_t)$	0.951	0.940	0.927

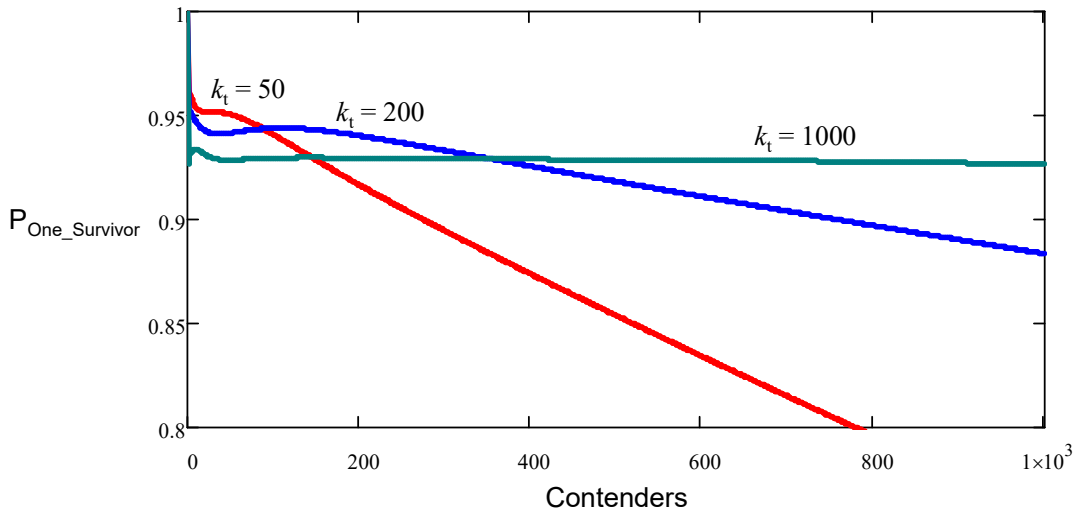


Figure 3-10. Performance of the Six-Contention-Phase CRS Designs




Figure 3-11. Contention Collision Resolution Success Ratio of the Six-Contention-Phase CRS Designs

Table 3-2. CRS Design for Seven Contention Phases

Phase	Design Density, k_t		
	50	200	1000
0	0.06	0.03	0.01
1	0.26	0.19	0.10
2	0.33	0.31	0.22
3	0.41	0.40	0.36
4	0.45	0.45	0.43
5	0.48	0.48	0.46
6	0.49	0.49	0.48
$P_{\text{One_Survivor}}(k_t)$	0.975	0.970	0.962




Figure 3-12. Performance of the Seven-Contention-Phase CRS Designs




Figure 3-13. Contention Collision Resolution Success Ratio of the Seven-Contention-Phase CRS Designs

Table 3-3. CRS Design for Eight Contention Phases

Phase	Design Density, k_t		
	50	200	1000
0	0.06	0.03	0.01
1	0.26	0.19	0.10
2	0.33	0.31	0.22
3	0.41	0.40	0.36
4	0.45	0.45	0.43
5	0.48	0.48	0.46
6	0.49	0.49	0.48
7	0.49	0.49	0.49
$P_{\text{One_Survivor}}(k_t)$	0.987	0.985	0.981




Figure 3-14. Performance of the Eight-Contention-Phase CRS Designs




Figure 3-15. Contention Collision Resolution Success Ratio of the Eight-Contention-Phase CRS Design

Table 3-4. CRS Design for Nine Contention Phases

Phase	Design Density, k_t		
	50	200	1000
0	0.06	0.03	0.01
1	0.26	0.19	0.10
2	0.33	0.31	0.22
3	0.41	0.40	0.36
4	0.45	0.45	0.43
5	0.48	0.48	0.46
6	0.49	0.49	0.48
7	0.49	0.49	0.49
8	0.50	0.50	0.50
$P_{\text{One_Survivor}}(k_t)$	0.994	0.992	0.990




Figure 3-16. Performance of the Nine-Contention-Phase CRS Designs

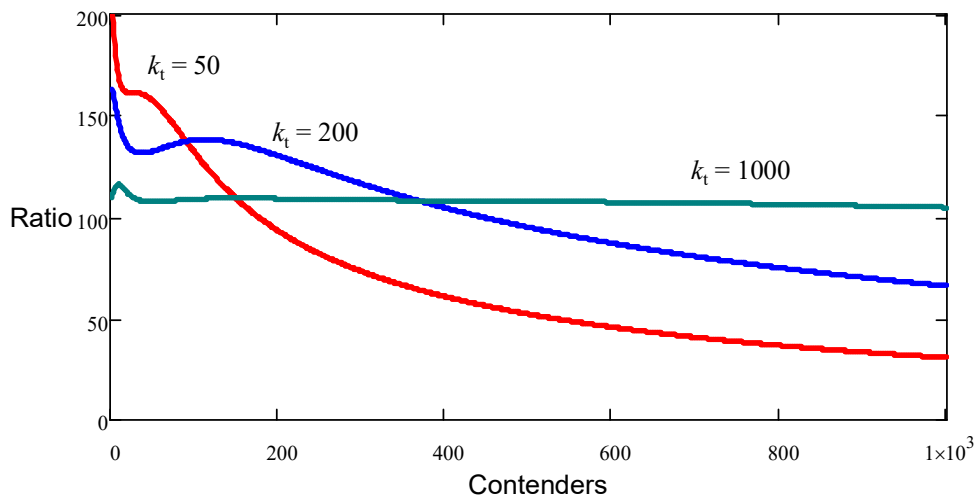
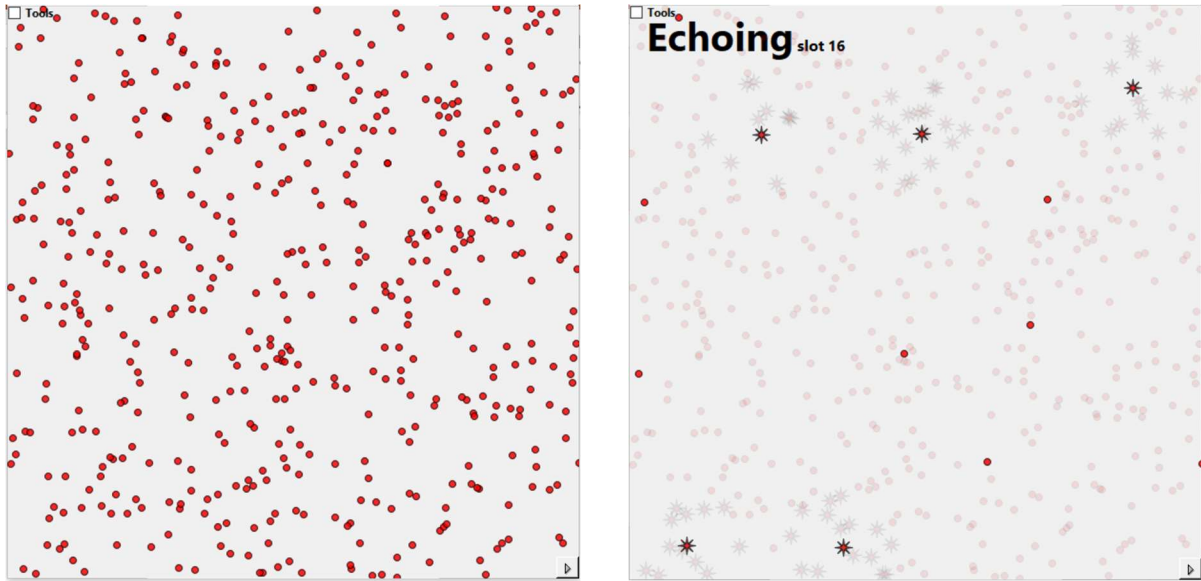


Figure 3-17. Contention Collision Resolution Success Ratio of the Nine-Contention-Phase CRS Design

3.3.3.3 Spatial Demonstration of CRS

The CRS mechanism is very effective at not only resolving collisions among a set of collocated contending radios but is also at orchestrating spatial reuse of a channel among a set of geographically distributed radios. Figure 3-18 illustrates a simulated scenario of 500 randomly placed radios that contend for access using eight contention phases with echoing and the final result. All 500 contend in this example and at the conclusion of CRS there are 13 remaining devices (the red circles and stars) spatially separated so that their transmissions do not interfere with each other.



a. Starting conditions of 500 contending devices

b. Final contention outcome with 13 surviving devices

Figure 3-18. Reduction of 500 Contending Devices to a Set of 13 Devices out of Range of Each Other

Figure 3-19 illustrates the simulation of the series of contention phases that resulted in the physical separation of the radios in Figure 3-18b. At the start of each phase each radio independently draws a pseudorandom number and uses the phase probability to determine whether or not to signal. For example, if the numbers drawn are between 0 and 1 and equally likely, a radio signals if the drawn number is less than the probability for the phase. The radios that hear the signal then echo that signal in the echo phase. This simulation shows the range of the signals to be the same in all directions. In actual applications, this type of signaling naturally works with whatever complicated propagation environment that might exist, taking advantage of terrain features that contribute to greater reuse. Further, in the rare cases when contending radios within two-hop range of each other end up surviving the contention they are often still successful in their use of the spectrum because of thinning out of the contenders and the differences in their use of the spectrum, for example, trying to communicate to destination nodes in opposite directions. A subsequent contention among the same set of contenders will result in a completely different set of survivors.

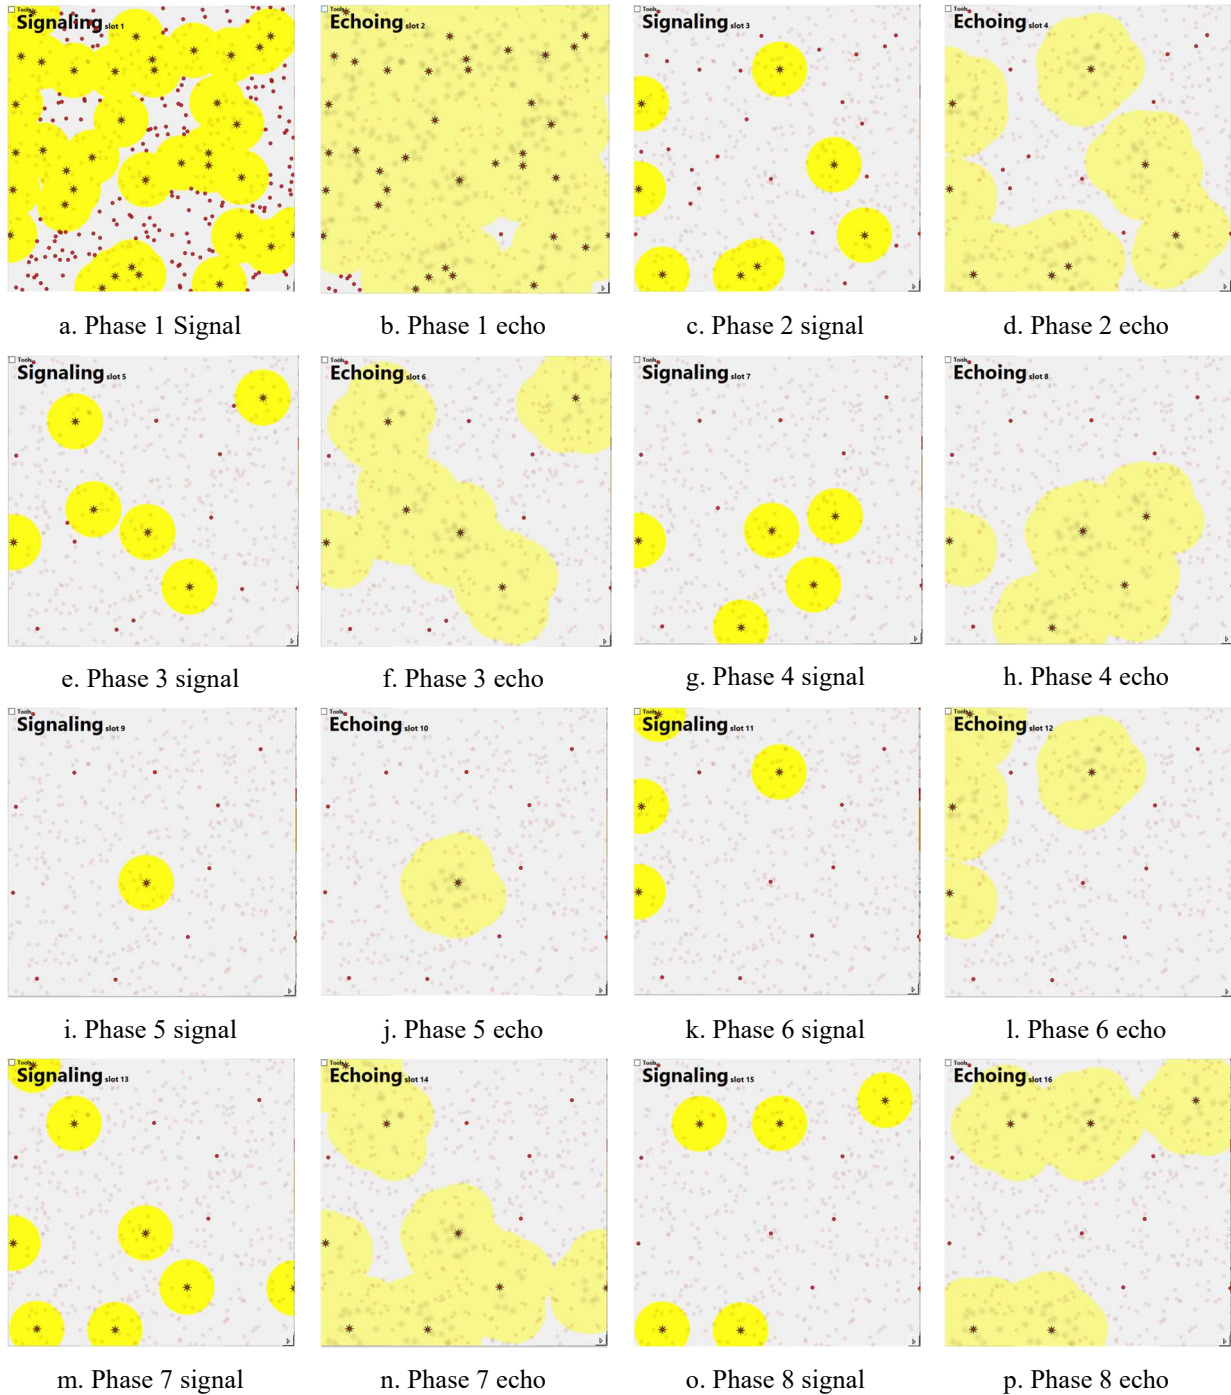


Figure 3-19. Simulation of an eight phase contention with echoing.

3.3.3.4 Designing Priority Phases

Priority phases are not based on a pseudorandom number draw but on a mapping to a particular access objective. That mapping can be to the user of the device (e.g. a user with precedence), the task the device is performing (e.g., transmission of higher priority traffic or jamming missions

over communications), or a particular performance objective (e.g., supporting a reservation for streaming traffic). The following subsections describe the techniques.

3.3.3.4.1 Precedence

Using precedence involves ranking the tasks performed by the devices and then mapping them to a priority. As has been demonstrated, whichever device signals first in CRS will have precedence. Priority phases are placed at the front of a contention and can be imagined as a binary number with the most significant bit being the first phase and signaling being associated with the 1 value. So with three priority phases, there can be eight priorities. The binary number 111 is the highest priority and corresponds to a contending radio sending signals in every priority phase. A contending device using a 110 priority will beat out a device using a 101 priority in the second phase. This losing device will learn that it does not have precedence in the second priority phase and will not signal in the third priority phase.

3.3.3.4.2 Reservations

Reservations involve the division of timeslots into repeating frames and the conditional use of an earlier priority in the set of priority phases within a timeslot of a frame based on a successful contention in the same ordinal timeslot in the previous frame. Consider the arrangement illustrated in Figure 3-20. Reservations are associated with two priority phases: one that is used to gain access the first time and then one to hold the reservation. In this four-phase design consider the use of priority phase P0 and P1 for reservations and then the priority phases P2 and P3 for precedence. So for routine traffic that does not require a reservation, differentiation for access is determined by the phases P2 and P3. These two phases allow differentiation of four priorities.

The reservation involves a two-step process. First, a device needing a reservation would send a signal using the priority phase P1. If that device successfully uses the timeslot as a result of that contention (e.g., exchanges a packet), it then can use the priority P0 for access in the same ordinal slot of all subsequent frames for as long as the device needs the reservation. The timeslot is immediately made available to other devices when the device stops signaling in this slot. If a device needs more channel capacity it can repeat the process for other ordinal slots of the frame.

Devices with reservations may be mobile. If two mobile reservations for the same timeslot move into range of each other their reservations may collide. The precedence priority slots can be used to indicate the precedence of one slot over another, so the lower precedence stream would yield the slot to the higher priority reservation. If the two reservations had the same precedence, then the contention phases can resolve which device keeps the reservation. The device that loses a reservation as a result of this contention can then try to set up another in another timeslot of the frame.

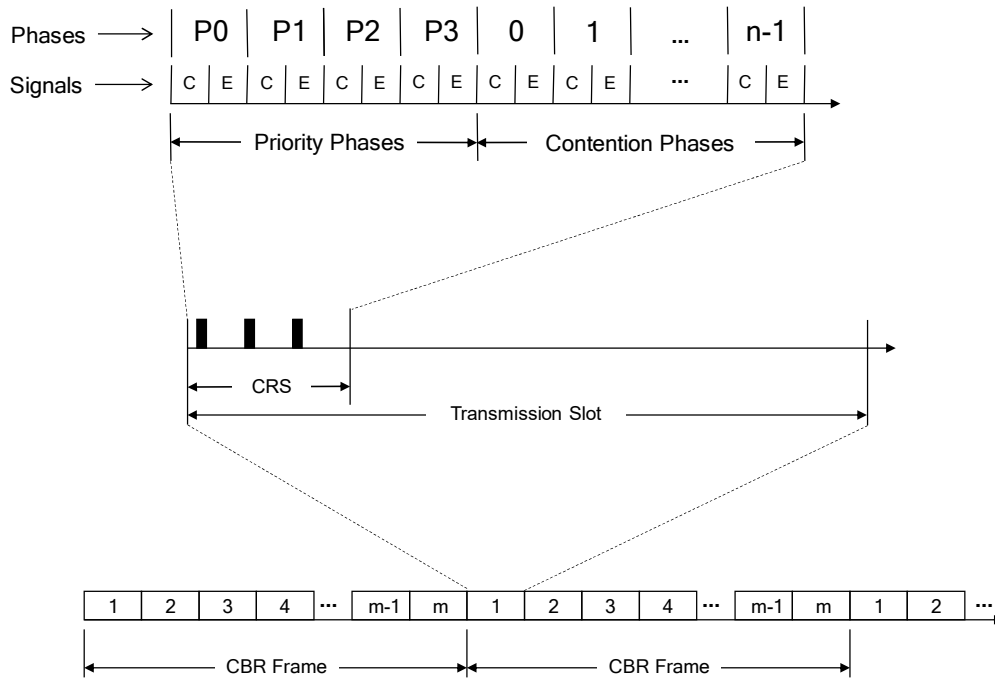


Figure 3-20. CRS and Frame Design for Periodic Slot Reservations.

3.3.3.4.3 Preemption

Some tasks may have higher precedence and urgency than others warranting the preemption of a reservation. Preemption can be built into signaling by giving devices a means to signal and beat a device with a reservation. Referring back to Figure 3-20 and the description of using precedence to distinguish reservation in Section 3.3.3.4.2, it should be noticed that the P1 phase is only used in the first contention of a reservation but is not used during the reservation itself. A radio with a need to preempt the stream can signal in both the P0 and P1 phases and through the signal in the P1 phase take the timeslot from the device with the reservation.

Signaling can be designed in other ways to get the preemption result. As an example, the two-step reservation may occur between the P2 and P1 phases; the P2 and P3 phases would still be used for differentiating precedence, and the P0 phase would be held for use by devices with tasks that require preemption.

Designs that support preemption may also establish rules governing whether the preempted device must rebuild the reservation or can assume that a preemption is short lived and attempt to maintain it. In the latter case, the device continues to contend as though it had retained reservation and the preempting device preempts until its needs are satisfied.

3.3.3.5 SCR Challenges

The primary drawback of SCR is its dependence on time of day synchronization. However, no particular level of synchronization is required until the slots of signaling are designed. A signaling slot must accommodate the maximum error in synchronization, the maximum

propagation time, the time to tune to the right frequency (if frequency hopping between signals), and the time it takes for a receiver to detect a signal, make a decision on the response and make any transition required to generate a response (e.g., changes in frequency and moving from a receive state to a transmit state or vice versa). Thus, the required synchronization is a function of design and the amount of overhead that the designer will allow.

The second drawback is the amount of overhead associated with signaling. As described above, a signaling slot must accommodate a number of activities and phenomena that can easily add up as more signaling slots are used in the contention. An advantage, however, is that once CS is designed, the overhead is fixed and the payload-to-overhead ratio of the design never changes and is not influenced by channel use as in the case of random access. This ratio can be increased by making the payload slots larger.

Third, the signaling range may not match the payload range. This may not be an issue at all. Generally signals would be designed to exceed the range of payload transmissions. This provides a more conservative spatial reuse that better avoids interference on the same channel.

When CRS designs do not use echoing there are cases where contentions can result in stalemated communications. Two nodes wanting to communicate to the same destination but out of range of each other could repeatedly win a contention and then interfere with each other in their payload transmissions. If the protocol is made reliable, (e.g., ACKs) are sent in reply, on not receiving an ACK these two nodes could repeatedly contend, win, and fail to reach their destination. This blocking situation is a motivation for using CRS designs with echoing.

3.4 Conclusions

A wireless channel access protocol will be a necessary part of a spectrum highway. Highways, will require a distributed contention access mechanism that provides the ability to differentiate access based on user, use, and QoS. Most contention access protocols provide no means to differentiate uses and those that do have variable duration of execution. The SCR access paradigm provides a mechanism that can support an arbitrary quantity of differentiation levels and is very effective at arbitrating contentions and orchestrating spatial reuse of an RF channel. The duration of its execution is fixed supporting the tight restrictions of the timeslot hierarchy of a spectrum highway. SCR is MITRE's choice for the access mechanism used in highway design; Section 4.1 summarizes the reasons for this choice.⁷

⁷ SCR is patented and has a government purpose rights clause. The author of this report is the inventor and holder.

4 Multi-System Contention Design

This section explores design options for arbitrating SDS access on a spectrum highway. Building on the review of channel access options presented in Section 3 and the spectrum highway requirements listed in Section 1.3.3, Section 4.1 quickly eliminates some standard MAC approaches from consideration because of their inability to meet the requirements and proposes the use of the SCR access paradigm because it does. Section 4.2 describes the design alternatives available to address the highway. Section 4.3 examines the physical layer design alternatives for the signaling used by SCR. With the SCR paradigm, many designs are feasible, but they have different capabilities and tradeoffs. Section 4.4 lists and describes measures that can be used to compare alternative designs. Section 4.5, and specifically Table 4-3, list the various highway design objectives, design approaches to achieve those objectives, and then the performance trades associated with each design approach.

4.1 Assessment of MAC Approaches for Spectrum Highway Use

Most contention based access mechanisms are unsuitable for managing access on a spectrum highway.

TDMA is not suitable because, by design, TDMA commits to the allocation of spectrum resource and does not support statistical multiplexing which is an intended outcome for a spectrum highway.

Aloha and Slotted Aloha are not suitable because they provide no means to create any of the features demanded by the spectrum highway. There is no ability to differentiate access either by bandwidth or by use. Success is based totally on happenstance and the traffic load. They are also a very inefficient access mechanisms compared to others.

Carrier sensing employing LBT and CA would not support a highway design. This protocol does not accommodate systems that listen only. It does not provide an effective means to differentiate service and it suffers from hidden terminal and exposed nodes. It strictly statistically multiplexes traffic providing no practical means to reserve capacity or to preempt such reservations. There is no clear way to differentiate access for lanes and combinations of lanes that may be used by particular systems.

PSMA (i.e., using the exchange of RTS and CTS packets) could mitigate hidden terminals and support listening SDSs (the RTS-CTS exchange reserves a time period for listening to occur) but it still suffers most of the other limitations that carrier sensing suffers. There is no clear way to differentiate service, reserve capacity, or support preemption. Proposed methods to differentiate access increase overhead and are not guaranteed. This technique cannot simultaneously support multiple lanes and then the combining of those lanes into larger lanes.

An alternative to a pure application of these access mechanisms is to create a hybrid approach, using them in combination with TDMA. This matches the highway time slotting discussion of Section 2.3.3. In this approach, times are set aside for contentions that match to the use of time slots. The overhead associated with random access contention can be very large. The design alternatives will either result in a large lag between contention and spectrum use as seen in WNW's access mechanism described in Section 3.3.1.2 or in more overhead than is practical. A

feature of random access is the random wait of CA. Mitigating collision effects that result from delays for propagation and signal processing usually results in discrete increments, aka minislots, for delays and accommodating a lot of users requires a lot of minislots. Add onto the requirement to be able to differentiate bandwidth of channel and precedence of use, there can be an inordinate amount of minislots. If minislots are constrained to a small number per contention, there is a risk of the access mechanism being very sensitive to congestion and congestion collapse being a problem.

As a final alternative, a TDMA scheme could be coupled with a separate orderwire channel for access. The advantage here is that none of the bandwidth of the highway would be consumed by the contention. All of that would occur on the orderwire. However, the fact that the contention is moved to an orderwire does not fix the risk of the problems above. The contentions for the multiple possible channels would have to fit in the duration of a transmission opportunity. Service differentiation and congestion collapse would remain challenges.

These deficiencies immediately drive us to considering the SCR access paradigm described in Section 3.3.3 as the best choice for a spectrum highway. It provides features that differentiate multiple precedence levels with confidence that higher precedence users will win and it can be used for both channel selection and for differentiating users. There is a good science concerning its performance in congestion (see Section 3.3.3.2). It is effective at arbitrating space and orchestrating special reuse (see Section 3.3.3.3). It is a very fair protocol among contenders with the same precedence.

4.2 Access Design Alternatives to Achieve Requirements

Given the use of the SCR access paradigm, this section describes the alternative signaling designs for achieving the requirements for the spectrum highway. Design tradeoffs are called out.

The CRS design for access must resolve access based on four criteria, precedence among contending devices on which should select the active lanes, precedence among overlapping lanes, precedence of the use, and resolution of contention among contending devices with the same precedence. Additionally, it must support different types of traffic classes and different physical layer capabilities.

Thus the signaling design may have three sets of priority phases, one set to resolve who selects the active lanes, one set to select the active lanes, and one set to resolve the precedence of use. Figure 4-1 illustrates a signaling design with these four sets of signaling phases. Table 4-1 describes the purpose, requirement, and design trades for each of these phases. The set to resolve precedence of use will also resolve QoS issues such as supporting reservations. The fourth phase will be the contention resolution described in Sections 3.3.3.1 and 3.3.3.2.




Figure 4-1. CRS Using Four Sets of Phases for Lane Precedence, Lane Selection, Use Precedence, and Contention Resolution

Table 4-1. Purpose, Requirements, and Design Trades for the Four Possible Phase Sets

Phase Set	Purpose	Requirement	Design Trades
Lane Precedence	Lane precedence is used to differentiate among contending devices and decide which has precedence in selecting the active lane.	This is an optional phase that is used for highway designs that have overlapping lanes and where some SDSs that use the highway have a clear precedence in selecting the active lanes.	The quantity of differentiation levels is a function of the quantity of phases.
Lane Selection	Lane selection is used to allow contending devices to indicate which lanes are active in the upcoming contention.	A lane selection phase is required for all highway designs that have an explicit definition of lanes that overlap other possible lanes. If there are no overlapping lanes this phase is not used.	The more complicated the lane structure, the more phases are necessary to provide the variety of signaling combinations that avoid ambiguity about which lanes are active at the conclusion.
Priority	Priority is used to differentiate among users, uses, and in certain applications, usually communications, QoS. It is also used to differentiate the duration of a contention access.	This is an optional phase that is used in any highway design where there is a need to differentiate access based on user, use, or QoS and in any design that has a hierarchy of transmission opportunities.	This is perhaps the most difficult phase to design, since it balances four independent factors for giving precedence to a device. As the levels of differentiation increase, so too does the quantity of phases in this phase set.
Contention Resolution	This type of signaling is used to resolve contention among multiple contenders that survive all preceding phases. This signaling also orchestrates the spatial reuse of lanes.	This is a required phase in all contention designs. It is the phase that resolves contention when there is more than one contender for the same lane.	The number of signaling phases and the selection of signaling probabilities determines the effectiveness of contention resolution. Effectiveness increases with the number of phases but this increases the overhead. Probabilities can be chosen to target the anticipated quantity of contenders with less effectiveness for smaller quantities.

This design places precedence of lane selection above use precedence. The rationale is that once the active lanes are decided then individual contenders will move to the separate contention channels that are matched to the particular highway lanes that remain active and in which they intend to operate.

The fundamental assumptions about the devices that participate are that they can all meet the same performance requirements that are necessary for effective signaling. The key performance criteria are that the devices be synchronized to the time of day within a specified tolerance used in the signaling design, can execute the type of signaling required, have the processing power to detect a signal and to decide the next signaling action within a specified tolerance, and, if the signaling requires frequency hopping, that they can change frequencies within a specified blanking interval.

4.2.1 Differentiation of Service and Precedence of Access

CRS provides a means to differentiate both service and precedence of access. Precedence of access consists of distinguishing between one user or use over another based on a privilege associated with the user or the use; for example, a more senior user in an organization may get precedence over junior users. A particular type of device may perform a function that has precedence over another; for example, a radar may have precedence over the routine communications of a networking radio. Differentiation of service refers to providing resources to achieve a performance goal for a system, such as ensuring high-priority packets are transmitted before low-priority packets, sufficient time is provided for a radar echo to be received, or resources are set aside to support streaming services such as voice and video without interruption.

4.2.1.1 Lane Precedence

The purpose of the lane precedence phase set is to overcome the natural precedence of overlapping lanes over narrow lanes in the lane selection phase set. This natural precedence can be wasteful. For example, a lower precedence user may select an overlapping lane for its use and then another higher precedence device gains access. The higher precedence device beats out the lower precedence device, but does not need the full bandwidth of the active lane. The lane precedence phase set avoids this waste by allowing the higher precedence device to preempt the lower precedence device from selecting the active lanes.


4.2.1.2 Lane Selection

Creating a mechanism to arbitrate lane selection allows certain devices to contend for the higher bandwidth lane in a single contention rather than placing the burden on the device to contend for and win all of the lanes that make up the larger bandwidth lane. The undesirable outcome of the latter approach is that the device could win some subset of lanes and then not be able to use them because the larger bandwidth lane was not won and so in the process waste the use of the lower bandwidth lanes that were won. The goal of supporting lane selection contention is to identify the lane that will be used early so that all users operating in that bandwidth can contend for that lane rather than for the lower bandwidth lanes. The lane selection contention would take place on a contention channel in which all SDSs using the spectrum highway participate. After the lane selection phases, the SDSs go to the contention channels of the active spectrum lanes. This is intended to prevent SDSs from contending for lower bandwidth lanes on their respective contention channels and not participating in the contention for the larger bandwidth overlapping lane.

Figure 4-2 illustrates an exemplar of a set of spectrum lanes that could be defined for a spectrum highway and identifies the particular precedence for those lanes. Lanes 1 through 4 are the lower bandwidth channels that cover all of the spectrum of the spectrum highway. Lane 5 covers the spectrum of Lanes 1 and 2 and when it is active it precludes Lane 1 and 2 from being used individually. Similarly, Lane 6 covers the spectrum of Lane 3 and 4 and when it is active precludes Lane 3 and 4 from being used individually. There is a similar story for Lane 7 and Lane 8. Lane 7's operation prevents all lanes except Lane 4 from operation and when lane 8 is established, no other combination of lanes of the highway can operate individually.

The channel precedence contention has three phases as depicted in Figure 4-1. Figure 4-2 illustrates both the channel contention used for each lane and the contention observations that indicate a particular lane is active. The three-bit binary number corresponds to the signal combination in the channel contention phase. The 1 bit indicates a contention signal was transmitted in a phase and the 0 bit indicates that no signal was sent. The bold binary number indicates the contention that an SDS would use to contend for and make active a particular lane. The asterisks indicate the observed signaling combinations that indicate that the radio making the observation can contend for the channel. The exception is for Lane 8. A radio may contend for the 110 combination if those signals are received in the first half of a signaling phase and should not contend if the 111 contention was observed.

In this process the standard rules of contention are slightly modified. A device contending to select a lane is not preempted by the presence of an earlier phase with a signal but by the presence of signal(s) in earlier phases that indicate an active lane that would preempt its own lane selection. For example, the presence of a signal in the first phase would preempt selection of lanes 1, 2, 3, 5, and 6, but not lanes 4 or 8. Also, the observation of a particular signaling combination that does not match the combination for selecting an active lane does not mean that the lane is not available. Thus the design specifies the full set of possible channel contention outcomes that indicate a lane is active. For example, an SDS knows



that Lane 4 is active if it observes any one of the signaling outcomes 000, 001, 100, and 101. A listening device may perceive signal combinations that are not associated with a specific lane because it is in a position between two contention areas that choose different active lanes. For example, there is no particular lane associated with the signaling combination 101. This signaling combination might be observed because the listening device straddles two regions where on one side Lane 7 is active and on the other Lane 5 is active. In this example of lane selection signaling design, region straddling is also the case for the combination 111.

4.2.1.3 Use Precedence

When multiple types of SDSs share spectrum, there may be a precedence associated with the SDS type (e.g. a radar may have precedence over a communications system) or with the mission (e.g. a fire control network may have a precedence over a logistics network). Differentiation of this type of precedence is a standard application of CRS. Given n phases, there are 2^n possible precedence levels. The highest precedence is associated with the larger number. A number in the range of 0 to $(2^n - 1)$ is assigned to each SDS that participates. The default is precedence 0. Planners of the highway decide which precedence each SDS has. In the contention, the higher precedence devices would signal early and eliminate the lower precedence devices which would not contend any further.

4.2.1.4 Quality of Service (QoS) Precedence

It is often desirable to distinguish access among similar spectrum users based on what they individually want to do after gaining access. The most common QoS application is to distinguish traffic types in a network; for example, ensure devices with a priority packet get access before devices with a routine packet. With a contention used for highway access the QoS precedence allows multiple SDSs with the same use precedence to compete based on the priority of their immediate task. Here again, the design is a strict CRS priority design where n -phases can support 2^n priority levels and the largest priority value has the highest precedence.

4.2.1.5 Reservations

A reservation allows an SDS to gain periodic access to a lane. Such an access can support streaming service, guarantee a throughput, or guarantee a device can execute a task that requires this type of periodic access. Section 3.3.3.4.2 describes how to use signaling to reserve a timeslot. This signaling is combined with a slot design that divides a fixed quantity of slots into repeated frames. Figure 3-20 illustrates an exemplar design. Access is a two-step process. Once a radio gains access using a standard priority, it can use the reservation priority for subsequent access to the same slot of each subsequent frame. This is a use it or lose it reservation and when the radio ceases to need the reservations it simply stops contending for the slot.

4.2.1.6 Preemption of Reservations

There are times when reservations must be preempted either because of the urgency of task to be performed in the slot or because the reservation has become dysfunctional. An example of a dysfunction is a hot mike where a radio is accidentally keyed for a voice transmission when there is none. A reservation scheme may support a streaming voice application. The radio would be indifferent to why a mike is keyed and the user of the device would be oblivious. Using a

preemption mechanism, a distant device on the same network could preempt the stream and thus remotely turn off the hot mike. This assumes that the user would have to rekey the mike to regain access.

Section 3.3.3.4.3 describes alternatives for designing preemption into a priority phase.

4.2.2 Contention Channels

Supporting contention for multiple lanes requires providing the means for multiple contentions. These contentions can occur simultaneously on separate channels or serially on the same channel. Contention channel designs must also consider the time required to react to the results of a contention and the resilience of the channel to bad actors who could disrupt a highway by attacking the contention process.

4.2.2.1 In-Band Channels

The term “in-band” describes channels that are the same channel for contention as those used for the SDS functions. Using an in-band channel consumes part of each timeslot in which it occurs. During a contention, that is all that occurs; thus its duration is an overhead that reduces the capacity of the highway.

Since multiple lanes could be active multiple contention channels will be required to support contention on each of those lanes. However, the identification of active lanes must take place on a channel in which all spectrum highway users participate, while the contention for each active lane must occur on a separate channel for that active lane. So a possible design, using the lane design of Figure 4-2, could be a contention channel assigned to each of Lanes 1, 2, 3, and 4 and then the rule would be that lane identification signaling and Lanes 5, 7, and 8 would use Lane 1’s contention channel and Lane 6 would use Lane 3’s contention channel. As a way to suppress contentions by devices unaware of the lane selection (this outcome would occur by exception), those devices contending for the larger combined lanes would simultaneously send signals on the contention channels of the overlapped lower bandwidth lanes but would only respond to signals on the designated contention channel.

Figure 4-3 illustrates an example of the in-band signaling and the access to the various lanes of the highway as depicted in Figure 4-2 and as described above. Contention happens immediately before the payload section of each timeslot. The signaling channel for Lane 1 is always used, initially for lane activation and then for whatever lane activation overlaps Lane 1. This illustration demonstrates all of the possible outcomes for lane selection as well as the corresponding contention channels used for those combinations. Although shown as being within the confines of the bandwidth of the lanes, it is perfectly practical to frequency hop the signals of the contention across the full bandwidth of all the lanes since no active transmission occurs during this time. The contentions for each of the lanes would frequency hop and be scheduled so that they are always on different frequencies.

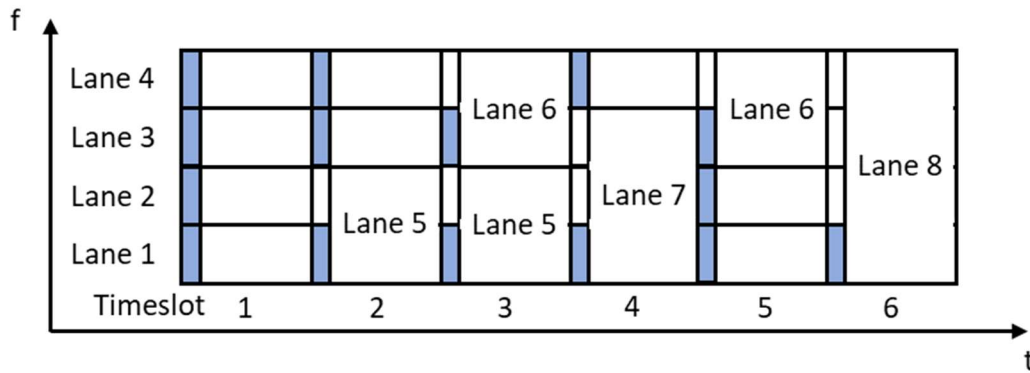


Figure 4-3. Example of In-Band Signaling

4.2.2.2 Out-of-Band Channels

The term “out-of-band” describes a channel that does not occupy the same spectrum as the highway lanes. The advantage is that there is no signaling overhead that consumes the capacity of the lanes; however, there is a spectral loss in supporting the contention channel. Contentions occur simultaneously with the timeslots but on the separate channel. With using a separate channel comes the option of signaling in a completely different band: one where the propagation is better or one where it is easier to hide the signaling to obtain greater resilience. Use of this approach requires devices to be coupled to a separate signaling transceiver in order to participate in the contentions. The contention channel would be time slotted in some manner that serializes the contentions for each lane. The contention to designate active lanes would first be followed by the contentions for each of the lower bandwidth lanes. An advantage of this approach is that all devices can be aware of the activities on all lanes. This allows them to rapidly change lanes as described in Section 4.2.5. The contentions themselves can take place in the timeslot immediately preceding the timeslot in which the access takes effect or some other specified time before it as necessary to allow SDSs to react to the results of the contention.

Figure 4-4 illustrates an example of an out-of-band contention channel. The contentions happen serially in time starting with the contention for the lane activation and followed by the contentions for each of the active lanes. This contention example matches the equivalent lane activation result illustrated in Figure 4-3. In this example, for the same duration of time an additional timeslot is available. Figure 4-5 illustrates the very same result for out-of-band signaling but with a shifted signaling design. Such a shift might be used to give SDSs more time to react to the outcome of the contention. With these benefits comes the trade that additional spectrum must be provided for the CRS and there is a gap between the time when the contention is won and the lane can be used.

Figure 4-6 illustrates the series of signaling phases used to arbitrate access to the lanes of a timeslot. At the start, contenders decide who has precedence in establishing the active lanes. The survivors of this set of phases then contend to select the active lanes in the lane selection phase. Once the active lanes have been established, contentions occur in the lane contention phases to decide which device gets access to those lanes. The illustrated priority and contention phase sets are the same for all four lanes.




Figure 4-4. Example of Serial Out-of-Band Signaling

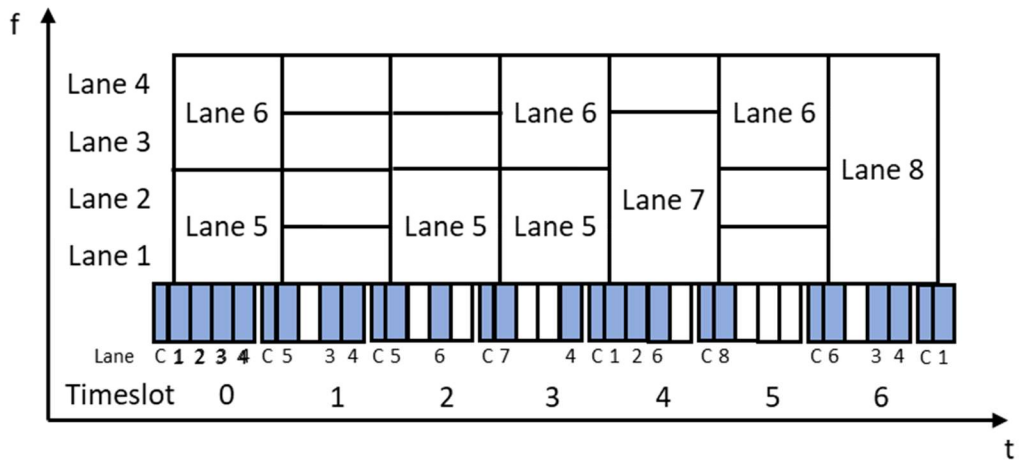


Figure 4-5. Demonstration of a Time Shift of the Serial Out-of-Band Signaling

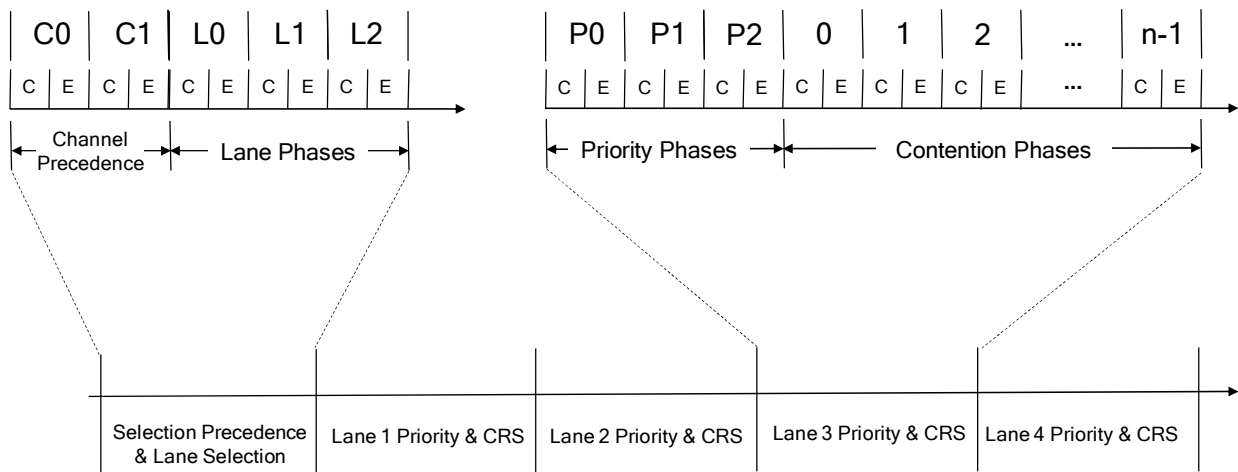


Figure 4-6. Example of Serial Contention for Spectrum Highway Lanes

Out-of-band signaling may also use multiple channels. The motivation for using separate channels for each of the lanes are to reduce the duration of the signaling so that the timeslot transmission opportunities can be shorter, gain more time in the signaling to have more precedence levels, or allow more contention phases for better contention resolution. Figure 4-7 illustrates this type of out-of-band signaling. Although the signaling channels are shown to be adjacent and of the same bandwidth as the lanes, this is not likely. In practice, the signals could be much narrower, be frequency hopped, and be interwoven in a much smaller band.⁸ (see Section 4.2.2.4.)




Figure 4-7. Example of Concurrent Out-of-Band Signaling

Since the contention device would be a common type of transceiver, the potentially large market for the transceivers (the SHADs described in Section 1.4) to be coupled with the SDSs of the highway would allow development and production at scale that would drive down prices and achieve high performance. SDS developers would buy and couple their SDS to a particular SHAD which gives their SDS access to a defined spectrum highway. As illustrated in Figure 4-8 the SDS would inform the SHAD of its access needs and authorizations for precedence and the SHAD would inform the SDS when access is gained. The SDS would then use the spectrum.




Figure 4-8. Coupling of an SDS with a SHAD to Arbitrate Spectrum Access

⁸ All SDSs operating on a highway use the same physical layer for their signaling.

4.2.2.3 Consolidated Signaling

Reacting to the results of signaling can take time because of the processing associated with preparing the SDS to take advantage of the access. For example, an SDS that transmits or receives data and can use variable bandwidth channels, would have to create a packet that can fit in the transmission opportunity it ultimately wins. Additionally, changes in channel bandwidth would also require the SDS to configure itself to the physical layer required for the lane that is won.

The concept underlying consolidated contention is to allow all the CRS for the timeslots of a frame to occur back to back before the frame begins. Every SDS would then have early knowledge of the activities of the lane and could schedule its own activities during the frame that allow it to participate in the lanes in which it assesses it should participate. Figure 4-9 illustrates an in-band consolidated CRS where the contentions are consolidated for the next six time slots. These signals and the results are identical to those of Figure 4-3 except for the consolidation in time when CRS takes place. This approach to in-band signaling would allow SDSs to access lanes in consecutive timeslots and to use those resources without interruption either as a transmitter or a receiver.




Figure 4-9. Example of in-band signaling using consolidated contentions

Part of the reason for using consolidated signaling is to give devices more time to react to the results of lane contentions. Figure 4-10 depicts two different slot designs with ten timeslot frames but with one where the contention occurs immediately at the beginning of the frame and a second where the contention for the timeslots of a frame occurs two timeslots before the start of the frame. This illustration shows that the set of consolidated contentions can take place at an arbitrary time before the first use that results from these contentions in order to support the planning and device transitions.




Figure 4-10. Alternative Schedules for Consolidated Contention in 10 Timeslot Frames

The disadvantage of this contention approach is the delay between contention and slot use. This reduces the ability to react quickly to changing needs, on a frame-by-frame basis rather than a slot-by-slot basis. It also reduces the ability to react to conflicts that result from mobility; resolution can only be made frame to frame.

4.2.2.4 Resilient Channels

The contention channels are an attractive target for bad actors wanting to disrupt highways by signaling unfairly and seizing channels or simply jamming the channel so signal detection is difficult. Creating resilient channels makes it more difficult for outsiders to deny service in these ways. The techniques to harden the signaling involve designs that make the signals themselves unique and more difficult for an adversary to replicate. The most obvious technique is to frequency hop the signals in a pattern that is unknown to the bad actor. The wider bandwidth across which the signals hop makes pure noise jamming more difficult and the natural short duration of the signals would make it hard for the bad actor to detect and quickly relay the signals. Assuming the hopping pattern is well randomized among many alternative frequencies, it should also be difficult for a bad actor to anticipate and send a signal on the correct frequencies of each contention slot.

An advantage of in-band signaling, since it is done simultaneously for all lanes of the highway, is that the signaling can be frequency hopped across the entire bandwidth of the highway. The signals themselves can be narrowband, giving them a power advantage against a jamming adversary and providing more alternative frequencies that they can hop across.

Out-of-band signaling has the advantage that it can occur on completely different bands from the highway itself. As a result, techniques could be used to make the signals resilient by using spectrum resources that are different from those of the highway and possibly exceed the bandwidth of the highway. Frequency hopping narrow band signals is an option. Alternatively, the signaling could possibly use a very covert means of signaling that involves spreading the signals across a very large bandwidth.

4.2.3 Contending for Alternative Duration Access

A disadvantage of time slotting is the expectation that all uses fit into fixed-sized timeslots. Five approaches support varying duration uses.

The first approach is to make timeslots rather large so that most uses are accommodated. The disadvantage of that approach is that it can result in considerable waste if the timeslots are often unfilled when used.

The second alternative is to design the time slots with a variety of durations as illustrated in Figure 2-6. Devices that require long-duration access would contend for and use the longer time slots and those requiring short-duration access would contend for and use the shorter time slots.

The third approach requires use of consolidated contention. Here all the contentions are completed before a frame allowing a radio to prepare to exploit the result of the contention. A device wanting a longer duration access would attempt to contend and win consecutive timeslots. The rules of the contention combined with rights given to the device may allow certain devices to have a higher precedence for some number of consecutive contentions after a first contention win using a normal precedence, thus biasing access to support merging consecutive timeslots.

The fourth approach can be implemented using out-of-band signaling. The expectation in out-of-band signaling is that CRS for the next timeslot occurs during the current timeslot, and so a device using a timeslot can win the contention for the next timeslot and not be interrupted regardless of how it is using the highway lane. This may not work for packet radio networks, since packetization must know whether the longer slot is available before creating the packet. Packetization usually occurs before contention. Signal designs may provide a means to contend for the subsequent consecutive timeslots at a higher priority than the first, so that an SDS can have greater confidence it can complete that task.

The fifth technique to gain longer access is to allow a hierarchy of contentions in which contentions can be for whole frames and epochs. This type of hierarchy is illustrated in both Figure 2-5 and Figure 2-6. After a successful contention for a frame of timeslots, the winning device can use the entire period of the frame as it wishes. This is very helpful for providing access to systems with multiple devices that manage access among themselves using another protocol. For example, during a frame a long-term evolution (LTE) system could operate fully within the specification of that LTE system. One device would contend for and win a frame's duration of access and then the LTE system would use its own protocol to arbitrate how devices in the system share the spectrum across the frame. Generally in this approach to access, the contention should also decide on the lane (i.e., the channel bandwidth) that is held for the frame. This hierarchy could also allow a reservation of frames across epochs supporting the periodic access to frames. Similarly, contentions could be designed to allow access to a whole epoch. As with a frame, the system gaining access to an epoch can use the full duration of the epoch, employing its own protocols for internally sharing the reserved lane. Epochs, however, cannot be reserved; each epoch access requires a successful contention.

Figure 4-11 illustrates a timeslot design. The colored periods at the beginning of each timeslot indicate the type of timeslot contention that would be available to contenders. The first timeslot of each frame supports a contention for an entire frame and the first timeslot of an epoch

supports a contention for an entire epoch. The same color coding was used in Figure 4-10 and shows the concept can be used in consolidated signaling as well.




Figure 4-11. Identification of Slots Allowing Contentions for Timeslots, Frames, and Epochs

A device contending for a frame or epoch would first establish the lane used. The contention could be designed to give a device precedence for establishing the lane when contending for a frame or epoch. The technique to enable this precedence would be implemented as a separate priority type phase that precedes the channel selection phase of the contention. Further, it is feasible to restrict the ability to reserve a frame or epoch to certain lanes and not all of them. Through this restriction, some lanes may be set aside so they can continue to support reservations and not have those reservations preempted by SDSs that need frames or epochs to function. Using the highway metaphor, the lanes that support reservations would be equivalent to high occupancy vehicle (HOV) lanes and the devices that require frames and epochs are like trucks that cannot use the HOV lanes.

4.2.4 Spatial Reuse

An advantage of using the SCR paradigm for access is that signaling naturally orchestrates the separation of systems as demonstrated and described in Section 3.3.3.3. That demonstration, however, assumes equal signaling power being used across all the devices. The range of the signals, which is based on the actual propagation conditions on the ground, may vary but it is still symmetric in both directions. Here we consider conditions where signaling range varies from the device mission range and where there is asymmetry. There are two types of asymmetry. In the first the maximum EIRP of the signals is set, but on account of directional antennas it is only observed in the direction in which the antenna is pointed. In all other directions the power is attenuated. In the second, the maximum EIRP varies among signals. Directional antennas can still be part of this scenario. Section 4.2.4.2 and 4.2.4.3 discuss the effects of these two approaches to asymmetry.

4.2.4.1 Choosing Spatial Separation

There is no rule that the range of signals must match the range of RF transmissions. In some cases a highway may perform better when the signal range exceeds the payload transmission range (e.g., to provide greater protection to the receivers) and in others the range can be less (e.g., when modulation and error coding are resilient to interference). However, the greater the signal range the less spatial reuse, and so the highway would have less capacity.

4.2.4.2 Supporting Asymmetric Access Because of Directional Antennas

When asymmetry is a function of directionality alone, and directionality is fixed and is used in both sending and receiving signals, there is no disadvantage to any contention participants. The use of the directional antennas has the beneficial result of increasing the spatial reuse.

There is a disadvantage to using directional contention to resolve access and final pointing solutions among multiple peers, as in a network that is seeking to identify the pairs that would communicate to each other [3]. In this case, although contention signals are directional, a receiving device must listen in all directions and echoes in contention must propagate in all directions to enable interaction among multiple directional contenders contending to reach the same device. Because the contending devices use directional signals, their earlier contention signal may not suppress a peer contender that is outside the main beam. Because these devices listen and echo omnidirectionally, in a subsequent phase, an out-of-beam contender that was not suppressed by the directional devices earlier contention signal may suppress the directional devices contention by its contention signal because the directional contender must listen omnidirectionally. When all devices behave this way and access is not prioritized, the result is greater spatial reuse of the channel. When there is a prioritization phase, this phenomenon can circumvent the desired result of prioritization. Although a directional contender may have a higher priority, a lower priority omnidirectional contender may gain access over the higher priority contender because of this asymmetry. When devices with different capabilities contend for access, those with only omnidirectional signaling and devices with adaptive signaling (i.e. directional in contention signaling and omnidirectional in listening and echoing), they are not equal. The signals of the directional contender will only reach the devices in the direction its antenna is pointed and so not beat the neighboring devices in the other directions in a phase the neighboring devices do not signal. These devices can survive and beat out the higher priority radio in a subsequent phase, demonstrating radios that transmit directional contention signals have a disadvantage in gaining access.

4.2.4.3 Support Asymmetric Access Because of Different Signaling Powers

The effect of power asymmetry depends on the cause of the asymmetry. The critical characteristic that determines fairness is that every pair of devices can hear each other's signals equally well. If the asymmetry is a function of antenna gain alone and that gain would make unamplified peer signals perceptible, then the access would remain fair despite the asymmetric ranges of the signals. It could be unfair when a device has greater transmission range because it uses more power in the transmission and its signals can be heard by others whose signals it cannot hear. This inability to interact in an equivalent manner would result in the unfairness. The device with the longer range signal would have an advantage possibly overriding higher precedence contentions because it did not hear the signals from the device with the higher precedence. Nevertheless, all cases will involve successful contentions and use of the highway.

4.2.5 Changing Lanes

Enabling SDSs to change lanes rapidly creates multiple benefits: quick balancing of traffic on a highway, improved performance of the SDS (e.g., an SDS can avoid a lane occupied by higher priority users with use patterns that conflict with its own), and active lane switching to mitigate adversary monitoring and countermeasures. For multiple device SDSs such as networking radios

that operate on a single channel, this involves moving all the radios in unison to the selected channel. Basic contention does not reveal the winner of a contention so that the multiple radios know which channel a peer has won. Thus, channel switching must be orchestrated by policies and procedures of the SDS itself. For example, a simple policy could be that if the standard lane is not supported after the lane contention, SDSs must contend for the lowest frequency lane that has the correct bandwidth. A peer-to-peer coordination procedure takes longer to execute and cannot be executed from contention to contention. Here the SDSs communicate to each other and decide which lane to choose. The selection could be supported by approaches that monitor the use of all the lanes and a decision that determines which lane would best support the SDS. The timing of the lane change would be communicated, acknowledged, and executed in the specified timeslot. Another alternative is for a single master device that can see all slaves to decide alone and direct slave devices when to switch to a new lane.

A more flexible alternative is available with serial out-of-band signaling if there is tolerance for additional signaling. Signaling slots can be added after the contention phases that allow the winning contender to identify itself, as shown in Figure 4-12. Given q two-minislot phases, there are $2q$ possible signals, and so a binary number from 0 to $2^{2q} - 1$ can be used to identify the system. There are $2q$ signals because these signals do not need to be echoed since only neighboring peers need to hear them so both signals of a phase are used to indicate the identification number (ID). Peers that observe their ID on a particular lane will know which lane to use. Of course there must be enough signaling slots to support conveying a unique ID for each of independently operated SDS allowed to operate on the highway. Due to an anticipated small number of possible ID signaling slots, some type of highway authority will be needed to assign and regulate the use of IDs as opposed to using a permanent ID set at the factory or assigned at an initial registration.




Figure 4-12. Using Additional Signaling Phases to Support Indicating an ID Number

As an alternative to the approach above, rather than set aside a series of signaling slots for indicating an ID, a period can be set aside for the contention winner to send in a modulated signal conveying its individual ID (see Figure 4-13). In this way, IDs can be factory set or set as part of registration, as we anticipate the modulated signal could send a much larger number of bits in the same time.




Figure 4-13. Using a Modulated Signal to Provide an ID

A possible third approach is simply to map an SDS ID to a unique frequency. At the end of the contention the winner uses a single phase to send a signal on that frequency (see Figure 4-14). Peer devices in the SDS listen to the designated frequency⁹ for all ID periods following contentions, and when they hear a signal they understand the lane is being used by a peer in their SDS.




Figure 4-14. Using a Single Signal on a Unique Frequency to Indicate the SDS ID

4.2.6 Effectiveness in Mobile Environments

Mobility becomes an issue when mobile devices win a contention for frames or epochs or reservation for slots or frames. Those contention results may conflict with other results that will merge in space before the use is terminated because of the device's mobility.

4.2.6.1 Mobility and Slot Reservations

An advantage of slot reservation is that there is still a contention for every slot. As a mobile device with a reservation moves it will retain the reservation so long as the reservation priority in the CRS design has precedence. It will immediately give way to a contention that uses a higher precedence. In the cases where two reservation merge, the contention phases will resolve the access. If the SDS loses a reservation, it is still free to go through the process and create another.

4.2.6.2 Mobility with Frame and Epoch Access

Mobility of devices with frame and epoch access may result in collisions. This mobile device with a frame or epoch access is not required to contend for each timeslot. As a result, the mobile device may move into territory where devices are unaware that a frame or epoch of timeslots had been won. The unaware devices will continue to contend without regard for the multi-timeslot result of the mobile device. There is no mechanism for in-channel contention resolution.

There are two approaches to resolve this problem, the first would be to restrict devices that contend for frames or epochs to be non-mobile during the frame or epoch they have won. This does not resolve the opposite issue of mobile devices moving into the territory of the device that has gained access to one of these time periods.

The second resolution to handle mobility is to use out-of-band signaling. In this case the device with a frame or epoch reservation would still contend on a slot by slot basis on the out of band contention channel. The signaling design could be made to ensure the frame or epoch use does

⁹ If the signaling design uses frequency hopping then this signal would likely hop as well from slot to slot and so prevent adversaries from identifying the network.

not receive interference. A rule would be established that this type of device could use a high precedence signal in each slot contention for the duration of the frame or epoch it has won.

4.3 Signaling Physical Layer Design Alternatives

The signaling physical layer design influences the overhead of CRS and its resiliency and in turn the effectiveness of CRS in arbitrating precedence and contention and achieving spatial reuse. Meanwhile features added to signaling to improve its effectiveness may limit the ability of some SDSs to participate on the highway. This section describes the trades in signaling physical layer design.

4.3.1 Signal Requirements

Two requirements apply to signals and signaling minislots.

1. Signals superimpose such that a receiver that hears multiple signals will still detect a signal.
2. Signaling minislots and signals are sized to account for synchronization accuracy, propagation delay for the maximum range, detection time, receive-to-transmit transition time, and frequency blanking interval (if frequency hopping is used) such that the minislot in which a transmitter sends and a receiver detects the signal is unambiguous.

4.3.2 Signal Bandwidth Shape, Duration, and Range

A signal is a pulse of some type that is easy to detect and designed to fit within a specified bandwidth and to achieve a specified range. Assuming the power used in signaling is constant, the narrower the bandwidth the greater the range. However, to achieve a narrow bandwidth, the signal must have a longer duration. The design of a signal shape (i.e. the power varies during the signal transmission) can improve a signal's ability to achieve bandwidth constraints in the shortest duration signal practical. These are the trades inherent in signal design. Generally, the design objective is to find the shortest signal possible that can be reliably detected at a specified range.

4.3.3 Minislot Sizing

The duration of a single signal slot must accommodate the duration of the signal itself, the time to prepare to send or to receive a signal based on the worst case transition, the maximum allowable difference in clock synchronization, the maximum range allowed for the signal, and the time to decide on the next action based on what was observed. The signal duration is a design choice balanced to obtain range and meet bandwidth restrictions. The transition time is a function of what must be done in the transition and the expected processing speed of the contending devices. It includes the processing of the data observed in the previous contention, the decision to transmit or receive next, and the transition from one state to another. Transition time from receiving a signal to transmitting a signal is normally used in the designs because it is typically longest. The time set aside for synchronization difference depends on the desired tolerance to devices being unsynchronized. Additional guard time is necessary to accommodate the maximum distance over which a signal is expected to propagate.

Figure 4-15 and Figure 4-16 illustrate the two scenarios that affect the sizing of the minislots. In both cases the full signal may not arrive within the period when another device is receiving with enough overlap that the signal is detected reliably. The goal of the sizing is to ensure that it detection. Figure 4-15 illustrates an early-arriving signal in the scenario where two radios are very close to each other and are offset by the maximum allowed difference in clock time with the source radio's clock leading the destination clock. The requirement for detection in this case is that

$$t_{\text{observation}} = t_{\text{signal}} - t_{\text{max_clk_offset}} \geq t_{\text{detection}}$$

A delay in transmission from the start of a minislot could fully mitigate this problem.

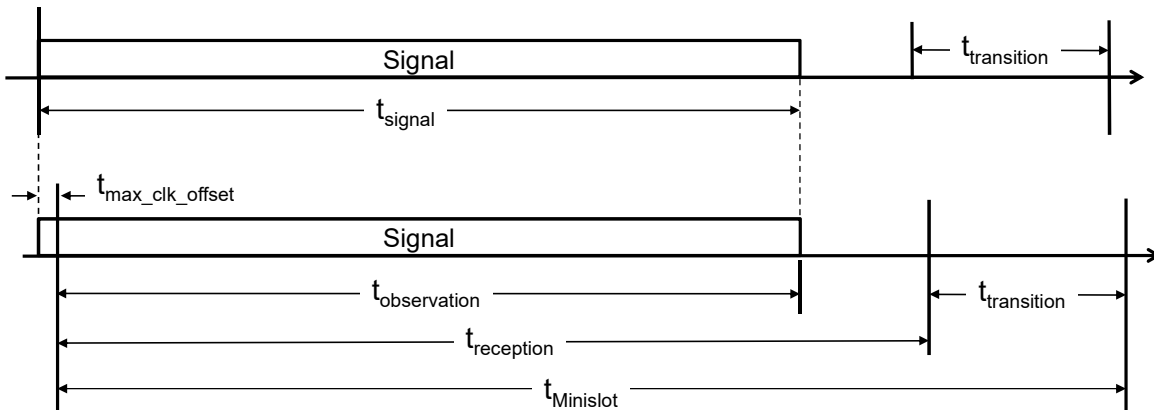


Figure 4-15. Signal Transmission with Maximum Lead Offset and Minimum Propagation

Late arrivals are more likely since the tolerances for clock errors is likely less than the expected propagation durations. Figure 4-16 illustrates a late arriving signal in the scenario where two radios are out of synchronization by the maximum allowed clock offset with the source radio's clock lagging the destination radio's clock and the radios are separated from each other by the maximum propagation range. The requirement for detection in this case is

$$t_{\text{observation}} = t_{\text{signal}} - t_{\text{max_clk_offset}} - t_{\text{max_prop_time}} \geq t_{\text{detection}}$$

A receiving device collects samples of what it receives throughout $t_{\text{reception}}$ and terminates its reception at the end and begins processing what it received on that schedule to determine if there was a signal. The transition time must also accommodate the decision making in response to the signaling observation and the transition to the next state.

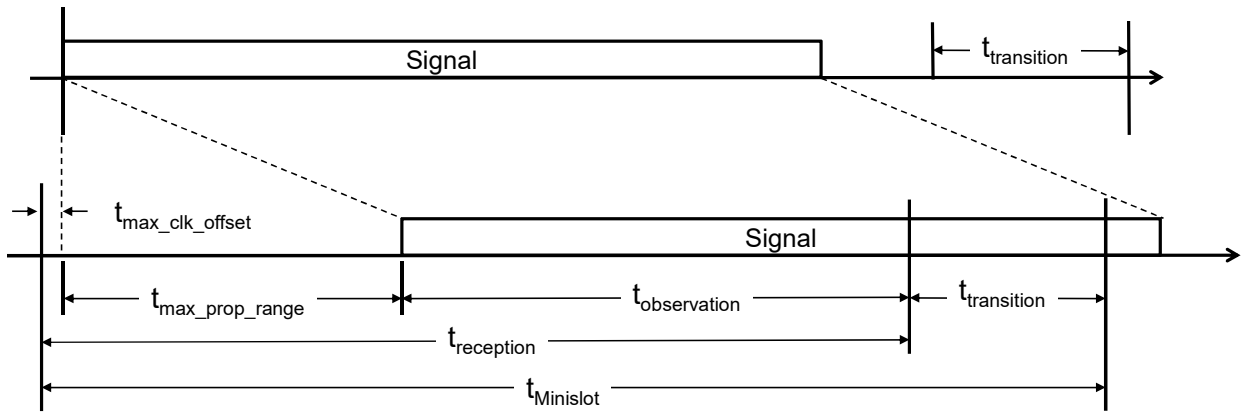


Figure 4-16. Signal Transmission with Maximum Lag Offset and Maximum Propagation

Assuming access to the Global Positioning System (GPS) or the requirement that GPS be accessible to participate on the lane, synchronization can be much less than a microsecond. Allowing a few microseconds may provide an effective guard to accommodate random and temporary losses of GPS access. Longer times may be selected to support the performance of other methods that achieve synchronization. Allowance for propagation time is easily determined. It takes 1 μsec for a signal to propagate 300 meters. Given a maximum range r , the delay that must be accommodated is $\left\lceil \frac{r}{300} \right\rceil \mu\text{sec}$.

4.3.4 Frequency-Hopping Alternatives

Frequency hopping offers several advantages. It makes it more difficult for an adversary to spoof the contention, it isolates signals ensuring that a late arriving signal is not perceived to be present in the wrong minislot, and it prevents devices that are grossly out of synchronization from disrupting the performance of the contention. It also has a cost. There is generally a period of time referred to as a frequency hopping blanking interval in which the frequency transition occurs and in which a radio cannot transmit or receive. Minislot sizing must accommodate the expected duration of the blanking interval of the devices participating on the highway. The blanking interval can dominate the transition time and it can be much greater than the clock tolerance and propagation times.

Designs that use frequency hopping may choose to hop per phase rather than per signal to halve this requirement to pause for the frequency change. This choice would give an adversary the opportunity to use the echo to disrupt the channel access (assuming the adversary can respond fast enough and it is not too far away). This sort of disruption, at best, would thin out the contenders but would have no effect on those devices that won the signaling phase.

Transition between frequencies takes place on an established schedule. It is desirable to design a device to execute the signal processing and decision making described above during the blanking interval.

Assuming a separate processor for this function, the decisions on whether to signal next could be made during other processing such as the changing of frequency. In an approach where the echo

uses the same frequency as the contention signal, the decision to relay a signal would be the dominant activity to precede execution of the transition.

4.4 Performance Measures

Appendix A describes multiple measures that may be used to compare highway designs. They capture what is traded when certain features are added to a design. Table 4-2 provides a brief list of those measures with a cursory description to assist in understanding Section 4.5 on design trades.

Table 4-2. Measures for Comparing Highway Designs

Measure	Definition
Temporal Overhead	The loss of capacity on the highway that results from the duration of the contention design
Spectral Overhead	The loss of spectrum use that results from setting aside spectrum for out-of-band contention
Use Efficiency	The average waste of resources of a timeslot because of the mismatch of bandwidth or duration from that of the SDSs that use the highway
Temporal Effectiveness	The fraction of transmission opportunities that can accommodate a device's use in a time period that matches the longest transmission opportunity
Congestion	The average latency of access as a measure to indicate congestion.
Responsiveness	The speed at which an ongoing use of spectrum can be redirected to a higher priority use which is a functions of access delay and slot delay
Access Delay	A responsiveness measure that is the average time from the end of a contention that garnered a timeslot to the beginning of the timeslot
Slot Delay	A responsiveness measure that is the time from the start of a transmission slot until the start of the next slot contention
Spatial Capacity	The average number of concurrent successful users of a lane per unit of area of the highway.
Spectral Agility	The average time it takes to move all devices of an SDS to a different lane
Asymmetry Risk	A measure of how the ideal performance of a highway is affected by differences in the range of signals or payload transmission among SDSs: typically captured using the measures below
Power Asymmetry	A measure of asymmetry among SDSs based on the asymmetry of their signaling power
Weighted power Asymmetry	A measure of asymmetry among SDS based on the distribution of signaling power differences across the SDSs operating on a highway
Directional Power Asymmetry	The ratio of area covered by a signal sent using a directional means versus the area covered by a signal sent omnidirectionally
Adaptive Directional Asymmetry	A measure of the unfairness that occurs when directional antennas adapt their pointing direction as a function of their use
Contention vs Device Asymmetry	Performance degradation that occurs because the SDS range is different than the signaling range
Multifunction Count	A count of the possible number of combinations of using a highway differentiated by duration, persistence, bandwidth, and volume
Collision Resolution	A measure of the ability of the highway to handle congestion and to ensure its capacity is fully used
Precedence Count	The quantity of differentiation levels that are provide by the signaling design
Reservation Count	A count of lanes and time periods that allow periodic reservations
Resilience	A measure to capture a systems vulnerability to countermeasures. This is a function of the signaling physical layer; does it frequency hop or use directional antennas
Mobility	The duration from the end of a contention to the end of the corresponding use of the lane

4.5 Design Trades

Table 4-3 lists a set of design alternative for highway access design, the capabilities they provide, and the trades that follow with the performance measures. Highway designs and their performance are dictated by the features desired and the accommodations necessary for the performance of the SDSs that operate on the spectrum highway.

Table 4-3. Alternative highway design options and their performance trades

Design Objective	Design Alternatives	Performance Trades
Precedence for Lane Selection – It uses contention to identify which devices select the active lanes. It ensures the active lane is not simply the largest but the one most needed by the user with precedence.	Place a lane precedence phase set before the lane selection phase set .	Improves: multifunction count, precedence count Degrades: temporal overhead, responsiveness (i.e., greater access delay)
Multiple Lane Bandwidth Alternatives – The goal is to allow the merging of a base set of lanes into wider bandwidth lanes	Require the inclusion of a set of lane selection phases. The greater the lane hierarchy, the more phases.	Improves: multifunction count, precedence count Degrades: temporal overhead, responsiveness (i.e., greater access delay)
Quality of Service – The design supports the differentiation of service among peers when their specific uses vary (e.g. transmitting different priority traffic).	Provide multiple precedence phases to support differentiation of service among peers.	Improves: multifunction count, precedence count Degrades: temporal overhead, responsiveness (i.e., greater access delay)
Reservation and Preemption – The design supports devices making periodic reservations.	Divide timeslots into frames and precedence slots are defined for both reservations and preemption.	Improves: multifunction count, precedence count, reservation count Degrades: congestion
Timeslot Duration – This is the base unit of a transmission opportunity. The goal is to provide a transmission opportunity that is sufficiently long to be useful.	Increase the timeslot duration to allow longer duration uses.	Improves: temporal effectiveness, temporal overhead Degrades: use efficiency, slot delay
	Decrease the timeslot duration to reduce wasted transmission time.	Improves: use efficiency, slot delay Degrades: temporal effectiveness, temporal overhead
Timeslot Hierarchy – This is the inclusion of frames and epochs, which enable features such as reservations and longer transmission opportunities.	Establish a frame consisting of a fixed number of timeslots.	Improves: multifunction count, reservation count
	Establish an epoch consisting of a fixed number of frames.	Improves: multifunction count, reservation count
Multi-timeslot Contention – This creates the ability to contend and gain access to either a whole frame or a whole epoch.	Given a timeslot hierarchy with frames, provide a precedence at the beginning of the frame to contend for the whole frame.	Improves: multifunction count, temporal effectiveness, temporal overhead Degrades: use efficiency, reservation count
	Given a timeslot hierarchy with epochs provide a precedence at the beginning of the epoch to contend for the whole epoch.	Improves: multifunction count, temporal effectiveness, temporal overhead Degrades: use efficiency, reservation count
User Differentiation – This provides a means to distinguish users where some users should have precedence over others.	Add contention phases in the priority phase set to distinguish the users. Provide a means to authorize devices to use those phases.	Improves: precedence count Degrades: temporal overhead

Design Objective	Design Alternatives	Performance Trades
Timeslot Merging – The design allows devices to contend for multiple consecutive timeslots so use can extend across timeslot boundaries.	In conjunction with using out-of-band contention, enable a device to reserve several contiguous timeslots on a lane and transmit contiguously across them.	Improves: multifunction count, temporal efficiency, temporal effectiveness Degrades: spectral overhead
Use Differentiation – Provides signaling phases to differentiate between uses.	Add contention phases in the priority phase set to distinguish the users. Provide a means to authorize devices to use those phases based on their function.	Improves: precedence count Degrades: temporal overhead
Network Identification for Rapid Lane Changing – This provides a means for contention winners to convey their SDS identity. Thus, a contention winner can contend on an alternative lane and inform peers in the SDS where to go.	As part of serial out-of-band signaling include a period the lane contention during which a device signals the SDS identity	Improves: spectral agility, temporal effectiveness Degrades: use efficiency
Reuse – The design supports a high density of concurrent users.	Tune the range of signaling to be very close to the range of the SDS that use the spectrum highway.	Improves: spatial capacity, congestion Degrades: contention versus device asymmetry
	Use directional signaling and SDS transmission.	Improves: spatial capacity, congestion Degrades: adaptive directional asymmetry
Long Range Contention – The goal is to support long range devices, for example radars.	Design signaling for long-range propagation so signaling slots are longer.	Improves: multifunction count, temporal effectiveness, access delay Degrades: spatial capacity, temporal overhead
Device Asymmetry – The design supports users with different antenna and use ranges and different signaling powers.	Require all signaling to use the same power and require that it exceed the range of the longest range SDS.	Improves: contention versus device asymmetry, multifunction count Degrades: spatial capacity, temporal overhead, responsiveness (i.e., greater access delay)
	Require all signaling to use the same power but allow SDS ranges that exceed the signaling power (may assume long-range use is rare).	Improves: multifunction count Degrades: contention versus device asymmetry
	Allow heterogeneous powers that match those of the SDSs contending.	Improves: multifunction count Degrades: power asymmetry, weighted power asymmetry, directional power asymmetry
Resilient Signaling – The signaling uses physical layer approaches to mitigate spoofing and jamming.	Frequency hop the signaling.	Improves: resilience Degrades: spectral overhead, temporal overhead
	Spread the signal and make it uniquely identifiable, maybe because of a code that is transmitted.	Improves: resilience Degrades: temporal overhead, spectral overhead

4.6 Conclusions

Section 4 has described an approach for defining access rules for the lanes of spectrum highways. However, developing the functionality that supports the many objectives for a highway requires design choices that generally involve engineering trades in highway performance. Section 4 also described the design techniques to achieve different performance objectives and where the impacts would be felt in other performance measures.

5 Specifying Highway Access Rules

The access rules are specified for each lane using the Policy or Protocol Construct of spectrum consumption modeling [1]. This chapter provides a standard approach using this construct to define access rules based on SCR which will support any of the alternative designs described in Chapter 4. Section 5.1 provides the IEEE 1900.5.2 Standard's description of the Protocol or Policy construct and what features the lane rules use. Section 5.2 defines how the rules of our proposed highway would be specified using this construct. Section 5.3 provides some examples of rule definition for spectrum highway lanes.

5.1 The Protocol or Policy Construct

The Policy or Protocol Construct of IEEE 1900.5.2 allows the specification of behaviors that enable spectrum sharing in the same band. The construct was intended to support the same goals as those for the spectrum highway. Since behaviors could support sharing in any number of ways, the construct was designed to be quite flexible in its use. Rather than predefining techniques to choose from, it simply provides a means to identify a name of a sharing approach and then gives a list of the parameters for its behavior options. In order to use the spectrum authorized with this construct, a device must understand the methods of the named sharing approach and be able to execute the approach to the standard that the performance parameters indicate. This section repeats the description of the data elements found in the standard and then identifies which of these data elements are used for the spectrum highway.

5.1.1 The IEEE 1900.5.2 Description of the Data Elements and their Meaning¹⁰

The basic data structure of a protocol or policy is a name for the protocol or policy, *PorPName*, followed by a set of parameters, *PorPPParameters*. The assumption made is that the name specifies the general behavior and the parameters fill in the details of device performance (e.g., sensitivity, response time), event timing, data structures, and counts. Each policy and protocol name would have an expected number of parameters associated with it that need to be provided for the policy or protocol to be complete. These are defined when the policy or protocol is defined. The definition of protocols and policies is beyond the scope of this standard. Each Parameter in *PorPPParameters* has at least a name and optionally a type and value. In some protocols the appearance of a parameter name is sufficient.

The *PorPIndex* is used to couple a policy or protocol with an underlay mask in a receiver model. Without this association, it has no meaning. The *PorPIndex* is only used in receiver models and it shall be used whenever a protocol or policy is included in a receiver model. The *PorPIndex* is not used in a transmitter model.

¹⁰ This Section is taken directly from the IEEE 1900.5.2 standard [1]. It is Subclause 8.12.2. Current work in the 1900.5.2a revision to the standard intends to support *PorPPParameters* with subparameters of different types.

Table 5-1. SCM Policy or Protocol Construct

Data element	Sub-elements	Types	Enumerated values / unit of measure	Occurrences	Notes
SCMPolicyOrProtocol		Structure			
	PorPName	String		1	
	PorPParameters	Structure		0..1	
	PorPIndex	Integer		0..1	Not used in transmitter models and is required in receiver models.
PorPParameters		Structure			
	Parameter	Structure		1..n	
Parameter		Structure			
	PPPName	String		1	May stand alone
	PPPType	Attribute	STRING, INTEGER, NUMBER, HEX, BINARY, BOOLEAN, DATETIME, DURATION	0..1	If type is used there is a value
	PPPValue	As specified		0..1	If value is used there is a type

This table is in the format used by the standard.

The use of a protocol or policy construct in a transmitter model indicates the specific protocol or policy that the modeled transmitting system is using, can use, or will be required to use depending on the context of the model, i.e., specifying a use, proposing a use, or identifying the conditions for use, respectively. The use of a protocol or policy construct in a receiver model indicates that to use the associated underlay in determining compatibility, the transmitting system shall use the specified protocol or policy.

Protocol or policy parameters may be levels of performance or specific criteria of the particular protocol or policy. For example, a device may be able to reliably sense signals in a band at a particular power level but the criteria of the policy or protocol may be something different, e.g., the sensor can detect signals as low as -100 dBm but the criteria for using the policy or protocol is only -80 dBm. The meaning of parameters in a transmitter model also depends on the context. A device seeking spectrum may convey to a management system that it can execute a protocol or policy and provide parameters that indicate how well it can perform. The management system may then provide policy and provide parameters that indicate the specific performance required to use the spectrum. A model of a transmitter's use of spectrum indicates the parameters of the protocol or policy it is using.

5.1.2 Highway Application

As described, the SCMPolicyOrProtocol construct has three sub-elements: PorPName, PorPPParameters, and PorPIndex. The highway uses the SCMPolicyOrProtocol to identify the conditions for spectrum use. In this role, only the PorPName and PorPPParameters are used; the PorPIndex is not used. A device receiving the SCMPolicyOrProtocol in a highway would first inspect the PorPName and if it recognizes it as a name of an access mechanism it can execute it will look further to the parameters listed in the PorPPParameters structure, verify that they can be executed, and if so, adjust its behavior to use the protocol with the specified parameters. All of the values are strings that allow the parameters to be either text or numbers.¹¹

5.2 Identifying Highway Access Feature and Rule Options

This section defines a standard and unambiguous way to describe the features of the SCR highway access mechanism described in previous chapters. It defines the parameters that can be used to describe the full range of design options for access.

5.2.1 PorPName

The single name for this set of highway access rules is “SCRHighwayAccess.”

5.2.2 PorPPParameters

The PorPPParameters identify the specific rules that are in place for the spectrum highway. Each of the six classes of parameters is included in every rule set. The classes and the rules follow. Although more verbose, in most cases, the parameters are provided in 3-tuples with the name of the parameter followed by the type and then the value to improve readability and to remove ambiguity. Some parameters used in defining the protocols are just a name when they are part of hierarchical data structure or serve as an implied Boolean, (i.e., the presence of the parameter name indicates the feature they identify is being used). It is expected that a protocol design does not contradict itself by defining parameters that are mutually exclusive, (e.g., a per slot signaling design that takes longer than a timeslot).

5.2.2.1 Timeslot Duration and Hierarchy

The first definition specifies the timing of the highway. The same timing is required for all lanes of the same highway and would be seen in every SCMPolicyOrProtocol construct specified for any lane of the highway. There are five parameters and they shall be presented in 3-tuples that include the name of the parameter, its type, and the value. Table 5-2 lists these parameters.

¹¹ The 1900.5.2a revision to this standard intends to add an additional parameter of number type to avoid the issues that may occur converting text to numbers. This document, however, follows the guidance of the current active standard.

Table 5-2. Parameters for the Timing Hierarchy

Parameter Name	Type	Units	Occurrences
TimeReference	UTC time		1
Frame	integer	Integer number of timeslots	1
Timeslot	structure		1..n
ConsolidatedSignalingStart	integer	Integer number of the timeslot in a frame before which consolidated signaling is performed	0..1
ConsolidatedSignalingDuration	double	nanoseconds	0..1
Epoch	integer	Integer number of frames	0..1
Timeslot			
TimeslotDuration	double	nanoseconds	1
TimeslotRange	double	meters	1

The TimeReference, Frame, and Timeslot are required values. TimeReference indicates the start of the first timeslot in the first frame of the first epoch. The TimeReference includes year, month, day, hour, minute, and second with an hour and minute displacement from the UTC time zone. The format for start time is $\langle YYYY, MM, DD, hh, mm, ss.s, \pm hh_o, mm_o \rangle$. *YYYY* is the Gregorian year, *MM* is the integer number of the month in the year, *DD* is the day of the month, *hh* and *hh_o* are integer numbers of hours on a 24 hour clock, *mm* and *mm_o* are an integer number of minutes less than 60, and *ss.s* is a decimal number of seconds less than 60. The combination $\langle hh, mm, ss.s \rangle$ indicates the time of day and the combination $\langle \pm hh_o, mm_o \rangle$ indicates the time zone. All values are provided. As an example, the string “2018,12,01,24,00,0,-05,00” is equivalent to saying “at 2400 hours on the 1st of December 2018, Eastern Time.”

Frame indicates the number of timeslots in a frame. If the number is greater than 1, then there are multiple timeslots. If the Frame value is followed by one Timeslot structure, then all Timeslots have the characteristics it describes. Otherwise Frame must be followed by a quantity of Timeslot structures equal to the Frame value. These Timeslots can have different TimeslotDuration and Range values.

TimeslotDuration is the length of a timeslot beginning to end. Range indicates the propagation distance that is used for determining a guard time. Given a range *r*, a transmission on a timeslot shall end $\left\lceil \frac{r}{300} \right\rceil$ μsec before the beginning of the next timeslot. Each timeslot begins immediately after the previous timeslot. The one exception is the timeslot that follows consolidated signaling, which starts ConsolidatedSignalingDuration nanoseconds after the previous timeslot.

Epoch is an optional parameter that indicates that frames are organized into epochs. The Epoch value is the integer number of frames that make up an epoch.

The ConsolidatedSignalingStart and ConsolidatedSignalingDuration are optional parameters used to set aside time in a frame for consolidated signaling when it is used. The period for consolidated signaling is placed before the timeslot indicated by Consolidated Signaling Start. (Numbering of timeslots begins with 1.) The duration of the frame includes the ConsolidatedSignalingDuration plus a Frame quantity of TimeslotDuration.

5.2.2.2 Contention Method

There are four contention methods: in-band-signaling per slot, consolidated in-band-signaling, serial out-of-band signaling, and concurrent out-of-band signaling. The purpose of calling out the contention method is to identify its effect on payload and response times, which can be a determining factor as to whether a device can participate in the highway. For example, in-band-signaling per slot affects the duration of payload transmissions. Table 5-3 lists the values.

Table 5-3. Parameters for the Contention Method

Parameter Name	Type	Units	Occurrences
ContentionMethod	enumerated string	PERSLOT CONSOLIDATED SERIAL CONCURRENT	1
PerSlotDuration	double	nanoseconds	0..1
OutOfBandOffset	double	nanoseconds	0..1
Order	integer	Integer ≥ 0	0..1

The ContentionMethod is an enumerated value with four options that shall appear in an SCRHighwayAccess SCMPolicyOrProtocol.

If the method is PERSLOT then a duration for the contention shall be provided. This is the time from the beginning of a timeslot in which contention occurs. The payload use starts PerSlotDuration after the beginning of each timeslot.

If either the SERIAL or CONCURRENT method is specified then an OutOfBandOffset shall be provided. This is the temporal offset from the beginning of a timeslot when the first signaling phase for a timeslot begins.

The Order is a required element when the Serial contention method is specified; it indicates whether the Priority Phase Set and Contention Phase Set follow the Lane Selection Phase Set or the contention for another lane. It identifies the number of contentions that separate it from the Lane Selection Phase Set. Referring to Figure 4-6, in this example, Lane 1 would have order 0, Lane 2 would have order 1, Lane 3 would have order 2, and Lane 4 would have order 3.

5.2.2.3 Lane Use Precedence Phase Set

The optional first phase set of a contention is the Lane Use Precedence phase set. It provides a means for certain users to have precedence in deciding the lanes that are used in a timeslot. This phase set is used to prevent devices that require wider bandwidth lanes from always defining the lane size simply because wider bandwidth lanes preempt the narrowband lanes they overlay. This can be wasteful when higher priority users with narrower lane requirements ultimately have precedence in the contention. In these cases the higher priority users should have precedence in establishing the active lanes. Table 5-4 lists the parameters that define the Lane Use Precedence Phase Set.

Table 5-4. Parameters for the Lane Use Precedence

Parameter Name	Type	Units	Occurrences
LaneUsePrecedencePhases	integer	A positive integer number of phases	1
Override	implicit Boolean	No value, use of parameter name means TRUE	0..1
LPChannel	string	Channel Name	0..1

LaneUsePrecedencePhases is a required field of the SCRHighwayAccess protocol. A value of 0 indicates it is not used otherwise it is a positive number. Given an integer greater than 0, there are $2^{\text{LaneUsePrecedencePhases}}$ precedence levels.

If LaneUsePrecedencePhases is a positive integer then an entry for Override may be provided. Override means there are $(2^{\text{LaneUsePrecedencePhases}} - 1)$ precedence levels for routine lane selection and the highest precedence is set aside to allow a critical device to preempt the selection.

LPChannel is the name of the contention channel that is used for this part of signaling, and shall be provided if LaneUsePrecedencePhases is a positive integer. The definition of the active contention channels that affect the access to a lane are provided at the end of the SCRHighwayAccess description (see Section 5.2.2.8).

5.2.2.3.1 Rules for Lane Use Precedence Phase Sets

Prior to the contention, each contending device determines its precedence and in which phases it is authorized to signal. A device shall use the precedence level it is granted at configuration time and if designated an override user it may use the override precedence when the conditions for override are present. Conditions for override include:

- The highway uses consolidated contentions per frame and supports timeslot reservations. A device that has a timeslot reservation shall use the Lane Use Precedence override.
- A device enters each phase of a Lane Use Precedence Phase Set as either a lane use contender or non-contender.
 - If a lane use contender and the device's precedence indicates a signal in the current phase, that device shall signal in the signaling slot of the phase. It survives that phase and moves onto the next one as a contender.
 - If a lane use contender and the device's precedence indicates it should not signal in the current phase, the device shall listen for the first signal of the phase and if it hears one, it shall echo the signal and mark itself as a non-contender.
 - If a lane use contender and the device's precedence indicates it should not signal in the current phase, and it did not detect the first signal in the phase, it shall then listen for the echo signal. If it detects an echo signal it shall become a non-contender and if it does not detect an echo it shall remain a contender.
 - If a device is a lane use non-contender, it shall not signal in the current phase but shall listen for the first signal of the phase and, if it hears one, it shall echo the signal in the same phase.

- A device gets to contend in the lane selection phase set if it survives all of the lane precedence phases as a lane use contender

5.2.2.4 Lane Selection Phase Set

The Lane Selection Phase Set identifies which lanes are active on a highway. The parameters specified for a lane indicate the signaling combination that a device shall use to specify the lane as active, identifies the signaling outcomes that indicate the lane is active and that the device can contend for access, and the outcomes that indicate the lane is active but the device may not contend.

Table 5-5. Parameters for the Lane Selection Phase Set

Parameter Name	Type	Units	Occurrences
LaneSelectionPhases	integer	A positive integer number of phases	1
LaneSelectionPause	double	nanoseconds	0..1
LaneSignalingSequence	binary	Digits correspond to phases in LaneSelectionPhases, the least significant digit the last	0..1
LaneOverride	binary	Digits correspond to phases in LaneSelectionPhases, the least significant digit the last	0..n
ContentionEligibleOutcome	binary	Digits correspond to phases in LaneSelectionPhases, the least significant digit the last	0..n
ActiveOnlyOutcome	binary	Digits correspond to phases in LaneSelectionPhases, the least significant digit the last	0..n
LSChannel	string	Channel name	0..1

LaneSelectionPhases is a required field of the SCRHighwayAccess protocol. A value of 0 indicates it is not used and is a legitimate entry if there are no lanes that overlap others on the highway. Otherwise, a positive integer indicates the number of phases that are used to identify which lanes are active.

The LaneSelectionPause is a period of time that precedes the first phase to allow a device to prepare for the signaling in the Lane Selection Phase Set. It is not used if the LaneSelectionPhases value is 0.

The contention rules for lane selection differ from those for other precedence phases since signaling combinations indicate active lanes and are not used to arbitrate access. There is no strict hierarchy of phases where losing an earlier phase as a contender transitions a user to non-contender status. Rather there is an explicit description of the behavior in signaling and an identification of outcomes that indicate a particular lane is active and whether a device can subsequently contend for access to the lane. There are four parameters that describe this behavior: LaneSignalingSequence, LaneOverride, ContentionEligibleOutcome, and ActiveOnlyOutcome. They are not provided when the LaneSelectionPhases value is 0.

LaneSignalingSequence, LaneOverride, ContentionEligibleOutcome, and ActiveOnlyOutcome are all binary values. The least significant digit corresponds to the last phase of

LaneSelectionPhases and each preceding more significant digit in order to the phases that precede it. A one value corresponds to a signal or echo occurring in that phase and a 0 corresponds to neither a signal nor an echo occurring in the phase.

LaneSignalingSequence indicates the sequence of signals that a device trying to establish this lane as active would use to indicate that choice. This value is a binary number as described above. In this case, a 1 in the digit position corresponding to a phase indicates the device should signal in that phase and a 0 indicates it should not signal in that phase.

LaneOverride identifies signaling that, if observed, indicates this lane cannot be reserved and that this device should cease efforts to establish this lane as active. This value is a binary number as described above. A 1 digit in a position that corresponds to a phase indicates a phase when a signal or echo is received and a 0 indicates a phase where a signal or echo is not received. If signals are received in the pattern of one of these lane override sequences during the lane selection phases, in phases prior to this radio sending its first signal according to the LaneSignalingSequence, then the contention is lost and the device should not signal in the lane selection phases.

ContentionEligibleOutcome identifies the signaling sequences that if observed allow this device to contend for this lane. This value is a binary number as described above. A 1 in a digit's place corresponding to a phase indicates a signal or echo is sent or received in that phase and a 0 in that place indicates a signal or echo is neither transmitted nor received in that phase. If the observation of the signaling matches one of these sequences then this device may contend for access to this lane.

ActiveOnlyOutcome identifies signaling sequences that if observed indicate the lane is active but the device may not contend to use the channel. This value is a binary number as described above. A 1 in a digit's place corresponding to a phase indicates a signal or echo is sent or received in that phase and a 0 in that place indicates a signal or echo is neither transmitted nor received in that phase. If this signaling sequence is observed, the device may not contend for access to this lane but is still expected to participate in the subsequent phase sets as a non-contender for the lane.

LSChannel is the name of the contention channel that is used for this part of signaling and shall be provided if LaneSelectionPhases is a positive integer. The definition of the active contention channels that affect the access to a lane are provided at the end of the SCRHighwayAccess description (see Section 5.2.2.8).

5.2.2.4.1 Rules for Lane Selection Phase Sets

Each active device enters the Lane Selection Phase Set as either a contender to establish an active lane or as a non-contender seeking to determine what lane is active. A device shall first determine if it may contend for the lane that it wants to access. Generally, a device may try to make any lane active with the following exceptions which seek to prevent a wideband lane contention when a long duration access is in effect on a narrowband lane it covers.

- In a highway using per-slot signaling where a lane overlaps other lanes in spectrum and those other lanes support frame contention, a device shall only contend to make the broadband lane active in contentions that support frame contention and in the contentions

for the timeslots between those contentions when the winning contention for the first timeslot of a frame was for the wideband lane.

- In a highway using per-slot or consolidated signaling where a lane overlaps other lanes in spectrum and those other lanes support epoch contention, a device shall only contend to make the broadband lane active in contentions that support epoch contention and in the contentions for the timeslots of the first frame when the winning contention for the first timeslot of an epoch was for the wideband lane. If frame contentions are not allowed then this rule applies to all timeslots of an epoch and if frame contentions are allowed then the previous rule must also be met for every timeslot within each frame.

A contender identifies the signaling sequence that is associated with the lane's selection and prepares to send the signals it indicates. Non-contenders shall echo contention signals. All devices record their observation whether a signal was received or sent in a phase and whether the first signal received or sent was a contention or an echo signal. Thus the outcome for each phase is contention, echo, or no signal. The collection of outcomes of all phases indicates whether this lane is active and whether the observing device may contend to use the channel.

- If a device is a lane-selection contender and the LaneSignalingSequence indicates a signal shall be sent in this phase, the device shall transmit a contention signal in the phase, record the phase result as contention, and remain a lane-selection contender going into the next phase if there is one.
- If a device is a lane-selection contender and the LaneSignalingSequence indicates a signal shall not be sent in this phase, the device shall listen for a contention signal, and if a contention signal is received it shall transmit an echo signal and then record the phase result as contention. The device shall then compare the sequence of phases in which signals were heard, whether contention or echo, and compare this to the LaneOverride sequences. If one of these matches the current observed result, the device shall change its state to a lane selection non-contender. Otherwise, it shall remain a lane-selection contender going into the next phase if there is one.
- If a device is a lane-selection contender and the LaneSignalingSequence indicates a signal shall not be sent in this phase, the device shall listen for a contention signal, and if a contention signal is not received it shall listen for an echo signal. If an echo signal is received it shall record the phase result as echo. The device shall then compare the sequence of phases in which signals were heard, contention or echo, and compare this to the LaneOverride sequences. If one of these matches the current observed result, the device shall change its state to a lane selection non-contender. Otherwise, it shall remain a lane-selection contender going into the next phase if there is one.
- If a device is a lane-selection contender and the LaneSignalingSequence indicates a signal shall not be sent in this phase, the device shall listen for a contention signal, and if a contention signal is not received it shall listen for an echo signal. If an echo signal is not received it shall record the phase result as no signal. It shall remain a lane-selection contender going into the next phase if there is one.

- If a device is a lane-selection non-contender, the device shall listen for a contention signal, and if a contention signal is received it shall transmit an echo signal and then record the phase result as contention.
- If a device is a lane-selection non-contender, the device shall listen for a contention signal, and if a contention signal is not received it shall listen for an echo signal and if an echo signal is received it shall record the phase result as echo. If no echo is received then the device shall record the phase result as no signal.
- At the conclusion of all phases, the device shall compare the signaling results to specified outcomes in the following order. Its status as a lane selection contender or lane-selection non-contender will have no effect.
 - The device shall convert the result to 1s and 0s converting any phase result with a contention signal to 1 and those with either an echo or no signal to 0. If the resulting sequence exactly matches any sequence identified as a ContentionEligibleOutcome then the lane shall be identified as active and this device may contend for its use. It need not make other checks.
 - The device shall convert the result to 1s and 0s converting any phase result with a contention or echo signal to 1 and those with no signal to 0. If the resulting sequence exactly matches any sequence identified as a ContentionEligibleOutcome then the lane shall be identified as active and this device may contend for its use. It need not make other checks.
 - The device shall convert the result to 1s and 0s converting any phase result with a contention or echo signal to 1 and those with no signal to 0. If the resulting sequence exactly matches any sequence identified as a ActiveOnlyOutcome, then the lane shall be identified as active, but the device may not contend for access. If the device is an eligible destination of a transmission on the lane it shall participate in the priority and contention phase sets.
- If a device has determined by the checks above that it may contend for access but it is not a contender, it shall assess whether an SDS peer may request it to participate on this lane and if so, it shall participate in the priority and contention phase sets.
- If a device has determined by the checks above that it may contend for access and it is a contender, then it shall enter the next phase set as a contender and participate in the signaling throughout the remaining phase sets regardless of whether it survives as a contender.

5.2.2.5 Priority Phase Set

The Priority Phase Set determines among contending devices for a lane which device in a geographical region has precedence over others. Reasons for precedence are user importance, use urgency, and QoS differentiation described in Section 3.3.3.4. Authorizations for precedence are granted to devices based on their function and their user and are not part of the lane Priority Phase Set definition. Figure 5-1 illustrates a priority phase example with all features, QoS, user/use, and reservation precedence. The order of priority phase segments, when used is illustrated. The reservation slot shall always precede the User/Use and QoS phases. The

User/Use phases shall always precede the QoS phases. Table 5-6 lists the parameters used to define the features of the precedence phase for a lane.

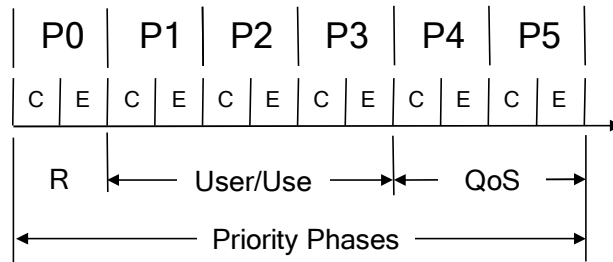


Figure 5-1. Example Priority Phase Set

Table 5-6. Parameters for the Priority Phase Set

Parameter Name	Type	Units	Occurrences
PriorityPhases	integer	A integer number of phases	1
PriorityPause	double	nanoseconds	1
QoSPhases	integer	A positive integer	0..1
UserUsePhases	integer	A positive integer	0..1
TimeslotReservations	implicit Boolean	No value; use of parameter name means TRUE	0..1
FrameContention	implicit Boolean	No value; use of parameter name means TRUE	0..1
EpochContention	implicit Boolean	No value; use of parameter name means TRUE	0..1
FrameReservation	implicit Boolean	No value; use of parameter name means TRUE	0..1
ReservationPrecedence	implicit Boolean	No value; use of parameter name means TRUE	0..1
PChannel	string	Channel name	0..1

PriorityPhases is a required field of the SCRHighwayAccess protocol. A value of 0 indicates it is not used; otherwise it is a positive number indicating the number of phases used in the Priority Phase.

PriorityPause is a period of time that precedes the first phase to allow a device to prepare for the signaling in the Priority Phase Set.

QoSPhases is an optional field and when used indicates the number of PriorityPhases that are set aside to differentiate QoS. There are $2^{QoSPhases}$ different QoS levels. The use of these phases is governed internally to the SDS of which the device is a part.

UserUsePhases is an optional field and when used indicates the number of PriorityPhases that are set aside to differentiate users or uses. There are $2^{UserUsePhases}$ levels of user/use precedence.

TimeslotReservations is an optional field and when used indicates PriorityPhases is set aside to support reservations. At most one phase is set aside for reservations.

FrameContention is an optional field and when used indicates that this lane allows contention for whole frames. The priority for a frame contention is the second highest priority that is able to be signaled using the user/use phases, $(2^{UserUsePhases} - 2)$.

EpochContention is an optional field and when used indicates that this lane allows contentions for whole epochs. The priority for an epoch contention is the highest priority that is able to be signaled using the user/use phases, $(2^{UserUsePhases} - 1)$.

FrameReservation is an optional field and when used indicates a priority phase is set aside to support reservations. At most one phase is set aside for reservations. If the lane allows both timeslot and frame reservations, then they both share the same phase for securing the reservation. Only one type of reservation can be active at a time.

ReservationPrecedence is an optional field and when used indicates a device shall not contend for a frame on a lane where there is an active timeslot reservation and shall not contend on an epoch when there is either an active timeslot or frame reservation. When not used, the design allows frame contention to preempt timeslot reservations and epoch contentions to preempt both timeslot and frame contentions.

PChannel is the name of the contention channel that is used for this part of signaling and shall be provided if PriorityPhases is a positive integer. The definition of the active contention channels that affect the access to a lane is provided at the end of the SCRHighwayAccess description (see Section 5.2.2.8).

In a proper priority phase set design the number of phases indicated by QoSPhases, UserUsePhases, and the activation of a reservation shall equal the number of PriorityPhases.

$$PriorityPhases = QoSPhases + UserUsePhases + \\ (if TimeSlotReservation or FrameReservations, 1, 0)$$

The rules for a device's use of the priority phases are described below, and privileges are granted to the devices separate from the definition of the highway.

5.2.2.5.1 Rules for Priority Phase Sets

Each active device enters the Priority Phase Set of an active lane as either a contender or non-contender to use the lane. The previous results place no restriction on a device's use of the lane. It is only necessary that the device can physically use some spectral portion of the lane.

At the start of the phase the device determines which priority phase signals it is authorized to use to access the channel. The authorization of priority follows a strict precedence signaling regimen.

- If a device intends to contend for access to a lane for a timeslot and the lane supports either frame or epoch contention, the device shall have participated in the previous contention for these long-period contentions. If the current timeslot is part of an accessed frame or epoch, it shall not contend.
- If the ReservationPrecedence field is present for the lane design, a device wanting to contend for a frame shall wait until after a frame in which there is no timeslot reservation and a device wanting to contend for an epoch shall wait for an epoch with no frame or timeslot reservation.

- If the ReservationPrecedence field is *not* present for the lane design that supports timeslot reservations and frame contention or these in addition to either or both frame reservations and epoch contention, then it is possible that devices that contend either for a frame or an epoch would preempt any timeslot reservation and devices that contend for an epoch would preempt any frame reservation. Devices with reservations shall monitor all contentions for frames and epochs and forgo their reservations during those frames or epochs that are made active as a single transmission opportunity by those contentions. They may immediately restart their reservation in subsequent frames and epochs that are not used as frame or epoch transmission opportunities.
- If a device is a contender and its precedence authorizes it to signal in the current phase, it shall transmit a contention signal and remain a contender.
- If the device is a contender and its precedence does not authorize a signal in the current phase, it shall listen for a contention signal and if one is received the device shall transmit an echo signal and transition to being a non-contender.
- If the device is a contender and its precedence does not authorize a signal in the current phase, it shall listen for a contention signal. If one is not received the device shall then listen for an echo signal and if one is received it shall transition to being a non-contender.
- If the device is a contender and its precedence does not authorize a signal in the current phase, it shall listen for a contention signal. If one is not received the device shall then listen for an echo signal and if one is not received it shall remain a contender.
- If the device is a non-contender it shall listen for a contention signal in the current phase and if one is received it shall transmit an echo signal.
- If a contending device survives all contention phases and leaves the last priority phase as a contender, then it shall participate in the contention phases.

5.2.2.6 Contention Phase Set

The contention phase set resolves which devices among those that survived the priority phase set ultimately gain access. The expectation is that devices of equivalent precedence will contend with each other. The design of the contention phases is driven by three factors: the anticipated quantity of devices that will likely contend with each other, the desired resolution confidence, and the desire to constrain the overhead of the contention resolution signaling. Section 3.3.3.2 provides a detailed explanation of contention resolution signaling design. Given this design approach and the proven designs it describes, the specification of the contention phase set can be reduced to the three parameters listed in Table 5-7.

Table 5-7. Parameters for the Contention Phase Set

Parameter Name	Type	Units	Occurrences
ContentionPhases	integer	Positive integer	1
ContentionPause	double	nanoseconds	1
DesignDensity	string	“50”, “200”, “1000”	1
CChannel	string	Channel name	1

ContentionPhases identifies the quantity of phases used in the contention phase set.

ContentionPause is a period of time that precedes the first phase to allow a device to prepare for the signaling in the Contention Phase Set.

DesignDensity identifies the targeted quantity of contenders in the CRS design. The possible choices are 50, 200, and 1000. In CRS, each phase has a specific signaling probability that a contender uses to determine whether it will signal. Table 5-8 lists the specific probabilities used for each signaling phase based on the phase number and the DesignDensity

CChannel is the name of the contention channel used for the contention phase set. The definition of the active contention channels that affect the access to a lane are provided at the end of the SCRHighwayAccess description (see Section 5.2.2.8).

Table 5-8. Contention Phase Set Signaling Probabilities

Phase	DesignDensity		
	50	200	1000
0	0.06	0.03	0.01
1	0.26	0.19	0.10
2	0.33	0.31	0.22
3	0.41	0.40	0.36
4	0.45	0.45	0.43
5	0.48	0.48	0.46
6	0.49	0.49	0.48
7	0.49	0.49	0.49
8	0.50	0.50	0.50
...	0.50	0.50	0.50

5.2.2.6.1 Rules for Contention Phase Sets

Each device that has the potential to use the lane either as a transmitter or a receiver shall participate in the contention phases. Those that have a use for the lane and survived the Priority Phase Set enter the Contention Phase Set as contenders. All others are non-contenders.

Prior to the start of each contention phase, contenders draw a uniformly distributed random number between 0 and 1 and if that number is less than the signaling probability provided for that phase the contender shall transmit a contention signal. Contention phase probabilities are listed in Table 5-8, where phase number and DesignDensity identify the cell with the probability. Given this determination the following rules define the behavior of devices in the contention phase.

- If a device is a contender and has determined it should send a contention signal in the current phase it shall transmit a contention signal and remain a contender.
- If the device is a contender and has determined it is not to send a contention signal in the current phase, it shall listen for a contention signal. If one is received the device shall transmit an echo signal and transition to being a non-contender.
- If the device is a contender and has determined it is not to send a contention signal in the current phase it shall listen for a contention signal. If one is not received the device shall then listen for an echo signal and if one is received it shall transition to being a non-contender.

- If the device is a contender and has determined it is not to send a contention signal in the current phase it shall listen for a contention signal. If one is not received the device shall then listen for an echo signal and if one is not received it shall remain a contender.
- If the device is a non-contender it shall listen for a contention signal in the current phase and if one is received it shall transmit an echo signal.
- If a contending device survives all contention phases as a contender then it has permission to use the lane for the time period contended for.

5.2.2.7 Contender ID Phase Set

The Contender ID Phase Set is an optional phase set used with serial out-of-band signaling only. It allows a device that survives the contention to identify the SDS of which it is part so that other devices within the SDS can move to and operate in the lane. It allows SDSs to move among lanes timeslot to timeslot. As described in Section 4.2.5, there are three approaches to conveying SDS IDs. Table 5-9 lists the parameters that govern the behavior of this phase set. Use of these parameters, indicate the approach and the approach’s parameters.

Table 5-9. Parameters for the Contender ID Phase Set

Parameter Name	Type	Units	Occurrences
ContenderIDMethod	enumerated string	“Binary”, “ModulatedID”, “Frequency”	0..1
ContenderIDPhases	integer	Positive Integer	0..1
IDPause	double	nanoseconds	0..1
IDWindow	double	nanoseconds	0..1
IDChannel	string	Channel name	0..1

The ContenderIDMethod is an optional element that serves two purposes when used. First it indicates that there is a Contender ID Phase Set and second it identifies the method used to transmit the ID. There are three choices: “Binary” meaning multiple signaling phases are used to send a binary number associated with the ID; “ModulatedID” meaning a time period is provided for the contention winner to send a modulated signal that contains the ID; and “Frequency” which means a single phase is used but the frequency used in the signaling is uniquely associated with the SDS. Peers on the SDS listen for signaling on that frequency.

ContenderIDPhases is used whenever the ContenderIDMethod is Binary or Frequency and indicates the number phases used in the Contender ID Phase Set.

IDPause is a period of time that precedes the first phase to allow a device to prepare for the signaling in the Contender ID Phase Set.

IDWindow is used whenever the ContenderIDMethod is ModulatedID. It indicates the duration of time set aside for the transmission of the SDS ID.

IDChannel is the name of the contention channel used for the Contender ID Phase Set. The definition of the active contention channels that affect the access to a lane are provided at the end of the SCRHighwayAccess description (see Section 5.2.2.8).

Since there are three methods for conveying the SDS ID there are three sets of rules for this phase set.

5.2.2.7.1 Rules for Contender ID Phase Sets When the Binary Method is Specified

Contending devices transmit a series of signals to provide a binary SDS ID consisting of bits conveyed by the use of each signaling slot. Both the contention and echo signaling slots are used independently. The presence of a signal shall indicate a 1 and the absence of a signal shall indicate a 0. When there are *ContenderIDPhases* phases there are $(2 \times \textit{ContenderIDPhases})$ bits and so $2^{2 \times \textit{ContenderIDPhases}}$ unique SDS IDs are possible. Figure 4-12 illustrates an example of the design of this phase set and the use of the contention phases to identify two bits each.

- The contending device knows its SDS ID and prepares to signal in the Contender ID Phases by identifying in which signal slots it will transmit a signal. It shall have a binary ID and shall transmit in the signaling slots the most significant bit first.
- At the start of a signal slot, whether the signal is a contention signal or an echo of the specified phase, a contending device shall assess whether there is a 1 or 0 in the SDS ID bit position that matches the current signaling slot position. If it is a 1 the contending device shall transmit a signal. If it is a 0, the contending device shall not transmit a signal.
- A non-contending device shall only listen for signals. It shall record any signal it receives in a signaling slot as a 1 in the corresponding bit position of the binary number and it shall record a 0 if no signal is received in that signaling slot.
- At the end of all contender ID phases the non-contender shall assess the value of the SDS ID indicated by the signaling, and if the binary number matches its own ID, it shall identify the lane as being used by its SDS.

5.2.2.7.2 Rules for Contender ID Phase Sets When the ModulatedID Method is Specified

- A contending device shall send a modulated version of its SDS ID within the specified IDWindow. Figure 4-13 illustrates that the ID window is simply a period of time long enough to support the transmission of the ID.
- A non-contending devices shall listen for a modulated signal during the IDWindow and if the demodulated output is an SDS ID that matches its own, it shall identify the lane as being used by its SDS.

5.2.2.7.3 Rules for Contender ID Phase Sets When the Frequency Method is Specified

- Contending devices shall send a contention signal in the contention phase on a designated frequency that is associated with its SDS ID.
- Non-contending devices shall listen on the designated frequency associated with its SDS ID and if it hears a signal it shall send an echo signal on the same frequency and it shall identify the lane as being used by its SDS.

5.2.2.7.4 Common Rules when Using a Contender ID Phase Set

Highways that use serial signaling support rapid movement by devices across lanes and also the aggregation of adjacent lanes through the successful contention for each of the adjacent lanes. The following rules support those features.

- Contending and non-contending devices shall participate in the contention signaling of all lanes and shall assess whether their SDS ID is used in all lanes.
- If a contending devices observes its SDS ID prior to its own attempt to contend for a lane, it shall transition from being a contender to a non-contender for the duration of the timeslot.
- If a contending device loses a contention for a lane, the winner has a different SDS ID, and there are subsequent contentions for other lanes for the same timeslot, the device may transition back to being a contender and contend for one of those lanes.
- If the highway design does not explicitly call out wideband lanes that overlap others and allows combining lanes in contention, then:
 - A contending device that has won a lane may contend for the adjacent lane in order to aggregate lanes for wideband use.
 - A non-contending device shall assess the SDS ID outcome for all lanes and if its SDS ID was used for multiple adjacent lanes, it shall listen for the broadband signal that would be associated with the bandwidth of the combined lanes.

5.2.2.8 Contention Channels

There are many design options for contention channels and the specific signals they use. This is likely to be an ongoing area of development and so the definition is vague on details such as specific frequencies, signal shapes, and signal content. The data elements provide a means for these types of definitions to be continuously updated as signaling characteristic are developed and improved. The data elements provide a general timing description that focuses on sequencing and whether the signaling is in-band or out-of-band as well as some indication of orthogonality to other contention channels used in the same highway.

Multiple contention channels will support a multilane highway. Each lane will participate on a subset of the channels used by the highway. The lane definition identifies the channels used by a lane, so the names specified in LPChannel, LSChannel, PChannel, CChannel, and IDChannel are provided at this point in the SCRHighwayAccess protocol. They may be repeated in other lane definitions.

Multiple contention channels may be defined for a lane. Each defined Channel will have the parameters listed in Table 5-10.

Table 5-10. Parameters for Contention Channel Definition

Parameter Name	Type	Units	Occurrences
Channel	structure		1..n
Channel			
ChannelName	string	Name of the channel	1
PhaseDuration	double	nanoseconds	1
ChannelParameters	structure		1..n
ChannelParameters			
ParameterName	string	A parameter name	1
ParameterType	string	STRING, INTEGER, NUMBER, BINARY, HEX, BOOLEAN, DATETIME, DURATION	0..1
ParameterValue	As defined		0..1

ChannelName is the name of the channel to which values of LPChannel, LSChannel, PChannel, CChannel, and IDChannel are matched.

PhaseDuration is the total time allotted to a phase. It is the time from the start of the phase until the end of the phase. In multiphase phase sets the duration of the phase set is pause time that precedes the phases plus the length of time for the specified number of PhaseDurations to occur end-to-end.

ChannelParameters are a list of parameters that define a channel. They follow the name-type-value approach to parameters with the repeated 3-tuple ParameterName, ParameterType and ParameterValue. This matches the 3-tuple of the PorPParameters.

5.2.2.9 Range Symmetry

The last set of parameters to define signaling are those that govern power level and range symmetry. Signaling will use a different modulation than that used by SDSs, and different frequencies when using out-of-band signaling. Signals are very likely to have a propagation advantage because they will have narrower bandwidth and support all power being applied in that bandwidth. These parameters define levels of performance expected of the devices that share the highway so that spectral resources are not underused spatially and the access to the spectral resources of the highway follows the intent of the design and is not unfair. Table 5-11 lists the parameters.

Table 5-11. Parameters for Power Asymmetry

Parameter Name	Value	Units	Occurrences
NominalPower	double	dBW	1
MaxPowerVariance	double	dB	1
NominalAntennaGain	double	dB _i	1
MaxAntennaGain	double	dB _i	1
TimingRange	double	meters	
SquareLawRange	double	meters	1
TerrestrialRange	double	meters	1

UseAsymmetry	double	fraction	1
--------------	--------	----------	---

NominalPower is the power that is ideally used in a signal transmission.

MaxPowerVariance is the maximum difference in power that signaling devices may deviate from the NominalPower.

MaxAntennaGain is the gain of the antenna used in estimating ranges.

TimingRange is the maximum range supported in signaling due to the timing of signaling. Transitions to the next signal may not support the detection of a signal arriving late due to propagation range.

SquareLawRange is an estimate of the range of the signal using NominalPower and the MaxAntennaGain in a free space propagation environment. This specified range can be a function of the transmitted power but also of the timing of the signals.

TerrestrialRange is an estimate of the range of the signal using NominalPower and the MaxAntennaGain across the terrestrial environment for which the highway is specified. This specified range can be a function of the transmitted power but also of the timing of the signals.

UseAsymmetry is a restriction placed on the maximum asymmetry (a use may not exceed this value) between the range of signaling and the range of a transmission in the use of the highway. This is a fraction of the signaling range and it may be either less or more than the signaling range.

5.2.2.9.1 Range Symmetry Rules

The objective of these rules is to allow signaling to achieve results that support fair spatial access to spectrum by SDSs and to prevent SDSs from using spectrum in a way that is so asymmetric from the signaling results that it is likely to cause harmful interference to peers on the same highway that are spatially reusing the same lane. It is equally desirable that SDSs that use high gain antennas be able to achieve the benefits those antennas provide.

In most cases highway participants will achieve compliance with these rules by the initial configuration of SDSs and signaling or in the pairing of SDS devices with SHADs. In cases where SDS use is adaptable, these rules assume that each spectrum-dependent device has a good faith estimate of its range across its various configurations and fixes the pairing of those configurations with signaling configurations that match and follow these rules.

- Signaling shall not use a power greater than $(NominalPower + MaxPowerVariance)$.
- Signaling may use a power lower than NominalPower as long as the device performing the signaling has a NominalAntennaGain or greater and its receiver sensitivity would allow it to receive a signal sent at NominalPower using NominalAntennaGain at SquareLawRange in free space.
- A signaling device shall use the same antenna for sending and receiving signals.
- A signaling device operating from NominalPower to $(NominalPower + MaxPowerVariance)$ may use an antenna of any gain.

- An SDS’s power spectral flux density at one meter on a lane shall not exceed that specified by the combination of the lane’s ReferencePower, SpectrumMask, and SCMPowerMap.¹²
- An SDS shall not operate at a range greater than the TimingRange.
- An SDS shall not operate at a range greater than UseAsymmetry of the range of signaling (either SquareLawRange or TerrestrialRange depending on where the SDS is operated) in any direction with the exceptions that:
 - If the signaling is done with a high-gain directional antenna thus extending its range in a direction, then an SDS may be operated within the range of that signaling. This requires the SDS to have similar antenna directionality.
 - An SDS that uses a displaced SHAD for access shall operate within UseAsymmetry of the signaling’s range and so the permitted range of the SDS will be smaller or larger depending on direction. The SDS shall operate within the footprint of the signaling.

5.3 Distinguishing SDS Precedence Authorizations

Separate from the highway description are the authorizations given to a device that allows it to operate on the highway, establishes its lane selection and user/use precedence, maps its various functions to QoS levels, and assigns an SDS ID. Table 5-12 list the proposed parameters. These parameters are not part of the highway description. As a result they are not part of the PorPPParameters list. Each parameter is a field and has a type and value.

Table 5-12. Parameters Used to Govern SDS Device Access to Highways

Parameter	Type	Units/Enumerated Values	Occurrences
SDSID	integer	≥ 0	1
LSPrecedenceLevels	integer	≥ 0	1
UserUseLevels	integer	≥ 0	1
QoSLevels	integer	≥ 0	1
UserPrecedence	structure		1..n
UsePrecedence	structure		1..n
QoSPrecedence	structure		1..n
UserPrecedence			
UserName	string		1
Precedence	integer	≥ 0	1
PreemptionCount	integer	≥ 0	1
LSPrecedence	integer	≥ 0	1
UsePrecedence			
UseName	string		1
Precedence	integer	≥ 0	1
PreemptionCount	integer	≥ 0	1
LSPrecedence	integer	≥ 0	1
QoSPrecedence			
QoSName	string		1
Precedence	integer	≥ 0	1

¹² Computed as though it were supporting a far field estimate with the maximum gain of the antenna being used by the SDS.

Parameter	Type	Units/Enumerated Values	Occurrences
ReservationEligible	Boolean		1
ReservationLSPrecedence	integer	≥ 0	1

SDSID is the unique identification number provided to each SDS to distinguish it from others and to support rapid lane changes in serial out-of-band signaling. The process for managing these numbers must still be defined. The intent is for each SDS to have its own SDSID. Individual devices, such as radios that move between networks, may have different SDSIDs depending in which network they participate. Further, signaling may provide only a small number to be used. Factory assignments are not practical.

LSPrecedenceLevels identifies the number of precedence levels used as the basis for the LSPrecedence values in the authorization. When operating on a highway with a different number of levels, the level of precedence used by this SDS will be adjusted. The highest lane selection precedence is always set aside for urgent preemptions. Given a highway with LaneUsePrecedencePhases, the scaling is defined as

$$AdjustedLSPrecedence = \frac{2^{LaneUsePrecedencePhases} - 1}{LSPrecedenceLevels - 1} \cdot LSPrecedence \quad (5-1)$$

UserUseLevels identifies the number of user/use precedence levels used as the basis for the Precedence values in the authorization under UserPrecedence and UsePrecedence. When operating on a highway with a different number of UserUseLevels, the level of precedence used by this SDS will be adjusted. Given a highway with UserUsePhases, the scaling is defined as

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 3}{UserUseLevels} \cdot Precedence \quad (5-2)$$

when frames and epochs are available for contention access (i.e. the parameters FrameContention and/or EpochContention are present). In cases where only timeslots are available, the scaling is defined as

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 1}{UserUseLevels} \cdot Precedence \quad (5-3)$$

In anticipation that access to frames and epochs by devices must be controlled by the permissions given to those devices, both the UserPrecedence and UsePrecedence must have values that make this distinction. However, in some cases the highway design can support SDS access requirements with just timeslots where standard precedence applies. To support this duality, values between $(UserUseLevels + 1)$ and $(2 \times UserUseLevels)$ indicate frame access is allowed and the standard precedence is $(Precedence - (UserUseLevels + 1))$. Values between $(2 \times UserUseLevels + 1)$ and $(3 \times UserUseLevels)$ indicate epoch or frame access is allowed and the standard precedence is $(Precedence - (2 \times UserUseLevels + 1))$.

Given UserUseLevels the valid values for standard precedence extend from 0 to UserUseLevels. If the user or use is authorized to contend for a frame then the precedence value will be between $(UserUseLevels + 1)$ and $(2 \times UserUseLevels)$. If a timeslot is adequate for the user or use, then the device shall use the precedence

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 3}{UserUseLevels} \cdot (Precedence - (UserUseLevel + 1)) \quad (5-4)$$

If contending for a frame or epoch is not supported, and a timeslot is a sufficient transmission opportunity, then the device shall use the precedence

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 1}{UserUseLevels} \cdot (Precedence - (UserUseLevel + 1)) \quad (5-5)$$

If a frame is required for the user or use and can be contended for, then the device shall use the precedence for access to a frame for that lane design, usually $(2^{UserUsePhases} - 2)$.

If the user or use is authorized to contend for an epoch, then the precedence value will be between $(2 \times UserUseLevels + 1)$ and $(3 \times UserUseLevels)$. If a timeslot is adequate for the user or use, then the device shall use the precedence

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 3}{UserUseLevels} \cdot (Precedence - (2 \cdot UserUseLevel + 1)) \quad (5-6)$$

and contend for any timeslot. Otherwise, if a frame is adequate for its use and frames can be contended for it shall use the frame contention level and contend for the next frame. Otherwise the epoch precedence for access, usually $(2^{UserUsePhases} - 1)$, shall be used for the next available epoch contention.

If contending for a frame or epoch is not possible, and a timeslot is a sufficient transmission opportunity, then the device shall use the precedence

$$AdjustedUUPrecedence = \frac{2^{UserUsePhases} - 1}{UserUseLevels} \cdot (Precedence - (2 \cdot UserUseLevel + 1)) \quad (5-7)$$

QoSLevels identifies the number of QoS precedence levels used as the basis for the Precedence values in the authorization under QoSPrecedence. When operating on a highway with a different number of QoSLevels, the level of precedence used by an SDS will be adjusted. Given a highway with QoSPhases, the scaling is defined as

$$AdjustedQoSPrecedence = \frac{2^{QoSPhases} - 1}{QoSLevels} \cdot Precedence \quad (5-8)$$

UserPrecedence is a structure allowing a device to receive multiple specifications for user precedence. Multiple users may use the same device. Each entry has four values. UserName identifies the user, Precedence is the User/Use precedence the device uses when this user operates the device, PreemptionCount is the number of consecutive contentions where a reservation prevents access over which a use must wait before it can execute a reservation preemption, and LSPrecedence is the lane selection precedence that accompanies the user.

UsePrecedence is a structure allowing a device to receive multiple specifications for the different ways the device may use the spectrum. For example a multi-aperture system may be able to act as both a communications device and a radar. Depending on which is used the use precedence

may be different. Each entry has four values. UseName identifies the use, Precedence is the User/Use precedence the device uses when it is operating in the stated use, PreemptionCount is the number of consecutive contentions where a reservation prevents access over which a use must wait before it can execute a reservation preemption, and LSPrecedence is the lane selection precedence that accompanies the use.

QoSPrecedence is a structure allowing a device to receive multiple specifications for different levels of QoS depending on what operations it is performing. The most common use is to distinguish among different types of traffic in a network based on the urgency of the traffic. QoSName identifies the service type. Precedence identifies the QoS precedence level that may be used in gaining access. ReservationEligible indicates whether a device servicing this type of traffic may reserve a slot. Typically, reservations are made to support streaming traffic or a standing level of service. ReservationLSPrecedence specifies the lane selection precedence the device may use when it has made a reservation and only for the timeslot in which it has the reservation.

As described above there is a special precedence for authorization to contend for frame and epoch access. One additional precedence level is assigned to urgent responses that must begin immediately. The precedence for this use would be given as a use precedence of $(3 \cdot \text{UserUseLevels})$. This precedence supports transmissions that must start immediately. This use authorizes the device to transmit immediately, even if transmission starts in the middle of a timeslot, but then the expectation is for the device to contend with the highest precedence values in all contentions so long as it needs to contend. This precedence is more practical when using out-of-band signaling.

5.3.1 Rules for SDS Use of a Spectrum Highway

An SDS is given authorization to operate on a highway through the transfer of the SCM Authorization Set that defines that highway. However, before an SDS may operate on a highway, it must assess whether that highway is suitable for its purpose and the SDS can execute the signaling and meet the constraints of the highway design. Once an SDS determines that it should operate on the highway it follows rules that govern its behavior based on the permissions given to the SDS and the design of the highway. Section 5.3.1.1 provides a list of considerations that may be used by SDSs to assess the suitability of the highway. Section 5.3.1.2 presents rules that govern the assessment an SDS must make to determine if it is capable of operating on a highway. Section 5.3.1.3 contains rules an SDS must apply in its use of a highway based on its precedence and QoS permissions and the design of the highway.

5.3.1.1 Considerations for Highway Selection

It is anticipated that an operator of an SDS and in some cases the SDS itself will make the decision whether it will try to operate on one spectrum highway or another. Certain criteria are likely to drive that decision. They include:

- The devices of the SDS can operate at the frequency and within the bandwidth of the specified lanes.

- The range constraints of the highway will not adversely affect the operation of the SDS the power limitations of the SDS give it a disadvantage in gaining access and using the highway.
- The duration of the transmission opportunities will support the operations that the SDS performs.
- The device has a use that the highway can support.
- The recurrence of reserved slots can support the rate at which the device needs to transmit.
- The responsiveness of the system will support the needs of the SDS.
- The lane access has sufficient differentiation – whether user, use, or QoS – to enable the performance necessary for the SDS to function as intended.
- The occupancy of the highway by other SDSs will not have an adverse effect on the performance of the SDS (e.g. the highway is not too congested).

5.3.1.2 Rules for Access to Spectrum Highways

The SDS shall have the ability to discern whether it can operate on a highway based on the highway definition. It shall execute following checks before operating on a highway:

- Devices shall assess whether they can execute the contention signaling in all channels required to gain access to the defined lanes and shall defer operating on the highway if they cannot. These performance requirements would be listed in the ChannelParameters. Likely characteristics that may be specified that define performance requirements include the signal shape, the requirement to frequency hop and the allocated blanking interval, the method to generate the hopping sequence, the synchronization reference and tolerance, and the allowed processing time between phases.
- Devices shall only operate on a highway lane if they can operate within the spectral boundaries of the lane.
- Devices shall only operate on highway lanes that provide transmission opportunities of sufficient duration for the device to function as intended.

5.3.1.3 Rules for Precedence Behavior on Spectrum Highways

When a device needs to use a highway it will assess what resources it needs to contend for its next use, identifying the lane, the transmission opportunity, and the QoS. The device will also determine its LSPrecedence level for access. The following rules govern its attempt to gain access.

- If frame access is allowed on the highway, when a device becomes active it shall wait for the passing of at least one contention used for frame access before it may attempt to gain access.
- If epoch access is allowed on the highway, when a device becomes active it shall wait for the passing of at least one contention used for epoch access before it may attempt to gain access.

- When using in-band signaling:
 - A device shall not attempt to gain access to a timeslot on a lane during a frame in which another device has gained full frame access to the lane.
 - A device shall not attempt to gain access to a timeslot on a lane during an epoch in which another device has gained full epoch access to the lane.
 - A device shall not attempt to gain access to a frame on a lane during an epoch in which another device has gained full epoch access to the lane
- When using out-of-band signaling
 - A device shall hold its long-duration, frame, or epoch access, by using the override lane use precedence signal to hold the lane and the long-duration use priority to hold the long-duration access.
 - A device given the authority to preempt long-duration access may use the override lane use precedence signal and all signals of the Priority Phase Set to preempt access.
- Each device shall select a lane for operation before a contention:
 - If a device needs to use the highway, it shall select the minimum bandwidth lane and shortest transmission opportunity required to support the spectrum use task. It shall attempt to gain access in the next available contention for the chosen lane.
 - If a device is part of an SDS and has no need for access, but may need to receive transmission from another device in the SDS in a timeslot, e.g. a radio in a communications network, it shall select the lane used by the SDS and participate on the contention channels as a non-contender for that lane.
- If the highway has a lane use precedence phase set, a contending device shall contend according to the rules given in Section 5.2.2.3.1 using the LPChannel for the selected lane and use the highest lane selection precedence it is permitted by either its user, use, or QoS precedence.
- If the highway has a lane use selection phase and the device is a contender for lane selection (i.e. it survived the lane use precedence phase set according to the rules specified in Section 5.2.2.3.1 or there is no lane use precedence phase set) it shall contend on the LSChannel to set its choice of lane to be the active lane using the rules given in Section 5.2.2.4.1.
- At the conclusion of the lane selection phase set, a device determines which lanes are active and on which lane it should participate. Active lanes are determined by the rules given in Section 5.2.2.4.1.
 - If the lane the device chose at the start of the phase set to participate in is active, the device shall operate on the PChannel and CChannel for that lane.
 - Otherwise, if the lane the device chose at the start of the phase set for its spectrum use is not active, the device may participate in the contention for any active lane that will

meet its spectrum use objectives and if it selects one it shall operate on the PChannel and CChannel for that lane.¹³

- Otherwise, the device may participate on the PChannel and CChannel of any active lane as a non-contender.
- Prior to the Priority Phase Set, a contending device shall determine if it may use a reservation phase and the highest precedence based on either the user or the use.
- A contending device shall use the reservation phase of a Priority Phase Set if:
 - The contention is for a timeslot, the lane supports timeslot reservations, the device successfully contended and survived as a contender in the same ordinal timeslot in the previous frame, the current use is an extension of the same service or stream, and the QoS precedence allows reservations.
 - The contention is for a frame, the lane supports frame reservations, the device successfully contended and survived as a contender in the same ordinal frame in the previous epoch, the current use is an extension of the same service or stream, and the QoS precedence allows reservations.
- A contending device shall use the greater of its user or use precedence values in the following situations:
 - If the pending use requires a frame and the user or use precedence indicates a frame or epoch contention is permitted, the SDS shall use the frame precedence to gain access to a frame in a contention designated for that frame.
 - If the pending use can be supported by a timeslot and the user or use precedence indicates a frame (not an epoch) contention is permitted, the device shall use the standard precedence for the traffic which is computed using Equation (5-4) in lanes where frame contention is enabled and using Equation (5-5) in lanes where frame contention is not enabled.
 - If the pending use requires an epoch and the user or use precedence indicates an epoch contention is permitted, the SDS shall use the epoch precedence to gain access to an epoch in a contention designated for that frame.
 - If the pending use can be supported by a timeslot and the user or use precedence indicates an epoch contention is permitted, the device shall use the standard precedence for the traffic which is computed using Equation (5-6) in lanes where epoch and/or

¹³ There may be rules that SDSs have established internally that govern how to select a lane to participate after lane selection. The lane selection process does not preclude any device from contending for access to one of the active lanes. It does preclude devices from contending for access to inactive lanes.

frame contention is enabled and using Equation (5-7) in lanes where epoch and frame contention is not enabled.

- If using out-of-band signaling and a frame has been successfully accessed through contention then the device shall use frame precedence in all contentions for all timeslots for the duration of the frame.
- If using out-of-band signaling and an epoch has been successfully accessed through contention, then the device shall use epoch precedence in all contentions for all timeslots for the duration of the epoch.
- The device shall preempt a reservation if the SDS has a need for access and it has observed PreemptionCount of consecutive reservations since the time when that need was identified.
 - A device shall use the smaller PreemptionCount found in the UserPrecedence or UsePrecedence for the user and use of the pending spectrum use.
- The device shall use the QoS precedence indicated for its particular use of the lane.
- The device shall participate in the Priority Phase Set contention using the authorizations for reservation, user/use, and QoS as determined by the rules above, applying the rules defined in Section 5.2.2.5.1 for its behavior in the Priority Phase Set.
- Devices that survive the Priority Phase Set as contenders shall begin their participation in the Contention Phase Set as contenders.
- Devices that are non-contenders at the end of the Priority Phase Set shall begin their participation in the Contention Phase Set as non-contenders.
- Devices that participate in the Contention Phase Set shall use the specified CChannel for the lane and follow the rules defined in Section 5.2.2.6.1 for their behavior in the Contention Phase Set.
- Devices that survive the Contention Phase Set shall use the transmission opportunity for the purpose used to select the user, use, and QoS precedence levels.

5.4 Conclusion

This section in essence constitutes the rules manual for operation on a spectrum highway. It falls into two main parts: the definition of the highway and the signaling it uses for access (Section 5.2) and the guidance given to SDSs on their authorization to use different signaling features (Section 5.3). These are two separate definitions, one being part of a highway design and the other being given to an SDS based on its functions. This approach to highway design allows SDSs to adapt to operate on highways with different signaling designs. Appendices B and C provide exemplar designs of highway access for in-band and out-of-band signaling respectively. Appendix B also provides an exemplar design of the physical layer of signaling. The next section describes the configuration of SDSs with their authorizations for precedence in contention.

6 Device Access Authorizations

An important part of implementing highways is conveying to SDSs the access precedence they have for selecting lanes and for their use of the priority phase set. The best ways to assign priorities is still to be determined. A goal is to make it universal so that SDSs can easily move from highway to highway and carry the same precedence. This section presents an examples that illustrate how the precedence design may balance the diverse needs of different users. In this example, SDSs have different functions. It include SDSs that provide services that are routine, mission-critical function, or support both levels of mission importance: these are known as reactive mission devices

6.1 Routine Mission SDSs

Routine mission devices are those that provide a service that can generally be supported by regular access to a channel based on the demand of users. Interruption of that use is tolerated without significant impacts on the mission. In some cases, based on the scenario of use, these systems may perform mission-critical functions, in which case their precedence may change, or may be elevated for some period of time.

6.1.1 Communications SDS

Communications systems typically have multiple types of users, uses, and QoS levels. The tables in this section define examples of precedence permissions given for this type of device. Table 6-1 lists the levels considered as the basis for each type of precedence. Table 6-2 identifies the user types with special access based on their identities. Here, four user types are identified: a commander (CDR) above a brigade combat team (BCT) CDR, a BCT CDR, a battalion (BN) CDR, and a company CDR. Table 6-3 identifies the use types with special access. They are closely aligned with the QoS types that follow. These tables show that although a user may be low on the user precedence list, the traffic the communications SDS supports may have a higher precedence than that of all user types. Devices contending to send Critical or Flash Override messages will beat out all types of users in their precedence. Table 6-4 maps each of the message types used in a communications SDS to a QoS precedence level. It also identifies which traffic types can make reservations. In these examples all of them can. With this set or authorizations, certain SDSs can still be given user, use, or QoS precedence that can beat the precedence of a communications SDS.

Table 6-1. Example Access Authorizations for a Communications SDS

Primary Data Element	Sub - elements	Value	Notes
LSPrecedenceLevels		4	Assumes two phases
UserUseLevels		8	Assumes three phases
QoSLevels		8	Assumes three phases
UserPrecedence		See Table 6-2	Has several instances
UsePrecedence		See Table 6-3	Has several instances
QoSPrecedence		See Table 6-4	Has several instances

Table 6-2. Example Set of User Precedence Guidance for a Communications SDS

Primary Data Element	Sub -elements	Value	Notes
UserPrecedence			
	UserName	CDR > BCT CDR	
	Precedence	3	
	PreemptionCount	20	
	LSPrecedence	2	
UserPrecedence			
	UserName	BCT CDR	
	Precedence	2	
	PreemptionCount	20	
	LSPrecedence	2	
UserPrecedence			
	UserName	BN CDR	
	Precedence	2	
	PreemptionCount	40	
	LSPrecedence	2	
UserPrecedence			
	UserName	Company CDR	
	Precedence	1	
	PreemptionCount	60	
	LSPrecedence	1	
UserPrecedence			
	UserName	Common User	
	Precedence	0	
	PreemptionCount	100	
	LSPrecedence	0	

Table 6-3. Example Set of Use Precedence Guidance for a Communications SDS

Primary Data Element	Sub -elements	Value	Notes
UsePrecedence			
	UseName	Critical	
	Precedence	6	
	PreemptionCount	10	
	LSPrecedence	2	
UsePrecedence			
	UseName	Flash Override	
	Precedence	4	
	PreemptionCount	20	
	LSPrecedence	2	
UsePrecedence			
	UseName	Flash	

Primary Data Element	Sub -elements	Value	Notes
	Precedence	2	
	PreemptionCount	30	
	LSPrecedence	2	
UsePrecedence			
	UseName	Priority	
	Precedence	1	
	PreemptionCount	100	
	LSPrecedence	1	
UsePrecedence			
	UseName	Routine	
	Precedence	0	
	PreemptionCount	200	
	LSPrecedence	0	

Table 6-4. Example Set of QoS Precedence Guidance for a Communications SDS

Primary Data Element	Sub -elements	Value	Notes
QoSPrecedence			
	QoSName	Critical	
	Precedence	6	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	2	
QoSPrecedence			
	QoSName	Flash Override	
	Precedence	3	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	2	
QoSPrecedence			
	QoSName	Flash	
	Precedence	2	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	2	
QoSPrecedence			
	QoSName	Priority	
	Precedence	1	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	1	
QoSPrecedence			
	QoSName	Routine	
	Precedence	0	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	0	

6.1.2 Sensing SDS

A sensing SDS may operate at will in any spectrum it can without coordination but sensing adversary activities on a spectrum highway that is busy may require the sensor to silence other SDSs operating on the highway. Under the rules governing highways, a sensing SDS can silence friendly users so adversary actions on the lanes can be observed. Because the signaling of spectrum highways seeks to orchestrate spectrum reuse, sensing may be of two types: local or global. In local sensing, the device will individually contend with the intent of gaining access. This has the effect of silencing other SDSs on the selected highway lane up to a two-hop range. In global sensing, sensing devices will coordinate among themselves contending at the same time to silence as much of the highway as possible for the duration of the transmission opportunity being contended. This collaboration amounts to signaling for the same timeslots in unison with their high precedence. This timing is not part of the authorizations but part of the sensing plan with these devices.

The effectiveness of silencing the highway will depend on the geographical distribution of the contending sensors across the spectrum highway. Table 6-5 defines precedence levels that are the basis of the precedence given to these devices. Table 6-6 lists the user precedence levels and Table 6-7 lists the uses and their precedence. There is only one user precedence level in Table 6-6 and in this case it has no relevance because all of the uses listed in Table 6-7 have a higher precedence level. Table 6-8 lists the QoS precedence that matches each of the uses.

Table 6-5. Example Access Authorizations for a Sensing SDS

Primary Data Element	Sub -elements	Value	Notes
LSPrecedenceLevels		4	Assumes two phases
UserUseLevels		8	Assumes three phases
QoSLevels		8	Assumes three phases
UserPrecedence		See Table 6-6	Has several instances
UsePrecedence		See Table 6-7	Has several instances
QoSPrecedence		See Table 6-8	Has several instances

Table 6-6. Example Set of User Precedence Guidance for a Sensing SDS

Primary Data Element	Sub -elements	Value	Notes
UserPrecedence			
	UserName	Any user	
	Precedence	0	
	PreemptionCount	200	
	LSPrecedence	0	

Table 6-7. Example Set of Use Precedence Guidance for a Sensing SDS

Primary Data Element	Sub -elements	Value	Notes
UsePrecedence			
	UseName	Global Sense	
	Precedence	7	

Primary Data Element	Sub -elements	Value	Notes
	PreemptionCount	10	
	LSPrecedence	3	
UsePrecedence			
	UseName	Local Sense	
	Precedence	1	
	PreemptionCount	50	
	LSPrecedence	2	

Table 6-8. Example Set of QoS Precedence Guidance for a Sensing SDS

Primary Data Element	Sub -elements	Value	Notes
QoSPrecedence			
	QoSName	Global Sense	
	Precedence	7	
	ReservationEligible	FALSE	
	ReservationLSPrecedence	2	
QoSPrecedence			
	QoSName	Local Sense	
	Precedence	1	
	ReservationEligible	FALSE	
	ReservationLSPrecedence	2	

6.2 Mission-Critical SDSs

Mission-critical devices are those that provide a service that depends on regular and reliable access. Examples may include radars or other SDSs that must periodically observe the environment. The sensing SDS defined in Section 6.1.2 could easily be classified as a mission-critical SDS in its global sensing mission.

6.2.1 Radar SDS

Table 6-9 provides the precedence levels used as the basis for the precedence definitions for a radar. A radar must be able to operate on a reliable basis to adequately detect and track what its targets. In this role, it is given higher precedence to achieve this critical mission function. Table 6-10 shows that in this authorization that there is no distinction based on the user. Here again there is no distinction between user types and the uses dictate the user/use precedence for gaining access. Tracking has the higher of the two precedence levels. Table 6-11 defines the use precedence and Table 6-12 defines the use QoS. Both types of uses are eligible to reserve slots for access

Table 6-9. Example Access Authorizations for a Radar SDS

Primary Data Element	Sub -elements	Value	Notes
LSPrecedenceLevels		4	Assumes two phases
UserUseLevels		8	Assumes three phases
QoSLevels		8	Assumes three phases
UserPrecedence		See Table 6-10	Has several instances
UsePrecedence		See Table 6-11	Has several instances
QoSPrecedence		See Table 6-12	Has several instances

Table 6-10. Example Set of User Precedence Guidance for a Radar SDS

Primary Data Element	Sub -elements	Value	Notes
UserPrecedence			
	UserName	Any user	
	Precedence	0	
	PreemptionCount	200	
	LSPrecedence	0	

Table 6-11. Example Set of Use Precedence Guidance for a Radar SDS

Primary Data Element	Sub -elements	Value	Notes
UsePrecedence			
	UserName	Track	
	Precedence	7	
	PreemptionCount	2	
	LSPrecedence	3	
UsePrecedence			
	UserName	Detect	
	Precedence	6	
	PreemptionCount	20	
	LSPrecedence	2	

Table 6-12. Example Set of QoS Precedence Guidance for a Radar SDS

Primary Data Element	Sub -elements	Value	Notes
QoSPrecedence			
	QoSName	Track	
	Precedence	7	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	3	
QoSPrecedence			
	QoSName	Detect	
	Precedence	6	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	2	

6.3 Reactive Mission SDSs

Reactive mission SDSs are those that have a routine or critical function that may trigger a very significant behavior that would preempt other types of uses. An example would be a counter-radio-controlled improvised explosive device (IED) electronic warfare system. The system would act as a mission-critical device in its sensing of the environment, but would then preempt all other uses if it needed to transmit a signal to interrupt the radio control of an IED.

6.3.1 Counter Radio-Controlled IED Electronic Warfare (CREW) SDS

A CREW device may be designed to listen periodically on a channel and, upon detection of a suspicious signal, to immediately attempt to disrupt the transmission. This device performs two activities: first is sensing and then jamming. Table 6-13 lists the precedence levels used as the basis for precedence of the CREW SDS. Table 6-14 shows the user precedence levels. There is no distinction made between users. The CREW use dictates the user/use precedence used in the Priority Phase Set. Table 6-15 list the precedence levels for the different functions of the CREW SDS. The jamming activity, which is triggered by sensing a suspicious signal, is executed with the highest use priority and the device is permitted to immediately preempt any timeslot reservation. Table 6-16 lists the QoS precedence for each of the uses; they align with the use precedence. Note that the sensing activity may make periodic reservations but jamming cannot. There is no reason for SDSs to jam periodically but periodic sensing provides greater assurance that trigger signals will be detected and the jamming precedence ensures jamming will be responsive to the detection.

Table 6-13. Example Access Authorizations for a CREW SDS

Primary Data Element	Sub -elements	Value	Notes
LSPrecedenceLevels		4	Assumes two phases
UserUseLevels		8	Assumes three phases
QoSLevels		8	Assumes three phases
UserPrecedence		See Table 6-14	Has several instances
UsePrecedence		See Table 6-15	Has several instances
QoSPrecedence		See Table 6-16	Has several instances

Table 6-14. Example Set of User Precedence Guidance for a CREW SDS

Primary Data Element	Sub -elements	Value	Notes
UserPrecedence			
	UserName	Any user	
	Precedence	0	
	PreemptionCount	200	
	LSPrecedence	0	

Table 6-15. Example Set of use Precedence Guidance for a CREW SDS

Primary Data Element	Sub -elements	Value	Notes
UsePrecedence			
	UseName	Jam	
	Precedence	8	
	PreemptionCount	0	
	LSPrecedence	4	
UsePrecedence			
	UseName	Sense	
	Precedence	7	
	PreemptionCount	10	
	LSPrecedence	3	

Table 6-16. Example Set of QoS Precedence Guidance for a CREW SDS

Primary Data Element	Sub -elements	Value	Notes
QoSPrecedence			
	QoSName	Jam	
	Precedence	8	
	ReservationEligible	FALSE	
	ReservationLSPrecedence	4	
QoSPrecedence			
	QoSName	Detect	
	Precedence	7	
	ReservationEligible	TRUE	
	ReservationLSPrecedence	3	

6.4 Conclusion

This section provided several examples of SDS precedence authorizations, showing how they might differ among multiple systems. These are notional because, at present, no attempts have been made to determine the best ways to distribute precedence levels among multiple system types and to assess their interactions on the same highway. There may be many precedence levels within an SDS operator and operational space and some fraction of these may be of greater importance than those of another SDS type. Establishing these levels should anticipate the interaction among the broad range of SDSs that may operate on a spectrum highway and ensure that the most important device missions receive precedence over those that have less importance.

7 Summary and Next Steps

This report provides a concept for a new DSA approach that follows the metaphor of a highway where bands of spectrum in space are subdivided into smaller bands called lanes. Devices that operate on the highway independently make choices as to the lanes they want to operate in and then use signaling to interact with other devices to resolve which of the devices gets access. SDSs and users of those SDSs support a varied set of missions and applications and the devices used by some will have precedence over others in gaining access to a lane on a highway. Precedence levels can be based on both the user of the device and the use of the device.

This report describes a comprehensive set of operational components required to support the advanced spectrum highway sharing construct.

- A method to define the geographical, spectral, and power limits of a spectrum highway and its subdivision into lanes, including the ad hoc ability to combine lanes into broader bandwidth lanes.
- A method to define the subdivision of lanes into a timing hierarchy for transmission opportunities
- A signaling approach that supports collaboration among device in settling
 - Which lanes are active
 - Which devices gain access based on a hierarchy of precedence levels that in turn are based on the user or use and the required QoS.
 - Ad hoc access to frames or epochs
 - Ad hoc combining of lanes for larger bandwidth use
 - Spatial reuse of the highway based on the collective demand for the highway and the current precedence of the individual devices that make up that demand.
- Rules that govern how a device behaves on a highway that determine
 - Which precedence levels a device may use in an access attempt
 - How the outcome of signaling determines the lanes that are made active
 - How devices interact to ensure the granting of access to the devices that have highest user/use precedence
 - How devices interact to ensure the granting of access to the devices that have the highest QoS authorization
 - How devices may use signaling to make an access reservation
 - How authorized devices can preempt a reservation
 - How devices become authorized to preempt a reservation
 - How SDSs and devices are given the precedence levels they may use to gain access based on their user, use, or the QoS they need

- Metrics to compare the performance of different highway designs and the engineering tradeoffs that follow from seeking particular performance objectives

Since this design is, thus far, conceptual, further activities are needed to demonstrate its efficacy, to examine design trade-offs in greater detail, and to create the mature technical components that are necessary for multiple SDS operation. The next several sections propose applications of spectrum highways and follow-on activities that would support creating the technical components for highways and demonstrating their operation on a highway.

7.1 Applications

Spectrum highways provide a workable and effective option for spectrum sharing that can replace other more expensive methods and create environments in which a greater variety of SDSs can share spectrum. The use of spectrum highways with their rules and contention mechanisms would result in an unmatched optimality and efficiency in DSA of heterogeneous SDSs in a dynamic environment. It is unlikely that highways would be used in short range spectrum use typical of the ISM bands where LBT protocols perform well. The sweet spot for highways are environments where intermediate range sharing is sought, so ranges from 300 to 20,000 meters. Highways can be used for all ranges but signaling cost increases with range since signaling must be designed to support longer propagation times. Several example uses are described below.

7.1.1 Three-Tier Spectrum Sharing

Currently, industry, under the rules for the Citizen's Band Radio Service (CBRS) in 3.5 GHz, is building a three-tier spectrum sharing system that requires all commercial and civilian devices, which make up two of the tiers, to connect to a Spectrum Access System (SAS). The SAS arbitrates the access of these Citizen Band Radio Service Devices (CBSDs), identifying the channels they may use and directing them to turn off if a higher precedence user needs the channel. It is a three tier system consisting of incumbent SDSs, Priority Access License (PAL) SDSs, and General Authorized Access (GAA) CBSDs. The incumbent SDSs do not connect to the SAS and operate at will. It is expected that an Environmental Sensing Capability (ESC) will monitor the environment and notify SASs of observed incumbent activity so the SASs can turn off and move PAL and GAA CBSDs into the bands that the incumbent SDS is using. The owners of PALs have the second status that they buy at auction that allows them to operate a collection of CBSDs on particular channels over an area without competition. There is a limit to the quantity of PAL licenses that can be purchased in any one area to ensure some spectrum remains for GAA CBSDs and to support the movement of PAL CBSDs to new channels when incumbents become active. The channels are managed through the internet and all PALs and GAA CBSDs are required to continuously connect to the SAS.

The alternative afforded by use of spectrum highway technology replaces the SAS and uses signaling to arbitrate access precedence and the sharing of spectrum among the CBSDs. The ESC could operate a collection of signaling stations to clear spectrum when incumbents arrive and in some cases incumbents may operate their own signaling stations for their protection from PAL and GAA CBSDs. PAL CBSDs would still get their status through an auction, which would authorize them to use a higher precedence on identified channels when they are in a

particular geographical area. GAA CBSDs would monitor the signaling on the highway and then choose the lane that they believe would best support their use and contend for access.

Advantages of this highway approach are that systems can be mobile and that signaling will naturally account for environmental effects that influence propagation; these are important effects that are hard for SASs to recognize and factor into their computations. It is likely to result in a greater density of users.

A benefit of the highway approach is that the technology can support any number of “tiers.” Given the technology, regulators can rapidly design multi-tiered access in new bands to accommodate any mix of legacy users and new users.

7.1.2 Intermediate-Range Spectrum Commons

An intermediate-range spectrum commons is the equivalent to the commons of ISM bands but at ranges up to tens of kilometers. While in the ISM bands sharing occurs because power is regulated and kept low, in the intermediate spectrum commons sharing occurs because SDSs would use the rules and methods of spectrum highways. This approach is most useful in environments where the range of transmission is further than that supported by the ISM bands and where a greater variety of devices use the highway and their use is sporadic.

7.1.3 Spectrum Maneuver Spaces

A spectrum maneuver space is a five dimensional space within which SDSs should be able to decide independently where to operate. Each SDS is provided the definition of multiple highways and then makes its own decision on which highway and lane to operate. In dynamic, congested, and contested environments the decision to maneuver and operate in other spectrum can be motivated by the activities of peers and those of adversaries. Operations can also be executed with the specific intent of maintaining initiative in the use of spectrum and avoiding an adversary’s attempt to deny access.

7.1.4 Efficient Spectrum Management

Today, spectrum access is centrally managed and provides limited ability for end systems to adapt to the environment. Central management of dynamic use does not allow the manager to anticipate the nuances of operational use that come from mobility and changes in operational intensity. As a result, the central manager is very conservative in assigning channels, thus limiting the number of users of spectrum and creating inefficiency when those users only access spectrum sporadically. The alternative to the current system is for the central manager to take some risk and overload assignments. The potential negative outcome is that SDSs that are assigned the same spectrum can intersect and disrupt each other’s use of that spectrum. The advantage of spectrum highways is that systems work collaboratively to share spectrum. Spectrum highways would allow these systems to interact in a fair way for access without the risk of one SDS blocking another from access. Spectrum managers would still have a means to control spectrum use through the design of the highways and the precedence given to SDSs. Using highways greatly simplifies spectrum management and allows SDSs to adapt to the local demand for spectrum.

7.2 Further Design Work

Creating a spectrum highway enterprise will require multiple technologies to be developed. The following sections describe the more obvious technologies.

7.2.1 Highway SDSs

SDSs must be designed to operate on spectrum highways. Highway SDSs should be designed to include the following features:

- Devices shall be able to receive precedence authorizations and be programmed to strictly follow them.
- Devices shall be able to receive highway definitions and to understand which of their operational modes can fall within the lanes of that definition to include the spectral, temporal, and range constraints of the lanes.
- Devices that perform their own signaling shall know which signaling designs they are capable of executing and in which bands and shall only operate in highways where they can comply with the signaling specification.
- Devices that use a SHAD shall be compliant with the communications standards for SDS-SHAD communications.
- Devices that use a SHAD shall be able to identify the transmission opportunities and lanes that are required for their function and be able to communicate that requirement to the SHAD.
- Devices that use a SHAD shall be able to receive information on highway activities that inform its decisions on lane selection.
- Devices may be designed to allow SHADs to make decisions on lane use; in this case the device shall be able to communicate to the SHAD its ability to operate on lanes and be able to react quickly to lane access notifications from the SHAD.
- Devices shall interface with SHADs and use signaling designs that provide sufficient time for all parts of this system to operate correctly. A system that cannot meet the timing constraints of the transmission opportunities shall not operate on the highway.
- A device that uses a SHAD shall be rated on its speed of reacting to access notifications from the SHAD
- Advanced devices may be designed to receive multiple highway definitions and then select the highway to use based on their location, congestion on the highways, and other factors that indicate the suitability of the highway for the device's spectrum needs.

7.2.2 Signal Designs

Section B.7 in Appendix B provides an example of a signal design. That design assumes signal frequency hopping and a narrowband signal that resulted in a long-duration signaling phase. Alternative designs may make other assumptions about the signal requirements. Single-channel signaling or the incorporation of multiple oscillators in signaling devices could both be used to

eliminate or reduce the time associated with the blanking interval. Wideband signals can reduce the duration of the signals themselves. Other timing factors can be adjusted to account for highway ranges.

In the military context, signal resilience will be important and research may be applied to creating designs that have a low probability of intercept (LPI) or low probability of detection (LPD). Efforts may be directed at making the signals very difficult to spoof.

7.2.3 SHAD Design

The goal of creating SHADs is to simplify SDS design and to enable very sophisticated signaling that the SHAD is specialized to in execute. If SHADs are used broadly, their production can provide cost savings through economies of scale. Further, use of SHADs avoids the cost of achieving signaling performance in the production of SDSs.

A SHAD design should include the following features:

- A SHAD shall be designed to execute some quantity of known signal designs.
- A SHAD shall be able to receive a highway definition and be able to assess whether it can perform the signaling defined in the Protocol or Policy constructs of the highways lanes.
- A SHAD shall be able to communicate to SDSs using a defined protocol (See Section 7.2.4).
- A SHAD shall be rated on its ability to quickly notify an SDS of its access success.

7.2.4 SHAD/SDS Communication

SHAD/SDS communication protocols will be necessary at the interface between SHADs and SDSs. This messaging of these protocols will convey information from the SDS about its precedence and pending requirements for a transmission opportunity including its precedence and QoS, and information from the SHAD about observations of highway activity, and outcomes of contentions that require the SDS to transmit or receive.

Defining the messaging will likely depend on the distribution of awareness and decision making. The awareness and decision making could be located in the SHAD, the SDS, or shared. Table 7-1 lists the differences in awareness and actions of the SDSs and SHADs based on the awareness strategy. These are presented as alternatives. Lowering the SDS requirements to manage operating on highways and placing as much access decision making in the SHAD probably provides both the most operationally and economically effective approaches to enabling highways. In this approach, SDSs would only have to communicate the access needs to the SDS and respond to the access outcomes reported by the SHAD. The harder highway access development can be built into a common SHAD that can be used by any SDS.

Nevertheless, some SDSs, specifically networking technologies, may want to have greater decision making roles so they can act in anticipation of their needs. They would already have the functionality to build the network topology, to route, and to manage traffic. An example of decision making they may want to own is the selection of the lane for operation. These networking SDSs may decide to distribute links to different highway lanes and want to direct the SHAD's movement among lanes to match the timing of the links.

Table 7-1. Strategies for Splitting Awareness and Decision Making between SDSs and SHADs

Locality of Awareness	SDS		SHAD	
	Awareness	Actions	Awareness	Actions
SDS	<ul style="list-style-type: none"> Receives the highway definitions and authorized precedence levels for signaling. Understands the SDS’s spectrum access needs and which highways, lanes, and transmission opportunities can support those needs. Tracks the state of the collective SDS and device movement among highways and lanes. 	<ul style="list-style-type: none"> Decides which highway, lane, and transmission opportunity duration to use. Directs the SHAD on contention actions. Guides SHAD in assessing highway use. Responds to SHAD feedback on gaining access and uses the highway either as a transmitter or a receiver. 	<ul style="list-style-type: none"> Learns the signaling design from and is guided by the SDS in contention. The SDS decides which lanes and which transmission opportunities to contend for and informs the SHAD on the precedence and the QoS to use. 	<ul style="list-style-type: none"> Contends for access as directed and informs the SDS of the outcomes. Indicates when a peer of the same SDS gains access to a lane. Provides feedback on highway use.
Shared	<ul style="list-style-type: none"> Receives the highway definition and authorized precedence levels for signaling. Understands the SDS’s spectrum access needs and which highways, lanes, and transmission opportunities can support those needs. 	<ul style="list-style-type: none"> Decides type of transmission opportunity that will support its needs and directs the SHAD to contend for its use. Configures itself to operate on the lane directed by the SHAD. Responds to SHAD feedback on gaining access and uses the highway either as a transmitter or a receiver. 	<ul style="list-style-type: none"> Receives the highway definition and the authorized precedence levels for signaling for the attached SDS. Tracks the state of the collective SDS and device movement among highways and lanes. 	<ul style="list-style-type: none"> Selects the highways and lanes to use. Contends for access as directed and informs the SDS of the outcomes.
SHAD	<ul style="list-style-type: none"> Is aware of the bandwidth and the required transmission opportunity durations that can support its operation. 	<ul style="list-style-type: none"> Informs the SHAD of the requirements for its next use of the highway. Configures itself to operate on lanes as directed by the SHAD. Responds to SHAD feedback on gaining access and uses the highway in the capacity the SHAD identifies, either as a transmitter or a receiver. 	<ul style="list-style-type: none"> Receives the highway definitions and authorized precedence levels for signaling for the attached SDS. Understands the SDS’s ability to operate in spectrum and which highways, lanes, and transmission opportunities can support its operation. Tracks the state of the collective SDS and their movement among highways and lanes. 	<ul style="list-style-type: none"> Selects the highway and lane. Contends for access to transmission opportunities that will support the SDS’s operation. Directs the SDS on which lanes to operate and when to operate as a transmitter or a receiver.

7.2.5 SDS Precedence Guidance

Both experimentation and modeling and simulation (M&S) will be needed to understand the likely interactions of SDSs based on the types of precedence they are given for both user/use and QoS. A goal of giving SDSs access to highways should be to avoid concentrating high-

precedence users on the same highway where their interaction leads to little service differentiation. Similarly, lower precedence users should not be starved of access because of the density of high-precedence users. The expected product of research would be guidelines on how to mix SDSs of different precedence and with different QoS requirements on the same highway.

7.2.6 Standardization

The access design guidance described in this report (see Section 5.2), physical layer signaling designs (see Section 7.2.2), SHAD/SDS communications (see Section 7.2.4), and the specification of access precedence to SDSs (see Section 5.3) are all candidates for standardization. Such standardization would make it easier for component manufacturers to develop systems that can operate on highways, for spectrum managers to implement highway designs, and for developers to apply innovation in SDS design that can exploit this access paradigm. Any serious attempt to make spectrum highway access real would have to be matched to these standardization efforts.

7.2.7 Regulation

Civilian use of spectrum highways must be matched with a regulatory regime that oversees the ecosystem. Regulators must certify that equipment is suitable for operation on the spectrum highways. SHADs and SDSs are suitable if they are able to operate according to the specifications of the standards and follow the intent of highway designs. Regulators must approve the matching of precedence to SDSs and certify the infrastructure that provides SDS precedence to SDS operators, and must guide the creation of the infrastructure that grants access to highways to SDSs. Finally, regulators must create an enforcement capability. This capability would be designed to verify that only authorized SDSs operate on the highways, that SDSs are using the precedence levels they are authorized to use, and that SDSs operate within the transmission restrictions that apply to the highways on which they are operating.

7.3 Approaches to Validate Spectrum Highway Efficacy

Creating spectrum highways will require demonstrating that they achieve the desired performance. This type of development can be evolutionary.

7.3.1 Modeling and Simulation

The quickest environment to stand up and to verify the functioning of these techniques would be in simulation. The first simulations would entail verifying the outcome of signaling and the temporal and spatial access that signaling renders based on the mix of devices and signaling designs. A higher order simulation would marry SDS models with signaling models and measure the performance of the SDSs as a result of highway designs, the mix of SDSs on the same highway, the precedence given to those SDSs, and the signaling designs.

7.3.2 Small Scale Demonstration

Small-scale demonstrations would use a propagation matrix simulator to demonstrate real signaling and the interactions between SDSs. The first demonstration would entail the implementation of signaling and verification of its performance in arbitrating access as the

signaling designs intend. These experiments may also address verification of the resilience of different physical layers used for the signaling. The next steps are to marry signaling with actual SDS access and verify that SDSs will perform as intended.

7.3.3 Spectrum Colosseum

The Spectrum Colosseum is the environment that DARPA created for SC2. It provides a large scale propagation matrix on which different environments can be simulated. The objective demonstration for a highway may use the same sharing scenarios developed for SC2 and have SDSs collaborating through signaling rather than a separate control channel. DARPA or another entity could create a similar project. A subset of performers would compete to design the physical layer of signaling and the SHADs to perform the signaling. Another subset of performers would provide SDSs that use the SHADs. These performers would collaborate on creating the highway designs that enable their systems and the signaling and precedence used to arbitrate access among SDSs that share spectrum. Performers would create their own algorithms for selecting the lanes of a multilane highway in which they would participate. The scoring of performance could be the same as that used for SC2 or expanded to measure the relative benefits afforded by signaling design, highway design, and SDS design.

7.3.4 Live Experimentation

After some level of development of physical components - SHADs and SDSs - that can operate on a highway and with the validation that they perform as intended either in the small-scale or Colosseum demonstrations, the next step would be to demonstrate their actual operation over the air.

7.4 Conclusion

This section has described the applications, the required development, and the experimentation that could be performed to demonstrate the efficacy of the spectrum highways that are proposed in this report. The great potential for this spectrum highway concept to dramatically improve spectrum sharing and to enable spectrum maneuver on an SDS's initiative provides the mandate that the DoD take action quickly to design highways and SDSs that can operate on them and to validate that these system will work. The DoD should undertake a significant effort to demonstrate these systems – an effort on the scale of a DARPA project. Such a project would progress through the activities described in Section 7.3.

8 References/Bibliography

- [1] IEEE Std 1900.5.2 2017, Standard for Method for Modeling Spectrum Consumption, Dec 2017.
- [2] J. A. Stine, G. de Veciana, K. Grace, and R. Durst, “Orchestrating spatial reuse in wireless ad hoc networks using Synchronous Collision Resolution,” *J. of Interconnection Networks*, Vol. 3 No. 3 & 4, Sep. and Dec. 2002, pp. 167 – 195.
- [3] K. H. Grace, J. A. Stine, R. C. Durst, “An approach for modestly directional communications in mobile ad hoc networks,” *Telecommunications Systems J.*, March/April 2005, pp. 281 – 296.
- [4] U.S. Patent No. 7,266,085, J. A. Stine, “An access and routing protocol for ad hoc networks using synchronous collision resolution and node state dissemination,” 2007
- [5] U.S. Patent No. 7,653,003, J. A. Stine, “Access protocol for wireless ad hoc networks using synchronous collision resolution,” 2010

This page intentionally left blank.

Appendix A Highway Design Performance Measures

The performance measures described in the appendix are designed to capture the relevant qualities of spectrum highway access. As described in the body of this report, the signaling approach to highway access provides an almost unlimited number of ways to arbitrate access based on design choices. Each design choice also requires trades in performance. These measures can be used to compare alternative designs.

A.1 Temporal Overhead

Temporal overhead is the loss of capacity on the highway that results from the duration of the contention design and the time to transition between users. It is applicable to both in-band and out-of-band contention designs although it is mostly an issue with in-band signaling. The measure is the fraction of time spent on average in contentions rather than payload transmission. It should be computed over an epoch as the fraction of time set aside for contentions during the epoch divided by the duration of the epoch.

A.2 Spectral Overhead

Spectral overhead is the loss of spectrum use that results from setting aside spectrum for out-of-band contention for a spectrum highway. The measure is the fraction of spectrum bandwidth set aside for CRS against all the spectrum used by the highway, the total bandwidth used to support the highway lanes and the bandwidth used for contention.

A.3 Use Efficiency

The efficiency metric captures the average wasted time and bandwidth in a timeslot. It measures the fraction of the spectral resource that is used. The measure of the use is the product of the bandwidth and the duration of the use. The efficiency metric is the ratio of that use to the product of the total bandwidth of the lane and duration of the timeslot. In a network with multiple types of uses that have different bandwidths and durations, the network efficiency metric would be the average efficiency of all SDSs on the highway as they use the highway collectively. This measure should apply to a highway in operation and would capture how well the highway is designed for the SDSs that use it.

A.4 Temporal Effectiveness

The temporal effectiveness metric identifies the fraction of transmission opportunities that can accommodate a device's use in a time period matching the longest transmission opportunity. As an example, consider a highway with a 10-slot frame and a four-frame epoch that can support contentions for timeslots, frames, and the epoch. There are 45 feasible transmission opportunities: one epoch transmission opportunity, four frame transmission opportunities, and 40 timeslot opportunities. If a device's use lasts longer than a timeslot and less than a frame then there are five opportunities and the metric is 5/45. The network temporal effectiveness is the average of the temporal effectiveness measures of all devices.

A.5 Congestion

The congestion metric assesses the operation of a highway. The measure will typically reflect access latency: the time from when a use is identified until the time that it is supported by the highway. In severe cases, uses will not be supported. The congestion metric therefore provides two values: the average access latency and the dropped use count.

A.6 Responsiveness

The responsiveness measure captures the ability to redirect the use of spectrum and to preempt reservations. It is the sum of the access delay and slotting delay. This measure provides a means to compare designs that can allow immediate jamming on a channel. It is also a measure of how much disruption jamming may cause.

A.6.1 Access Delay

Access delay is the time from the contention to the actual use of spectrum, measured from the beginning of a contention to the beginning of the use of the spectrum that the contention resolves. In in-band designs with consolidated signaling it will vary from ordinal slot to ordinal slot. Separately, it is a measure of how far in advance the decision must be made to redirect the use of spectrum. The highway design performance is provided by giving the longest and shortest duration delay.

A.6.2 Slot Delay

Slot delay is the time from the start of a transmission slot and to the start of the next slot contention. For a per-slot in-band CRS design it amounts to the duration of the transmission slot where it begins after a contention to the start of the next contention. For a consolidated in-band design, it is the average time from the beginning of the transmission slots of an epoch to the beginning of the contention for the same ordinal slot of the next frame. For serial out-of-band contention, it is the average time from the beginning of a transmission slot to the CRS for each of the primitive lanes. The lane selection signaling is skipped in these measures. The slot delay is a measure of the time it takes to observe a spectrum use and then to change that use through a contention.

A.7 Spatial Capacity

This is a measure of the average number of concurrent successful users (i.e. transmitters) of a lane. It is defined as the density of user per unit of area. The role of CRS is to arbitrate access and to orchestrate reuse. The density of contention survivors will be a function of the range of the signals. Greater spatial capacity comes from choosing a range that results in a close separation of transmitters that reliably supports their effectiveness, (e.g., closing a link with a neighbor). Reference [2] provides simulation results of CRS in action that can be used to estimate the spatial capacity given the ranges of signals. Reuse is a function of both the ranges of the signals and the operational use of the lane, which includes density of users and demand for use. The best reuse occurs when there is a high density of demand across the entire space of the highway.

A.8 Spectral Agility

This is a measure of how long it takes to move all participants in an SDS to a different lane. In most cases this is a function internal to the SDS. Devices that act independently can move much more rapidly than networked devices, which require coordination among the devices. However, some signaling techniques, e.g. serial out of band signaling with SDS ID, can support this agility from timeslot to timeslot.

A.9 Asymmetry Risk

The asymmetry risk is a measure of the extent to which a particular highway instantiation with its collection of participants will result in contentions with a non-ideal outcome in which a lower priority contender can preempt a higher priority contender because of asymmetry.

A.9.1 Power Asymmetry

With power asymmetry, an SDS with greater transmit power will be able to signal and reach a larger area than an SDS with lower power. The low-power SDS will be able to hear the signals of the high-power SDS but in some cases the high-power SDS will not be able to hear the signals of the low-power SDS. In these cases, the higher power SDS will have an advantage over lower power devices and will be able to suppress them with a lower precedence signal in the precedence and priority phases and with later signals in the contention phases. The highway will still function but it will not be as fair.

The measure for this risk is the ratio of the additional area covered by the high-power device's signal compared to the area it would cover if there were power parity with the lower powered device. The value of this measure depends on the range at which a signal attenuates and cannot be heard. In practice, SDSs are mobile and so the pathloss will vary because of differences in the environment that come with differences in location. To simplify this measure it can be complemented with one parameter: a pathloss exponent, n , for a log-distance pathloss model, making it agnostic to the actual environment. Given the high and low power, p_h and p_l , in Watts (not dBW) and assuming equivalent omnidirectional antennas in two dimensions, the power asymmetry risk (PAR) measure is computed as

$$PAR = \left(\frac{p_h}{p_l} \right)^{\frac{2}{n}} - 1 \quad (\text{A-1})$$

A.9.2 Weighted Power Asymmetry

Weighted power asymmetry is a measure of the effect of asymmetry as a function of the distribution of SDSs that contend using different power levels. The assumption underlying this metric is that the demand for the channel use is the same for all SDSs and their spatial distribution is uniform, so the power asymmetry fully accounts for access unfairness. Here we want to understand the extent of unfairness and measure it from two perspectives: the unfairness to a particular class of SDS categorized by its contention power level and then the unfairness across the highway. For these measures, let i be the index of a particular class of SDS, q_i be the

quantity of SDSs in the highway, p_i be its contention power level, and m be the number of power levels, where the SDSs are ordered such that $p_i < p_{i+1}$. The measure of unfairness to an SDS is

$$WPAR_i = \sum_{j=i+1}^m \left(\frac{q_j}{q_i} \right) \left(\left(\frac{p_j}{p_i} \right)^{\frac{2}{n}} - 1 \right). \quad (\text{A-2})$$

Given these measures, the unfairness on the highway is

$$HPAR = \frac{\left(\sum_{i=1}^{m-1} (q_i) WPAR_i \right)}{\sum_i q_i}. \quad (\text{A-3})$$

A.9.3 Directional Power Asymmetry¹⁴

Direction power asymmetry like omnidirectional power asymmetry captures the ratio of additional area covered by the higher power directional signals over the area covered by the lower powered signals, all assumed to be omnidirectional. To simplify the computation we assume the main beam in the computation has a uniform gain across its beam width, θ . The directional antenna gain, G_D , affects receiving the signals and so the low power coverage in the beam has greater range than when both ends of the signaling use omnidirectional antennas with gain G_O . The directional power asymmetry risk measure is:

$$DPAR = \frac{\theta/360 \cdot \left([G_D \cdot G_O \cdot P_h]^{\frac{2}{n}} - [G_D \cdot G_O \cdot P_l]^{\frac{2}{n}} \right)}{[G_O^2 \cdot P_l]^{\frac{2}{n}}} \quad (\text{A-4})$$

A.9.4 Adaptive Directional Asymmetry

Adaptive directional asymmetry is a measure of unfairness that occurs when directional antennas adapt their pointing direction as a function of their use, e.g. point to the intended destination node. In this situation, destinations with adaptive antennas have to listen for and echo signals omnidirectionally in case other terminals are trying to reach them. The measure is the fraction of SDSs on a highway that are omnidirectional or fixed directional SDSs. The assumption is that all other SDSs use adaptive directional access.

A.10 Contention versus Device Asymmetry

Contention versus device asymmetry does not affect the fairness in contention but possibly the performance of the device. When the device range is less than the contention range then there is no adverse effect on the device's performance. When the opposite is true, degradation could

¹⁴ This and other directional measures assume a 2-dimensional rather than 3-dimensional space. Omnidirectional antennas on a 2-dimensional plane usually have a gain greater than 0 dBi, for example the gain of a half wave dipole is 2.14 dBi.

occur. The use of echoes in contention signaling will tend to mitigate the problem in most cases other than those where the device's range is significantly larger than the contention range. The measure is the ratio of harmful collisions resulting from this asymmetry that occur on a highway to the number of contentions winners total. Obtaining a value for this measure of performance will likely require using M&S. In situations where designs are compared using this measure, an M&S scenario and measurement approach will have to be defined.

A.11 Multifunction Count

The goal of a highway is to provide a means for a diverse set of SDSs and users to share the same spectrum dynamically yielding access to the uses that are most urgent. Differentiation in use is associated with duration, persistence, bandwidth, and volume of use. The multifunction count captures the quantity of possible alternative combinations of these characteristics in the highway.

A.12 Collision Resolution

The collision resolution metric captures the ability of a highway to handle congestion and to ensure its capacity is fully used. The method used to measure the effectiveness of collision resolution depends on the method used to arbitrate access among contending users. In the case of SCR and CRS it is a design choice with an effectiveness that depends on the number of signaling phases assigned to CRS. Section 3.3.3.2 describes the design of CRS contention. There are two inputs to the design: number of phases and design density, (i.e., the upper quantity of contenders that CRS must resolve to one). Two CRS designs with different quantities of phases but the same design density can be compared using the probability of one survivor at the design density.

A.13 Precedence Count

The precedence count captures the quantity of differentiation levels that are provided by the access design. Differentiation can be based on three criteria; user, use, and quality of service (QoS).

- User differentiation gives precedence to one user over another because of who they are. It is a precedence that stays with the SDS and does not change. T
- Use differentiation gives precedence to an SDS because of the particular use involved. For example, consider a multipurpose system, say a combined jammer and communications system. In practice, the jammer use may get higher precedence than the communications use when the system is responding to a perceived threat.
- Finally, QoS precedence serves to support differentiation of performance within systems. For example, networks frequently differentiate between traffic categories such as flash override, flash, priority, and routine traffic. Reservations for streaming services are also a type of QoS precedence. Complementing reservations is a counter precedence to preempt the reservations. If the precedence level varies by lanes then the measure is a sum of the precedence count of each lane weighted by the ratio of each lane's band width to the total bandwidth of all lanes, even if they overlap and preclude each other.

A.13.1 Reservation Count

This measure is a count of lanes and time periods that allow reservations. If a lane supports the reservation of frames then its count is the number of frames in an epoch. If it supports the reservation of timeslots then it is the number of timeslots in a frame, and if it does not support a reservation its measure is 0. The ability to contend for a frame or epoch on a lane, however, can make it impractical to use timeslot reservations. Similarly a frame reservation is impractical when SDSs can contend for an epoch. To enable both, special measures would have to be in place to prevent contentions for long-duration access when a shorter duration transmission opportunity reservation is in effect. This may not be practical. Note, however, practicality does not affect the count.

A.14 Resilience

The resilience measure captures the contention mechanism’s vulnerability to countermeasures. The quantity used for this measure is the product of bandwidth and power spectral flux density of random noise that a jammer must achieve at the receiver to disrupt the contention when the separation of signaler and receiver is half the range of the signals with ambient noise alone. This measure depends mostly on the design of the signaling channel.

The resilience computation for a directional device uses the same product of the bandwidth and power spectral flux density as used by the omnidirectional antennas in the main beam direction but then increases in the directions outside the main beam by the difference in gain from the standard omnidirectional antenna and an assumed gain “ball” in the out-of-beam directions assuming the ball has the gain of the largest sidelobe of the directional antenna. An adversary will need greater power to attack the contention in the out-of-beam directions. Given the omnidirectional resilience measure of R_O , the omnidirectional gain of G_O , a main beam width of θ , and ball gain of $G_{1st-lobe}$, the adjusted resilience measure of the directional device is

$$R_D = \left(\frac{\theta}{360} + \frac{360 - \theta}{360} \cdot (G_O - G_{1st-lobe}) \right) \cdot R_O \quad (8-5)$$

If some fraction of SDSs use directional antennas then the measure is a weighted sum where the directional device measure is adjusted by the fraction of directional devices and the omnidirectional measure is weighted by the fraction of omnidirectional devices.

A.15 Mobility

Spatial mobility affects the performance of spectrum highways when contention winners can move and affect the performance of another contention winner because of that movement. The likelihood of this problem will be a function of the speed of the SDSs and the duration of time that spectrum use follows from the time a contention is won. Thus the measure used for mobility is the duration of time from the end of a contention to the completion of the use. A highway design that allows contention for timeslots, frames, and epochs or one that consolidates contentions will have multiple metrics for each alternative contention – transmission opportunity outcome. In practice, there must also be a mapping of allowed physical speeds to mobility measures. A fast-moving device may not be permitted to operate on a lane that allows long-term

transmission opportunities if there is a possibility the device will move far enough to interfere with a neighboring user during that transmission opportunity. Also, on lanes with multiple measures because of variable timeslot size or consolidated contentions, a fast-moving device may be permitted to operate on only particular transmission opportunities.

Appendix B Examples of In-Band Signaling Designs

This appendix describes an example of a highway access design that uses in-band signaling. This access design supports the eight-lane highway illustrated in Figure 4-2, the contention transmission opportunity designs illustrated in Figure 4-3 for per-slot contention and the frame consolidated contention approaches illustrated in Figure 4-9 and Figure 4-10.

An SCM Authorization Set is used to define a highway. There is a transmitter model for each lane of the highway so in the eight lane case used in this example, there would be eight transmitter models that each define the boundaries of one of the lanes. The access design governed by the rules provided in Chapter 5 is found in the SCMPolicyOrProtocol construct of the transmitter model. Each transmitter model uses the same protocol, “SCRHighwayAccess,” but the designs of access for each lane may vary and so are described separately. Though described separately, the access designs must operate in a complementary way as previously described.

B.1 Access Rule Policy

The SCMPolicyOrProtocol construct of SCMs is used to define the contention design and the rules for access to each lane. It begins with the PorPName which identifies the protocol or policy associated with the rest of the model. In this case the PorPName is “SCRHighwayAccess” to indicate we are using the protocol described in Chapter 5. The next sub-element of the SCMPolicyOrProtocol construct is the structure PorPPParameters. It consists of an unspecified number of parameters that are known and understood by SDSs that are designed to participate on the highway using the “SCRHighwayAccess” protocol.

Table B-1 lists the two sub-elements of the SCMPolicyOrProtocol construct that are part of the transmitter model of each lane.

Table B-1. Highway Access Policy Description

Primary Data Element	Sub - elements	Value	Notes
SCMPolicyOrProtocol			
	PorPName	SCRHighwayAccess	This is the name given to the design and complementary rule sets described in this report. An SDS designed to operate on a spectrum highway will understand what to look for in the subsequent elements to know how to behave on the highway.
	PorPPParameters	No value, a series of “Parameter” follow that define the access	The PorPPParameters construct is a structure with two sub-elements. Only the Parameter sub-element is used but it is used many times over to define the contention design.

The PorPPParameters define the specifics of the access design of each lane. In expressing the details, each “Parameter” has a name and optionally a type and value. Parameters are to be provided in the order used to describe the data elements in Chapter 5: Timing Hierarchy,

Contention Method, Lane Use Precedence Phase Set, Lane Selection Phase Set, Priority Phase Set, Contention Phase Set, Contender ID Phase Set, Contention Channels, and Range Symmetry.

B.2 Timing Hierarchy and Contention Method

The start of the rules defines the timing used for the lanes and the contention method. All lanes of a highway will have the same timing hierarchy and the same contention method. The contention method may be per slot or consolidated per frame. Table B-2 provides the data for defining a per-slot in-band contention and Table B-3 provides the data for defining a consolidated contention per frame in-band contention. These lists of parameters would be the same for all lane models of the same highway.

Table B-2. PorPPParameters for the Highway Timing Hierarchy (in-band per timeslot contention)

Primary Data Element	Sub-elements	Value	Notes
PorPPParameters			
Parameter 1			
	PPPName	TimeReference	
	PPPTYPE	TIME	
	PPPValue	2018,12,01,24,00,0,-05,00	UTC Time for 2400 1 December 2018, EST
Parameter 2			
	PPPName	Frame	
	PPPTYPE	INTEGER	
	PPPValue	10	
Parameter 3			
	PPPName	Timeslot	
Parameter 4			
	PPPName	TimeslotDuration	
	PPPTYPE	NUMBER	
	PPPValue	50000000	In nanoseconds, so 50 milliseconds duration
Parameter 5			
	PPPName	TimeslotRange	
	PPPTYPE	NUMBER	
	PPPValue	15000	15 kilometers is a 50-microsecond duration
Parameter 6			
	PPPName	Epoch	
	PPPTYPE	INTEGER	
	PPPValue	4	
Parameter 7			
	PPPName	ContentionMethod	
	PPPTYPE	STRING	
	PPPValue	PERSLOT	
Parameter 8			
	PPPName	PerSlotDuration	
	PPPTYPE	NUMBER	
	PPPValue	10320000	Signaling ends at 1305000 and then 15000 nanoseconds for transition to a transmission start

Table B-3. PorPPParameters for the Highway Timing Hierarchy (in-band consolidated contention)

Primary Data Element	Sub-elements	Value	Notes
PorPPParameters			
Parameter 1			
	PPPName	TimeReference	
	PPPType	DATETIME	
	PPPValue	2018,12,01,24,00,0,-05,00	UTC Time for 2400 1 December 2018, EST
Parameter 2			
	PPPName	Frame	
	PPPType	INTEGER	
	PPPValue	10	
Parameter 3			
	PPPName	Timeslot	
Parameter 4			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	39689000	In nanoseconds, so 39.6905 milliseconds duration
Parameter 5			
	PPPName	TimeslotRange	
	PPPType	NUMBER	
	PPPValue	15000	15 kilometers is a 50-microsecond duration
Parameter 6			
	PPPName	Epoch	
	PPPType	INTEGER	
	PPPValue	4	
Parameter 7			
	PPPName	ConsolidatedSignalingStart	
	PPPType	INTEGER	
	PPPValue	0	
Parameter 8			
	PPPName	ConsolidatedSignalingDuration	
	PPPType	NUMBER	
	PPPValue	103110000	$(10 \times 10305000) + (9 \times 5000)$, a 5000 nanosecond pause is added between contentions and a 15000-nanosecond pause after the last contention.
Parameter 9			
	PPPName	ContentionMethod	
	PPPType	STRING	
	PPPValue	CONSOLIDATED	

The difference between these designs is the value for ContentionMethod, either PERSLOT or CONSOLIDATED, and the presence of the parameters ConsolidatedSignalingStart and ConsolidatedSignalingDuration when consolidated contention per frame is used. The timeslot durations are different to achieve designs with the same frame duration. The time consumed for contention is part of the timeslot duration in the per timeslot contention design and is explicitly defined by PerSlotDuration. The time used for contention is the ConsolidatedSignalingDuration

when using consolidated contention. In both cases, the transmission opportunity is (50 – 10.305) = 39.695 msec and the frame is 500 msec long.

In the designs of Table B-2 and Table B-3 the timeslots are all identical. It is also possible to define timeslots in a frame to be of different lengths. In this case, the durations of each timeslot are provided in the order that they occur. Table B-4 provides an example of a consolidated in-band contention design with varying timeslot durations. There is an entry for each timeslot of a frame. This design has the same frame period as that of Table B-3 but with fewer timeslots and a shorter ConsolidatedSignalingDuration since there are fewer contentions. The shorter ConsolidatedSignalingDuration results in more time for the transmission opportunities in aggregate than the ten timeslot frame.

Table B-4. PorPPParameters for the Highway Timing Hierarchy (in-band consolidated contention with varying timeslot durations)

Primary Data Element	Sub-elements	Value	Notes
PorPPParameters			
Parameter 1			
	PPPName	TimeReference	
	PPPType	DATETIME	
	PPPValue	2018,12,01,24,00,0,-05,00	UTC Time for 2400 1 December 2018, EST
Parameter 2			
	PPPName	Frame	
	PPPType	INTEGER	
	PPPValue	6	
Parameter 3			
	PPPName	Timeslot	
Parameter 4			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	39532500	39.5325 milliseconds duration
Parameter 5			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	110000000	110 milliseconds duration
Parameter 6			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	39532500	39.5325 milliseconds duration
Parameter 7			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	170000000	170 milliseconds duration
Parameter 8			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	39532500	39.5325 milliseconds duration
Parameter 9			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	39532500	39.5325 milliseconds duration

Primary Data Element	Sub-elements	Value	Notes
Parameter 10			
	PPPName	Epoch	
	PPPType	INTEGER	
	PPPValue	4	
Parameter 11			
	PPPName	ConsolidatedSignalingStart	
	PPPType	INTEGER	
	PPPValue	0	
Parameter 12			
	PPPName	ConsolidatedSignalingDuration	
	PPPType	NUMBER	
	PPPValue	61870000	(6 × 10305000) + (5 × 5000), a 5000-nanosecond pause is added between contentions and a 15000-nanosecond pause after the last contention.
Parameter 13			
	PPPName	ContentionMethod	
	PPPType	STRING	
	PPPValue	CONSOLIDATED	

B.3 Lane Use Precedence Phase Set

The lane use precedence phase set is used to allow certain devices to have precedence in the ability to establish the active lanes. This design will be the same for all lanes of a highway. Table B-5 illustrates the parameters that define the phase set for each of the lanes.

Table B-5. PorPPParameters Defining the Lane Use Precedence Phase Set

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter LU1								
PPPName	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	2	2	2	2	2	2	2	2
Parameter LU1								
PPPName	Override	Override	Override	Override	Override	Override	Override	Override
Parameter LU1								
PPPName	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

This design establishes two phases for asserting precedence. The highest signaling level (i.e., signal in all phases) is reserved for override, so that devices with urgent use, as established by

the rules, may override the routine precedence differentiation. As a result, this design has three routine precedence levels and the override precedence level.

B.4 Lane Selection Phase Set

The lane selection phase set signaling design is the same for all lanes. The signaling that results in a lane becoming active differs by lane. The parts of the lane selection phase set description that concern the number of phases, the pause that precedes this phase set, and the channel used for selecting a lane are the same for all lanes of the same highway. The signaling that results in a lane becoming active may vary. Table D-6 illustrates that these differences center around what signaling is used to select a lane and the observed signals that indicate a lane is active and whether the observing device can contend to gain access to that lane. As seen in this table, the signaling and the number of signal combinations that result in a lane being active varies by lane.

Table B-6. PorPPParameters Defining the Lane Precedence Phase Set

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter LP1								
PPPName	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter LP2								
PPPName	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter LP3								
PPPName	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	000	000	000	000	001	010	100	110
Parameter LP4								
PPPName	LaneOverride	LaneOverride	LaneOverride	LaneOverride	LaneOverride	LaneOverride	ContentionEligibleOutcome	ContentionEligibleOutcome
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	100	100	100	110	100	100	100	110
Parameter LP5								
PPPName	LaneOverride	LaneOverride	LaneOverride	LaneOverride	ContentionEligibleOutcome	ContentionEligibleOutcome	ActiveOnlyOutcome	ActiveOnlyOutcome
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	001	001	010	010	001	010	101	111
Parameter LP6								
PPPName	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	LSChannel	LSChannel
	1	1	1	1	1	1	1	1

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	STRING	STRING
PPPValue	000	000	000	000	011	011	C1	C1
Parameter LP7								
PPPName	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	LSChannel1	LSChannel1		
PPPType	BINARY	BINARY	BINARY	BINARY	STRING	STRING		
PPPValue	010	010	001	100	C1	C1		
Parameter LP1								
PPPName	LSChannel1	LSChannel1	LSChannel1	ContentionEligibleOutcome				
PPPType	STRING	STRING	STRING	BINARY				
PPPValue	C1	C1	C1	001				
Parameter LP1								
PPPName				ContentionEligibleOutcome				
PPPType				BINARY				
PPPValue				101				
Parameter LP1								
PPPName				LSChannel1				
PPPType				STRING				
PPPValue				C1				

There are several key observations. Each lane may have a different number of parameters listed and different types of parameters. Each lane has a signaling sequence that a device would use to establish the lane as active, called LaneSignalingSequence. Each lane definition may provide signaling that indicates the lane cannot be won so cease contending for it, called LaneOverride. Each lane definition provides one or more signaling outcomes that indicate a lane is active and that the device may contend for access to that lane, called ContentionEligibleOutcome. Some lane definitions may provide a signaling combination that indicates that the lane is active but that the device hearing that sequence may not contend for the lane itself, called ActiveOnlyOutcome. Of note in the table are the number of ContentionEligibleOutcomes for Lane 4. This is because there are many lanes that can be active that do not overlap Lane 4: specifically, Lanes 1, 2, 3, 5, and 7. Lane 7 and Lane 8 have the fewest parameters because there is one and only one signaling outcome that permits contention for these lanes and no signaling that will override their selection. Given this collection of signaling outcomes for all lanes, devices can use their observations to determine all the lanes that are active in their location.

B.5 Priority Phase Set

The priority phase set accomplishes many goals. It differentiates which user or use has precedence among users of devices on the highway and which traffic has precedence among devices of the same or different SDSs. It is also used to contend for frames and epochs and to establish reservations. In the highway that has been designed so far, where lanes can be merged

for use, contentions for frames or epochs and reservations can be very confusing and places a burden on devices to track the status of lanes. Table B-7 provides a simple per-slot signaling design where there are no contentions for frames or epochs or reservations on any of the lanes. Part of the challenge of allowing frame and epoch contention with per-slot contention is that if a lane is captured for a frame or epoch then there should be no signaling in a contention channels during that frame or epoch for the lanes covered by the bandwidth used by the captured lane. Therefore in this lane design, allowing frames or epochs to be captured on lanes that cover the bandwidth of the signaling channel used for the lane use precedence and lane selection phase sets would disrupt the performance of the contentions of other lanes that may co-exist.

Table B-7. PorPPParameters Defining the Priority Phase Set (per-slot signaling)

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPTType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	6	6	6	6	6	6
Parameter PP2								
PPPName	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse	PriorityPhaseUse
PPPTType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP3								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPTType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP4								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPTType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP4								
PPPName	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel
PPPTType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C2	C3	C4	C1	C3	C1	C1

As a potential alternative, one that supports frame reservation on the highway, Table D-8 shows a design that allows frame contention on Lanes 3, 4, and 6. None of the other lanes require the use of Lane 3, 4, or 6 bandwidth for any part of their contention. This design does not support all the time agility it seems to imply. The limitation is that Lane 8 cannot be used during any of these longer term contentions and Lane 7 cannot be used when there is longer term access for Lanes 3 or 6. It is necessary for devices with that intent to access Lane 7 or 8 to track the state of these overlapping lanes. This is not practical unless devices that use Lane 7 or 8 listen to the contentions on channels C3 and C4 to know when this type of access is made. For this reason the rules require devices with larger bandwidth access needs in per-slot contention to defer their

access until the contention where frames are contended for. If a device wins that contention for the overlapping broadband channel, then and only then may it contend for a broadband channel in timeslots of the next frame up until the next frame contention.

Table B-8. PorPPParameters Defining the Priority Phase Set (per-slot signaling with long duration contentions)

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	6	6	6	6	6	6
Parameter PP2								
PPPName	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP3								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP4								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP5								
PPPName	PChannel	PChannel	FrameContention	FrameContention	PChannel	FrameContention	PChannel	PChannel
PPPType	STRING	STRING			STRING		STRING	STRING
PPPValue	C1	C2			C1		C1	C1
Parameter PP6								
PPPName			PChannel	PChannel		PChannel		
PPPType			STRING	STRING		STRING		
PPPValue			C3	C4		C3		

With consolidated contention, no transmissions take place during the contentions. Devices that gain access for a frame can assert their access across the contentions of the frame. Table B-9 shows a possible per-frame signaling design. This design allows three types of long term access: timeslot reservation, frame access, and frame reservation.

The expectation in this design is that those devices that contend and win access to a whole frame will subsequently contend in the follow-on contentions of the frame using the override lane use precedence to preserve the designated lane as active, and would use the frame priority to

preserve their access to the frame and to cause other devices in their vicinity to defer to another frame for their timeslot contentions.

This design also supports timeslot reservations and frame reservations. Timeslot reservations and any type of frame contention on the same lane are mutually exclusive. Thus, a device wanting to make a timeslot reservation on a lane must wait for a frame in which no frame contention occurs. When frame reservations are permitted, then the device must wait an Epoch number of consecutive frames with no frame contention before making a timeslot reservation. The use of the ReservationPrecedence parameter and the rule it invokes prevents a frame contention from interrupting a timeslot reservation.

Table B-9. PorPPParameters Defining the Priority Phase Set (per-frame signaling with long duration contentions)

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	6	6	6	6	6	6
Parameter PP2								
PPPName	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause	PriorityPause
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP3								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	2	2	2	2	2	2	2	2
Parameter PP4								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP5								
PPPName	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations
Parameter PP6								
PPPName	FrameContention	FrameContention	FrameContention	FrameContention	FrameContention	FrameContention	FrameContention	FrameContention
Parameter PP7								
PPPName	FrameReservation	FrameReservation	FrameReservation	FrameReservation	FrameReservation	FrameReservation	FrameReservation	FrameReservation
Parameter PP8								
PPPName	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence

Sub-elements	Value							
	1	2	3	4	5	6	7	8
Parameter PP9								
PPPName	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C2	C3	C4	C1	C3	C1	C1

A highway design without overlapping lanes is much more supportive of long duration contentions for frames and epochs. Turning away from the previous highway design, consider a highway that only has Lanes 1, 2, 3, and 4. There would be no need for a lane use precedence phase set nor a lane selection phase set. The active lanes do not change. A possible design could be as shown in Table B-10. This is feasible for both per-slot and consolidated contentions. In this design, with all possible long duration contentions, devices would have to monitor the contentions for timeslots, frames, and epochs and follow the rules before making these types of long duration contentions themselves

Table B-10. PorPPParameters Defining the Priority Phase Set with no Overlapping Lanes (with long-duration contentions)

Sub-elements	Value			
	1	2	3	4
PorPPParameters				
Parameter LU1				
PPPName	LaneUsePrecedencePhases	LaneUsePrecedencePhases	LaneUsePrecedencePhases	LaneUsePrecedencePhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	0	0	0	0
Parameter LS1				
PPPName	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	0	0	0	0
Parameter PP1				
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	6	6
Parameter PP2				
PPPName	PriorityPause	PriorityPause	PriorityPause	PriorityPause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000
Parameter PP3				
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	2	2	2	2
Parameter PP4				
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3
Parameter PP5				
Parameter	TimeslotReservations	TimeslotReservations	TimeslotReservations	TimeslotReservations
Parameter PP6				
Parameter	FrameContention	FrameContention	FrameContention	FrameContention

Sub-elements	Value			
	1	2	3	4
Parameter PP7				
Parameter	EpochContention	EpochContention	EpochContention	EpochContention
Parameter PP8				
Parameter	FrameReservation	FrameReservation	FrameReservation	FrameReservation
Parameter PP9				
Parameter	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence	ReservationPrecedence
Parameter PP10				
PPPName	PChannel	PChannel	PChannel	PChannel
PPPType	STRING	STRING	STRING	STRING
PPPValue	C1	C2	C3	C4

B.6 Contention Phase Set

The contention phase set is one of the simpler ones to define. Table B-11 provides a possible design for this eight lane highway. Of interest is that lanes have the same number of contention phases but use different contention designs. The wider bandwidth lanes assume fewer contenders and use lower density designs. Using the lower density designs improves the probability of having one survivor at the end of the contention when the number of contending devices is less than that design density.

Table B-11. PorPPParameters Defining the Contention Phase Set

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter C1								
PPPName	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	7	7	7	7	7	7	7	7
Parameter C1								
PPPName	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter C1								
PPPName	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	1000	1000	1000	1000	200	200	50	50
Parameter C1								
PPPName	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C2	C3	C4	C1	C3	C1	C1

B.7 Contention Channels

The previous tables identified four different contention channels for the phase sets, C1, C2, C3, and C4. These channels are described in this, the contention channel description part of the lane specification. Since our design uses the Lane Use Precedence Phase Set and the Lane Selection Phase Set, every lane would include the description of Channel C1. In addition Lane 2 would have Channel C2 defined, Lanes 3 and 6 would have Channel C3 defined, and Lane 4 would have Channel C4 defined. Currently there are no particular designs for signaling. It is anticipated that once designs are established most details can be collectively identified by a name. Additional parameters may be provided for variable parts of the design. For example, if the signaling frequency hops a key, the bandwidth of hopping, and a tuning interval might be provided.

As an example a 25 kHz pulse was designed for this type of signaling. In the physical layer specification of this design, the sizes of fields in the time domain are expressed as a number of time units $T_s = 1/1,200,000$ seconds, which is a period of approximately 0.83 microseconds, corresponding to a sampling rate of exactly 1,200,000 samples per second. The contention pulse slot in the design is precisely $155 * T_s$. This corresponds to a period of approximately 129.167 microseconds. The pulse was designed to be minimum phase to maximize the energy near the beginning of the signal minislot. The sample values of the pulse are shown in Table B-12; Figure B-1 illustrates its shape and the anticipated reception at a receiver. The length of the pulse is longer than the length of the minislot: therefore, a receiver must stop detecting before the entire minislot pulse has arrived. Additionally, when frequency hopping is used between minislots, it is possible that there will be some time during which the receiver must tune to the new hop frequency. This is known as tuning delay and it causes the receiver to stop detecting even earlier, as illustrated in Figure B-1.

Table B-12. Contention Pulse Sample Values

h(0) = 0.000882	h(39) = 0.084091	h(78) = 0.124470	h(117) = 0.053727
h(1) = 0.001302	h(40) = 0.086908	h(79) = 0.123767	h(118) = 0.050863
h(2) = 0.001812	h(41) = 0.089667	h(80) = 0.122991	h(119) = 0.048453
h(3) = 0.002420	h(42) = 0.092358	h(81) = 0.122142	h(120) = 0.046068
h(4) = 0.003132	h(43) = 0.094975	h(82) = 0.121222	h(121) = 0.043713
h(5) = 0.003952	h(44) = 0.097509	h(83) = 0.120232	h(122) = 0.041393
h(6) = 0.004885	h(45) = 0.099954	h(84) = 0.119175	h(123) = 0.039113
h(7) = 0.005936	h(46) = 0.102303	h(85) = 0.118052	h(124) = 0.036875
h(8) = 0.007104	h(47) = 0.104552	h(86) = 0.116864	h(125) = 0.034685
h(9) = 0.008391	h(48) = 0.106694	h(87) = 0.115614	h(126) = 0.032547
h(10) = 0.009798	h(49) = 0.108727	h(88) = 0.114305	h(127) = 0.030464
h(11) = 0.011323	h(50) = 0.110648	h(89) = 0.112937	h(128) = 0.028440
h(12) = 0.012963	h(51) = 0.112456	h(90) = 0.111513	h(129) = 0.026477
h(13) = 0.014718	h(52) = 0.114148	h(91) = 0.110035	h(130) = 0.024580
h(14) = 0.016582	h(53) = 0.115726	h(92) = 0.108506	h(131) = 0.022749
h(15) = 0.018550	h(54) = 0.117629	h(93) = 0.106927	h(132) = 0.020989
h(16) = 0.020618	h(55) = 0.118806	h(94) = 0.105302	h(133) = 0.019300
h(17) = 0.023395	h(56) = 0.119911	h(95) = 0.103633	h(134) = 0.017685
h(18) = 0.025595	h(57) = 0.120944	h(96) = 0.101922	h(135) = 0.016145

$h(19) = 0.027885$	$h(58) = 0.121904$	$h(97) = 0.100172$	$h(136) = 0.014680$
$h(20) = 0.030259$	$h(59) = 0.122788$	$h(98) = 0.098385$	$h(137) = 0.013292$
$h(21) = 0.032713$	$h(60) = 0.123595$	$h(99) = 0.096564$	$h(138) = 0.011605$
$h(22) = 0.035241$	$h(61) = 0.124326$	$h(100) = 0.094711$	$h(139) = 0.010343$
$h(23) = 0.037839$	$h(62) = 0.124978$	$h(101) = 0.092305$	$h(140) = 0.009159$
$h(24) = 0.040500$	$h(63) = 0.125550$	$h(102) = 0.090192$	$h(141) = 0.008053$
$h(25) = 0.043217$	$h(64) = 0.126043$	$h(103) = 0.088022$	$h(142) = 0.007026$
$h(26) = 0.045985$	$h(65) = 0.126456$	$h(104) = 0.085795$	$h(143) = 0.006080$
$h(27) = 0.048796$	$h(66) = 0.126787$	$h(105) = 0.083515$	$h(144) = 0.005212$
$h(28) = 0.051643$	$h(67) = 0.127038$	$h(106) = 0.081185$	$h(145) = 0.004422$
$h(29) = 0.054520$	$h(68) = 0.127207$	$h(107) = 0.078808$	$h(146) = 0.003708$
$h(30) = 0.057419$	$h(69) = 0.127294$	$h(108) = 0.076390$	$h(147) = 0.003069$
$h(31) = 0.060332$	$h(70) = 0.127300$	$h(109) = 0.073936$	$h(148) = 0.002502$
$h(32) = 0.063251$	$h(71) = 0.127225$	$h(110) = 0.071451$	$h(149) = 0.002005$
$h(33) = 0.066170$	$h(72) = 0.127069$	$h(111) = 0.068941$	$h(150) = 0.001574$
$h(34) = 0.069081$	$h(73) = 0.126833$	$h(112) = 0.066413$	$h(151) = 0.001205$
$h(35) = 0.071975$	$h(74) = 0.126517$	$h(113) = 0.063874$	$h(152) = 0.000894$
$h(36) = 0.074847$	$h(75) = 0.126121$	$h(114) = 0.061328$	$h(153) = 0.000636$
$h(37) = 0.078321$	$h(76) = 0.125648$	$h(115) = 0.058783$	$h(154) = 0.000427$
$h(38) = 0.081226$	$h(77) = 0.125097$	$h(116) = 0.056248$	

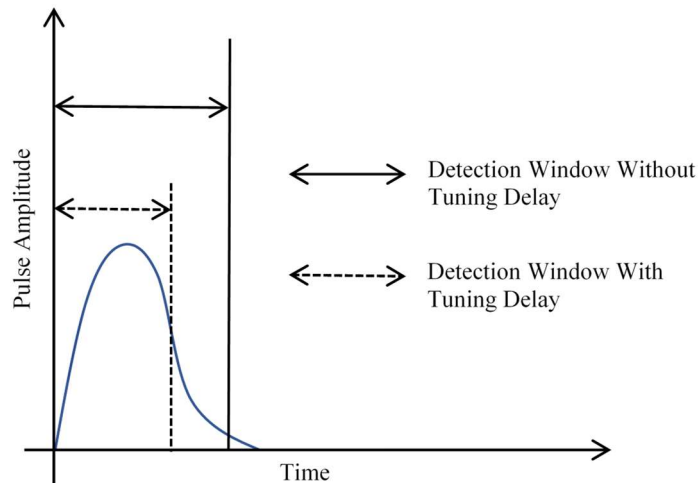


Figure B-1. Contention Pulse

Figure B-2 illustrates a contention phase where we assume it is preceded by a frequency hop. It allocates a period for frequency transition at the beginning and a period for the transition between signals.

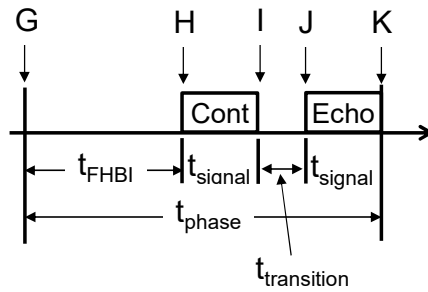


Figure B-2. Contention Phase

The signals are pulses of energy and receivers simply try to detect their presence. There is no period allocated for propagation. Delayed signals decrease the time available for receivers to detect their presence. In this design it was desirable that contending radios ignore long delayed signals. Signal receiver designs shall seek to reliably detect signals that are delayed less than 30 μsec (9 km propagation time) and reject signals that are delayed more than 45 μsec (13.5 km propagation time). Faster frequency hops and transitions increase the time available for detection but they shall not affect the times at which signals are transmitted. Times G, H, I, J, and K are requirements. Time G matches the start of the contention hop. Time H is the start of the contention signal. Time I is the end of the contention signal. Time J is the start of the echo signal. Time K is the end of the echo signal and the end of the contention phase. The requirement for the signal transmissions is that the times H, I, J, and K shall be met at the output of the antenna of each transmitting radio based on its understanding of time. Table D-13 lists both the durations of activities within a contention phase and the specific times relative to the start of the phase when events occur.

Table B-13. Contention hop activity durations and event start times

Parameter	Samples @ 1.2 MHz	Duration (approx.)
t_{FHBI}	300	250.000 μs
t_{signal}	155	129.167 μs
$t_{\text{transition}}$	76	63.333 μs
t_{signal}	155	129.167 μs
Total (t_{phase})	686	571.667 μs

Event	Sample Index	Time
G	0	0 μs
H	300	250.000 μs
I	455	379.167 μs
J	531	442.500 μs
K	686	571.667 μs

Let us name this phase design the JS25. Given the JS25 phase design, a possible set of parameters to define a channel would be the start and stop tuning frequency, the tuning interval, the key for choosing the frequency hop sequence, a reference start time for hopping to begin and a reset interval. The reset interval identifies a time interval when the pseudo random number generator is reset with the key and some other changing parameter related to the interval count from the reference start time. This interval time is defined as some number of timeslots.

Table B-14 provides an example set of contention channels definitions using the JS25 contention phase design. In this example, Channel C1 is defined for every lane since it will be used by contenders for all lanes for at least the lane use precedence and lane selection phase sets. The contention phase lasts 571667 nanoseconds, the phase design is JS25, the start frequency is 400.015 MHz, and the stop frequency is 400.985 MHz with a 5 kHz tuning interval. (Each channel has 195 individual tuning frequencies.) The key for the pseudo random number generator is the 64 bit number in hex format, A5F1890B3439CD54. The start time is 2400 1 December 2018 EST and the reset interval is every 200 timeslots, which corresponds to once every five epochs. Lanes 2, 3, 4, and 6 each have a second channel defined because they use different channels for the User/Use Precedence Phase Set and the Contention Phase Set than they use in the Lane Use Precedence Phase Set and the Lane Selection Phase Set. These second channel definitions differ in name, frequency hopping range, and key.

Table B-14. PorPPParameters Defining the Contention Channels

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter CC1								
PPPName	Channel	Channel	Channel	Channel	Channel	Channel	Channel	Channel
Parameter CC2								
PPPName	ChannelName	ChannelName	ChannelName	ChannelName	ChannelName	ChannelName	ChannelName	ChannelName
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1
Parameter CC3								
PPPName	PhaseDuration	PhaseDuration	PhaseDuration	PhaseDuration	PhaseDuration	PhaseDuration	PhaseDuration	PhaseDuration
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	571667	571667	571667	571667	571667	571667	571667	571667
Parameter CC4								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC5								
PPPName	PhaseDesign	PhaseDesign	PhaseDesign	PhaseDesign	PhaseDesign	PhaseDesign	PhaseDesign	PhaseDesign
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	JS25	JS25	JS25	JS25	JS25	JS25	JS25	JS25
Parameter CC6								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC7								
PPPName	StartFreq	StartFreq	StartFreq	StartFreq	StartFreq	StartFreq	StartFreq	StartFreq
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	400015	400015	400015	400015	400015	400015	400015	400015
Parameter CC8								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC9								
PPPName	StopFreq	StopFreq	StopFreq	StopFreq	StopFreq	StopFreq	StopFreq	StopFreq

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	400985	400985	400985	400985	400985	400985	400985	400985
Parameter CC10								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC12								
PPPName	TuneInt	TuneInt	TuneInt	TuneInt	TuneInt	TuneInt	TuneInt	TuneInt
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5	5	5	5	5	5	5	5
Parameter CC13								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC14								
PPPName	Key	Key	Key	Key	Key	Key	Key	Key
PPPType	HEX	HEX	HEX	HEX	HEX	HEX	HEX	HEX
PPPValue	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54	A5F1890B3439CD54
Parameter CC15								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC16								
PPPName	StartTime	StartTime	StartTime	StartTime	StartTime	StartTime	StartTime	StartTime
PPPType	DATETIME	DATETIME	DATETIME	DATETIME	DATETIME	DATETIME	DATETIME	DATETIME
PPPValue	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00	2018,12,01,24,00,0,-05,00
Parameter CC17								
PPPName	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters	ChannelParameters
Parameter CC18								
PPPName	Interval	Interval	Interval	Interval	Interval	Interval	Interval	Interval
PPPType								
PPPValue	40	40	40	40	40	40	40	40
Parameter CC19								
PPPName		Channel	Channel	Channel		Channel		
Parameter CC20								
PPPName		ChannelName	ChannelName	ChannelName		ChannelName		
PPPType		STRING	STRING	STRING		STRING		
PPPValue		C2	C3	C4		C3		
Parameter CC21								
PPPName		PhaseDuration	PhaseDuration	PhaseDuration		PhaseDuration		
PPPType		NUMBER	NUMBER	NUMBER		NUMBER		
PPPValue		571667	571667	571667		571667		
Parameter CC22								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC23								
PPPName		PhaseDesign	PhaseDesign	PhaseDesign		PhaseDesign		

Sub-elements	1	2	3	4	5	6	7	8
PPPType		STRING	STRING	STRING		STRING		
PPPValue		JS25	JS25	JS25		JS25		
Parameter CC24								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC25								
PPPName		StartFreq	StartFreq	StartFreq		StartFreq		
PPPType		NUMBER	NUMBER	NUMBER		NUMBER		
PPPValue		401015	402015	403015		402015		
Parameter CC26								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC27								
PPPName		StopFreq	StopFreq	StopFreq		StopFreq		
PPPType		NUMBER	NUMBER	NUMBER		NUMBER		
PPPValue		401985	402985	403985		402985		
Parameter CC28								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC29								
PPPName		TuneInt	TuneInt	TuneInt		TuneInt		
PPPType								
PPPValue		5	5	5		5		
Parameter CC30								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC31								
PPPName		Key	Key	Key		Key		
PPPType		HEX	HEX	HEX		HEX		
PPPValue		11BA097C 478D53B2	96D4A9E2 08C4C671	BA54789E C710743A		96D4A9E2 08C4C671		
Parameter CC32								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC33								
PPPName		StartTime	StartTime	StartTime		StartTime		
PPPType		DATETIME	DATETIME	DATETIME		DATETIME		
PPPValue		2018,12,01 ,24,00,0,- 05,00	2018,12,01 ,24,00,0,- 05,00	2018,12,01 ,24,00,0,- 05,00		2018,12,01 ,24,00,0,- 05,00		
Parameter CC34								
PPPName		ChannelParameters	ChannelParameters	ChannelParameters		ChannelParameters		
Parameter CC35								
PPPName		Interval	Interval	Interval		Interval		
PPPType								
PPPValue		40	40	40		40		

B.8 Range Symmetry

The power symmetry is defined for each lane. It is possible for each lane to have a different definition, although signaling through the Lane Use Precedence and Lane Selection Phase Sets affects all lanes and so a fairness issue arises if they are not the same. Table B-15 provides an example design and shows no difference in the lanes. The signal transmit power is 7 dBW, about 5 Watts and the antenna is assumed to have 3 dBi gain. The transmit power may vary by 0.5 dB which corresponds to a maximum power of 5.6 Watts. These combine to determine the range based on the signaling receiver sensitivity.

Table B-15. PorPPParameters Defining the Range Symmetry

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter RS1								
PPPName	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5	5	5	5	5	5	5	5
Parameter RS1								
PPPName	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	0.5	2	2	2	2	2	2	2
Parameter RS1								
PPPName	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	3	3	3	3	3	3	3	3
Parameter RS1								
PPPName	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	15000	15000	15000	15000	15000	15000	15000	15000
Parameter RS1								
PPPName	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	40000	40000	40000	40000	40000	40000	40000	40000
Parameter RS1								
PPPName	TerrestrialRange	TerrestrialRange	TerrestrialRange	TerrestrialRange	TerrestrialRange	TerrestrialRange	TerrestrialRange	TerrestrialRange
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	9000	9000	9000	9000	9000	9000	9000	9000
Parameter RS1								
PPPName	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	0.9	0.9	0.9	0.9	0.9	0.9	0.9	0.9

Despite the ability to derive a range from the power the tables shows three signaling ranges. The first is the Timing Range. This is the range assumed in the design of the signaling to accommodate the propagation of the signal. There is risk the contention signal will be detected beyond that range because of transitions in the transceiver. Since signaling uses echoing, the effect of signaling is usually greater than the timing range. The TerrestrialRange and SquareLawRange provides a maximum propagation range in those two environments. The UseAsymmetry value of 0.9 indicates that the device of the SDS gaining access has a range no more than 0.9 that of the signal range. In this case a SDS device should not seek to be effective beyond 0.9 of 9000 meters in terrestrial applications and 0.9 of 40,000 meters in free space propagation environments. Regardless, the SDS is also constrained by a spectrum mask used for the lanes, and that mask is the overriding limit on SDS transmit power.

B.9 Design Metrics

Design metrics help in comparing highway designs. Some metrics are drastically different depending on the contention band. With in-band signaling, temporal overhead is significant; with out-of-band signaling, spectral overhead is significant. Some metrics require estimates of the characteristics of the SDSs that will operate on the highway (e.g., temporal efficiency). The following subsections describe the metrics applicable to the in-band signaling designs defined above. Many of the metrics listed in Appendix A depend on the characteristics, mix, density, and activity of SDSs that use the highways.; these are not included below.

B.9.1 Temporal Overhead

In the per-slot design there are a total of 18 phases per contention with three spacers of 5,000 nanoseconds. Adding up the durations for the phases of all phase sets and the duration of the spacing between phase sets, the total duration of all contention signaling for a timeslot is 10,305,000 nanoseconds. Adding onto this the 15,000 nanoseconds before the start of the payload, and the 50,000 nanoseconds before the start of the next contention to account for 15 kilometers of propagation, the total overhead is 10,370,000 nanoseconds. Finally, since all timeslots have the same design there is no difference in the ratio of temporal overhead to temporal payload capacity. The temporal overhead measure is

$$\frac{10,370,000}{50,000,000} = 0.2074 \cdot \tag{B-6}$$

The consolidated contention design of Table B-3 reduced the timeslot duration to 39,689,000 nanoseconds to retain a 500,000,000-nanosecond frame that includes the time for the consolidated contention. The consolidated contention has 10 consecutive contentions with a 5,000-nanosecond spacer between contentions and a 15,000-nanosecond break at the end for a total contention duration for a frame of 103,110,000 nanoseconds. The transmission of each slot must end 50,000 nanoseconds before the next for a total of 500,000 nanoseconds lost to slot transition. The temporal overhead measure for this design is

$$\frac{103,110,000 + 500,000}{103,110,000 + 396,890,000} = 0.2072 \cdot \tag{B-7}$$

The consolidated contention design of Table B-4 with time-varying timeslots has fewer timeslots and therefore fewer contentions. This design again retains the 500,000,000 nanosecond frame. The consolidated contention, however has six consecutive contentions with a 5,000 nanosecond spacer between contentions and a 15,000-nanosecond break at the end for a total contention duration for a frame of 61,870,000 nanoseconds. The transmission of each slot must end 50000 nanoseconds before the next for a total of 300000 nanoseconds lost to slot transmission. The temporal overhead measure for this design is

$$\frac{61,870,000 + 300,000}{61,870,000 + (4 \times 39,532,500) + 110,000,000 + 170,000,000} = 0.1360 \cdot \quad (\text{B-8})$$

Any increase in the size of the transmission opportunities will increase the temporal overhead. Those designs that permit frame and epoch contentions will have less temporal overhead but the exact measure is a function of the operation of the highway and the resulting proportion of the epoch, frame, and timeslot access attempts. When long-duration access is used, the per slot design has a temporal overhead advantage because the time used for contention of the intermediate timeslots becomes available for payload transmission.

B.9.2 Spectral Overhead

Spectral loss may be computed as the proportion of a lane that is lost to guard bands between lanes. The bandwidth of the guard band would be measured at some point where the spectrum mask of the lane is some power level below the maximum power level to the point that divides adjacent lanes. Using the lane design of Table 2-10 as an example and the -3 dB point from the maximum power level as the transmission bandwidth, the bandwidth of the signal is 815 kHz while the total spectrum allocated to the lane is 1 MHz. The spectral overhead is

$$\frac{1,000,000 - 815,000}{1,000,000} = 0.1850 \cdot \quad (\text{B-9})$$

The actual spectral overhead may be less if combined lanes are used. The measure, however, would depend on the operational conditions and the resulting proportion of time the broadband lanes are used in lieu of the narrowband lanes.

B.9.3 Use Efficiency

Use efficiency is a measure of the match of the highway design and SDSs that are authorized to operate on the highway and their ability to fit their spectrum use to the limits of the transmission opportunities. This is a measure that should be made when a highway is in operation. It can be estimated by detailed simulation.

If all uses of timeslots in the per-slot signaling design exactly filled the 39,830,000 nanosecond duration allowed by that design and filled the 815 kHz of the lane then the use efficiency would be 1 for that lane. The same would be true if each of the uses were the full allowed bandwidth of the lane and were 39,639,000 nanoseconds long in the aggregate design or within each of the allowed use durations for the six-timeslot design with variable durations or within the long duration frame or epoch slot when they are contended for. It is unlikely that uses will exactly fit within these lanes and timeslots in practice.

B.9.4 Access Delay

The access delay is the average time from the start of the contention to the start of the use of the highway lane. In the per-slot signaling design this is the duration of the signaling and the end-of-signaling pause, so 10,320,000 nanoseconds. The access delay for the consolidated signaling design that had 10 timeslots per frame is

$$\frac{1}{10} * \left(\sum_{n=1}^{10} (103,110,000 + (n-1) \cdot 39,689,000) - (n-1) \cdot 10,310,000 \right) = 235,315,500 \text{ nanoseconds.} \quad (\text{B-10})$$

The computation for the 6 slot frame is similar but is not as tidy to write and so the computation is not provided. Its average access delay is 245,393,750 nanoseconds. Ordering the timeslots from shortest duration to longest duration improves the access delay measure and in this case would be 146,670,833 nanoseconds.

Per-slot signaling has a much shorter access delay than consolidated signaling.

B.9.5 Slot Delay

The slot delay is the time from the start of the use of a timeslot to the time of the contention for the next slot. In the per-slot signaling design it is 39,680,000 nanoseconds. For the 10 slot consolidated signaling design it is

$$\frac{1}{10} * \left(\sum_{n=1}^{10} (396,890,000 + (n-1) \cdot (10,310,000 - 39,689,000)) \right) = 264,684,500 \text{ nanoseconds.} \quad (\text{D-11})$$

The slot delay for the six slot frame design is 254,606,250 nanoseconds.

B.9.6 Responsiveness

The responsiveness measure is the average of the sum of the access delay and the slot delay and therefore is 50,000,000 nanoseconds for the per-slot signaling design and 500,000,000 nanoseconds for both the consolidated signaling designs. The access delay and slot delay are the same for all slots when using per-slot signaling. The access delay and slot delay each vary by ordinal position of the timeslot in a frame when using consolidated contention but the responsiveness remains the same.

B.9.7 Spatial Capacity

Simulation results in [2] demonstrate that the survivor density of CRS is a function of contender density and ranges between 0.5 to 0.75 of the area covered by the range of the signal used in the contention. The density peaks at about five contenders per area covered by a signal and then converges to 0.5 survivors as the density increases. Using the converged value and the TerrestrialRange of the signals in Table D-15 of 9 kilometers, the anticipated spatial capacity is one user per 509 square kilometers. Actual spatial capacity is likely to be larger because signal propagation ranges are less and the contender density will be somewhat less than that where the

survivor density converges to 0.5. There can be further gains if directional antennas are used in contention and SDS demand is less intense.

B.9.8 Spectral Agility

When in-band signaling is used, spectral agility is a function of methods internal to SDSs. It is not a measure of highway design or highway operation.

B.9.9 Asymmetry Risk

The asymmetry risk is a measure of the extent to which a particular highway instantiation with its collection of participants will result in contentions with a non-ideal outcome where a lower priority contender can preempt a higher priority contender because of asymmetry. This is entirely a function of the technologies that use a spectrum highway and thus not a function of the highway design but of deciding which users can operate on the same highway.

B.9.10 Contention versus Device Asymmetry

This metric is usually a measure of the operation of a highway and requires an understanding of the devices that operate on the highway. The parameter UseAsymmetry in Table D-15 restricts the device range to 0.9 of the signal range. Assuming all devices follow this guidance, the design precludes any degradation because of this type of asymmetry.

B.9.11 Multifunction Count

The multifunction count captures the different variety of uses of the highway resources. Each lane is a category of use and then each duration and persistence of the duration identifies the variety of use. Tables B-16 through Table B-19 list the different functions for each lane of the highways whose priority phase sets are described in Section B.5 and tallies the count. The per slot design of Table B-7 has a multifunction count of 8. The per slot design of Table B-8 has a multifunction count of 11. The consolidated signaling design of Table B-9 has a multifunction count of 32. The consolidated signaling design of Table B-10 has a multifunction count of 20.

Table B-16. Functions and Their Count for the per-Slot Design of Table B-7

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation								
Frame Access								
Frame Reservation								
Epoch Access								
Total	1	1	1	1	1	1	1	1

Table B-17. Functions and Their Count for the per Slot Design of Table B-8

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation								
Frame Access			✓	✓		✓		
Frame Reservation								
Epoch Access								
Total	1	1	2	2	1	2	1	1

Table B-18. Functions and Their Count for the Consolidated Signaling Design of Table B-9

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation	✓	✓	✓	✓	✓	✓	✓	✓
Frame Access	✓	✓	✓	✓	✓	✓	✓	✓
Frame Reservation	✓	✓	✓	✓	✓	✓	✓	✓
Epoch Access								
Total	4	4	4	4	4	4	4	4

Table B-19. Functions and Their Count for the Consolidated Signaling Design of Table B-10

Function	Value			
	1	2	3	4
Timeslot	✓	✓	✓	✓
Timeslot Reservation	✓	✓	✓	✓
Frame Access	✓	✓	✓	✓
Frame Reservation	✓	✓	✓	✓
Epoch Access	✓	✓	✓	✓
Total	5	5	5	5

B.9.12 Collision Resolution

When using CRS, the collision resolution metric follows from the contention design chosen. The contention design for all lanes use the same number of phases, namely seven, but the lanes with broader bandwidth assume a smaller population of contenders and so use designs driven by smaller design densities. They have greater success at contender densities less than and or equal to those design densities. The comparison of the collision resolution performance of these design can be found in Table 3-2, Figure 3-12, and Figure 3-13.

B.9.13 Precedence Count

The precedence count captures the quantity of differentiation levels. This is driven by the number of phases allocated to the priority phase set, their distribution to user/use priority or QoS, and then whether long-duration accesses are permitted and whether there are timeslot reservations. Table B-20, Table B-21, and Table B-22 list the count of precedence levels for each lane for the priority phase set designs of Table B-7, Table B-8, and Table B-9 respectively. They illustrate that if frame or epoch access is permitted two of the user/use levels are borrowed to make that distinction. These levels are used only if both frame and epoch access are

permitted. In these cases only frame access is used and so only one of the two precedence levels is counted. Second, if reservations are enabled then a phase must be borrowed from either those used for user/use differentiation or QoS differentiation. In the case of the design of Table B-9 the slot was borrowed from the QoS differentiation, resulting in half the QoS levels compared to when reservations are not enabled.

Table B-20. Precedence Count for the per-Slot Signaling Design of Table B-7

Function	Value							
	1	2	3	4	5	6	7	8
Frame Access								
Epoch Access								
User/Use Levels	8	8	8	8	8	8	8	8
QoS Levels	8	8	8	8	8	8	8	8
Timeslot Reservation								
Frame Reservation								
Total	16	16	16	16	16	16	16	16

Table B-21. Precedence Count for the per-Slot Signaling Design of Table B-8

Function	Value							
	1	2	3	4	5	6	7	8
Frame Access			1	1		1		
Epoch Access								
User/Use Levels	8	8	6	6	8	6	8	8
QoS Levels	8	8	8	8	8	8	8	8
Timeslot Reservation								
Frame Reservation								
Total	16	16	15	15	16	15	16	16

Table B-22. Precedence Count for the Consolidated Signaling Design of Table B-9

Function	Value							
	1	2	3	4	5	6	7	8
Frame Access	1	1	1	1	1	1	1	1
Epoch Access								
User/Use Levels	6	6	6	6	6	6	6	6
QoS Levels	4	4	4	4	4	4	4	4
Timeslot Reservation	1	1	1	1	1	1	1	1
Frame Reservation	1	1	1	1	1	1	1	1
Total	13	13	13	13	13	13	13	13

B.9.14 Reservation Count

The priority phase sets of Table B-7 and Table B-8 do not support reservations and so their reservation count is 0. The priority phase set design of Table B-9 allows reservations on all lanes for both timeslots and frames. Since there are ten timeslots per frame and four frames per epoch, each lane has a total reservation count of 14. This is the same count as that for the lanes of the priority phase set of Table B-10

B.9.15 Resilience

This measure requires knowledge of a receiver's ability to detect a signal in the presence of jamming. This typically requires some level of experimentation to discern how well a receiver can detect a signal and thus is a function of both the signaling design and the receiver used in the execution of the contention. This information is not available to provide a metric.

B.9.16 Mobility

The measure used for mobility is the duration of time from the end of contention to the end of use. Shorter times are more tolerant to the mobility of the devices. In the case of per-slot signaling where the transmission starts immediately after the contention, the mobility metric is the same as the duration of the transmission opportunity, which is 39,630,000 nanoseconds for the designs of Table B-2. When this design allows the frames access, the mobility metric for a frame access is 489,645,000 nanoseconds.

In the case of consolidated signaling, Table D-3, the mobility metric will be different for each timeslot based on its ordinal position in the frame. Equation (B-12) provides the metric as a function of the ordinal position of the timeslot denoted by the parameter n . Earlier timeslots have better mobility metrics.

$$\text{mobility} = (1031100000 + n \cdot 39689000) - 50000 - (n) \cdot 10305000 - (n-1) \cdot 5000 \text{ nanoseconds.} \quad (\text{B-12})$$

The mobility metric for frame contentions for this consolidated design, if used, would also be 489,645,000.

B.10 Conclusion

This appendix has provided examples of the timing and signaling design of spectrum highways that use in-band signaling. In practice, a signaling definition is provided with each lane definition. However, the design must consider how participants on each lane interact collectively to determine the active lane. In these designs, all devices contend on the same channel to establish the active lanes and then move to separate channels to contend for access to one of the lanes that is active.

The appendix also provided an example of a signal design. This signal design attempts to limit the bandwidth of the signal to just 25 kHz, which is achieved by having a relatively long pulse. Combined with the blanking interval for changing frequency, a signal phase a duration of nearly 572 microseconds. The total duration of a timeslot contention in the highway contention designs is over 10 milliseconds. The signaling duration consumes a large portion of the highway's capacity.

Appendix C Examples of Out-of-Band Signaling Designs

This appendix provides describes a highway access design that uses out-of-band signaling. This access design supports the eight-lane highway illustrated in Figure 4-2 and the contention transmission opportunity designs illustrated in Figure 4-4 and Figure 4-5 for serial out-of-band contention and Figure 4-7 for concurrent out-of-band contention. The design of this exemplar involved a deliberate effort to replicate many of the features of the in-band highway design presented in Appendix B but in a way that highlights the performance trades between in-band and out-of-band signaling. Several sections below are identical to those in Appendix B.

An SCM Authorization Set is used to define a highway. There is a transmitter model for each lane of the highway, so in the eight-lane case used in this example, there would be eight transmitter models that each define the boundaries of one of the lanes. The access design governed by the rules provided in Chapter 5 is found in the SCMPolicyOrProtocol construct of the transmitter model. Each transmitter model uses the same protocol, “SCRHighwayAccess,” but the designs of access for each lane may vary and so are described separately. Though described separately, the access designs must be designed to operate in a complementary way as previously described.

C.1 Access Rule Policy

The SCMPolicyOrProtocol construct of SCMs is used to define the contention design and the rules for access to each lane. It begins with the PorPName which identifies the protocol or policy associated with the rest of the model. In this case the PorPName is “SCRHighwayAccess.” The next sub-element of a the SCMPolicyOrProtocol construct is the structure PorPPParameters. It consists of an unspecified number of parameters that are known and understood by SDSs that are designed to participate on the highway.

Table C-1 lists the two sub-elements of the SCMPolicyOrProtocol construct that are part of the transmitter model of each lane.

Table C-1. Highway Access Policy Description

Primary Data Element	Sub - elements	Value	Notes
SCMPolicyOrProtocol	PorPName	SCRHighwayAccess	This is the name given to the design and complementary rule sets described in this report. An SDS designed to operate on a spectrum highway will understand what to look for in the subsequent elements to know how to behave on the highway.
	PorPPParameters	No value, a series of “Parameter” follow that define the access	The PorPPParameters construct is a structure with two sub-elements. Only the Parameter sub-element is used but it is used many times over to define the contention design.

The PorPPParameters define the specifics of the access design of each lane. In expressing the details, each “Parameter” has a name and optionally a type and a value. They are to be provided in the order used to describe the data elements in Chapter 5: Timing Hierarchy, Contention Method, Lane Use Precedence Phase Set, Lane Selection Phase Set, Priority Phase Set, Contention Phase Set, Contender ID Phase Set, Contention Channels, and Power Symmetry.

C.2 Timing Hierarchy and Contention Method

The start of the rules defines the timing used for the lanes and the contention method. All lanes of a highway will have the same timing hierarchy and the same contention method. The contention method may be executed serially on a single channel or concurrently on parallel channels. Table C-2 provides the data for defining serial contention and Table C-3 provides the data for defining concurrent contention. These lists of parameters would be the same for all lane models of the same highway except for the order parameter of serial contention as shown in Table C-2.

Table C-2. PorPPParameters for the Highway Timing Hierarchy (out-of-band serial contention)

Primary Data Element	Sub-elements	Value	Notes
PorPPParameters			
Parameter 1			
	PPPName	TimeReference	
	PPPType	DATETIME	
	PPPValue	2018,12,01,24,00,0,-05,00	UTC Time for 2400 1 December 2018, EST
Parameter 2			
	PPPName	Frame	
	PPPType	INTEGER	
	PPPValue	12	
Parameter 3			
	PPPName	Timeslot	
Parameter 4			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	40000000	In nanoseconds so 40 milliseconds duration
Parameter 5			
	PPPName	TimeslotRange	
	PPPType	NUMBER	
	PPPValue	15000	15 kilometers is a 50 microseconds duration
Parameter 6			
	PPPName	Epoch	
	PPPType	INTEGER	
	PPPValue	4	
Parameter 7			
	PPPName	ContentionMethod	
	PPPType	STRING	
	PPPValue	SERIAL	
Parameter 8			
	PPPName	OutOfBandOffset	
	PPPType	NUMBER	
	PPPValue	0	

Primary Data Element	Sub-elements	Value	Notes
Parameter 9			
	PPPName	Order	
	PPPType	INTEGER	
		1 2 3 4 5 6 7 8	Each lane would have its own PPPValue
	PPPValue	0 1 2 3 0 2 0 0	

Table C-3. PorPPParameters for the Highway Timing Hierarchy (out-of-band parallel contention)

Primary Data Element	Sub - elements	Value	Notes
PorPPParameters			
Parameter 1			
	PPPName	TimeReference	
	PPPType	DATETIME	
	PPPValue	2018,12,01,24,00,0,-05,00	UTC Time for 2400 1 December 2018, EST
Parameter 2			
	PPPName	Frame	
	PPPType	INTEGER	
	PPPValue	25	
Parameter 3			
	PPPName	Timeslot	
Parameter 4			
	PPPName	TimeslotDuration	
	PPPType	NUMBER	
	PPPValue	20000000	In nanoseconds so 20 milliseconds duration
Parameter 5			
	PPPName	TimeslotRange	
	PPPType	NUMBER	
	PPPValue	15000	15 kilometers is a 50 microseconds duration
Parameter 6			
	PPPName	Epoch	
	PPPType	INTEGER	
	PPPValue	4	
Parameter 7			
	PPPName	ContentionMethod	
	PPPType	STRING	
	PPPValue	CONCURRENT	
Parameter 8			
	PPPName	OutOfBandOffset	
	PPPType	NUMBER	
	PPPValue	0	

The difference between these designs is the value for ContentionMethod, either Serial or Concurrent, and the presence of the order parameter when serial contention is used. The timeslot durations are different to demonstrate one of the trades between serial and concurrent signaling. Concurrent signaling supports timeslots that can be as small as the time it takes to do a single contention. TimeslotDuration for serial contention must accommodate the time required for multiple contentions, which is 32.63 msec for this highway. Lane selection occurs only once in the signaling and then there are four contentions. As with in-band signaling, it is possible to

have designs with timeslots that have different durations. This would require a TimeslotDuration definition for each of the timeslots as demonstrated in Table B-4. The timing of out-of-band contention is always referenced to the start of the timeslot being contended.

C.3 Lane Use Precedence Phase Set

The lane use precedence phase set is used to allow certain devices to have precedence in the ability to establish the active lanes. This design will be the same for all lanes of a highway. Table C-4 illustrates the parameters that define the phase set for each of the lanes.

Table C-4. PorPPParameters Defining the Lane Use Precedence Phase Set

Sub - elements	Value							
	1	2	3	4	5	6	7	8
Parameter LU1								
PPPName	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases	LaneUsePrecedence Phases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	2	2	2	2	2	2	2	2
Parameter LU2								
PPPName	Override	Override	Override	Override	Override	Override	Override	Override
Parameter LU3								
PPPName	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1	LPChannel 1
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

This design establishes two phases for asserting precedence. The highest signaling level (i.e., signal in all phases) is reserved for override, so that devices with urgent use, as established by the rules, may override the routine precedence differentiation. As a result, this design has three routine precedence levels and the override precedence level.

There is an implied timing associated with this definition. When used it starts at the offset relative to the start of the timeslot being contended. It, together with the lane selection phase set are executed once at the start of the contention signaling and in serial signaling only the precedence and contention phases sets are executed serially.

C.4 Lane Selection Phase Set

There is no difference between how lane selection signaling is executed in out-of-band signaling versus in-band-signaling. The same signaling combinations are used to establish lanes. The following design and description is the same as that found in Section B.4.

The lane selection phase set signaling design is the same for all lanes, but the signaling that results in a lane becoming active differs by lane. The parts of the lane selection phase set description that concern the number of phases, the pause that precedes this phase set, and the channel used for selecting a lane are the same for all lanes of the same highway. The signaling that results in a lane becoming active may vary. Table C-5 illustrates that these differences center around what signaling is used to select a lane and the observed signals that indicate a lane

is active and whether the observing device can contend to gain access to that lane. As seen in this table, the signaling and the number of signal combinations that result in a lane being active varies by lane.

Table C-5. PorPPParameters Defining the Lane Precedence Phase Set

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter LP1								
PPPName	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases	LaneSelectionPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter LP2								
PPPName	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause	LaneSelectionPause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter LP3								
PPPName	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence	LaneSignalingSequence
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	000	000	000	000	001	010	100	110
Parameter LP4								
PPPName	LaneOverride	LaneOverride	LaneOverride	LaneOverride	LaneOverride	LaneOverride	ContentionEligibleOutcome	ContentionEligibleOutcome
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	100	100	100	110	100	100	100	110
Parameter LP5								
PPPName	LaneOverride	LaneOverride	LaneOverride	LaneOverride	ContentionEligibleOutcome	ContentionEligibleOutcome	ActiveOnlyOutcome	ActiveOnlyOutcome
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY
PPPValue	001	001	010	010	001	010	101	111
Parameter LP6								
PPPName	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	LSChannel1	LSChannel
PPPType	BINARY	BINARY	BINARY	BINARY	BINARY	BINARY	STRING	STRING
PPPValue	000	000	000	000	011	011	C1	C1
Parameter LP7								
PPPName	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	ContentionEligibleOutcome	LSChannel1	LSChannel1		
PPPType	BINARY	BINARY	BINARY	BINARY	STRING	STRING		
PPPValue	010	010	001	100	C1	C1		
Parameter LP8								

Sub-elements	1	2	3	Value				
PPPName	LSChanne 1	LSChanne 1	LSChanne 1	Contentio nEligible Outcome				
PPPTType	STRING	STRING	STRING	BINARY				
PPPValue	C1	C1	C1	001				
Parameter LP9								
PPPName				Contentio nEligible Outcome				
PPPTType				BINARY				
PPPValue				101				
Parameter LP10								
PPPName				LSChanne 1				
PPPTType				STRING				
PPPValue				C1				

There are several key observations. Each lane may have a different number of parameters listed and different types of parameters. Each lane has a signaling sequence that a device would use to establish the lane as active, called LaneSignalingSequence. Each lane definition may provide signaling that indicates the lane cannot be won so cease contending for it, called LaneOverride. Each lane definition provide one or more signaling outcomes that indicate a lane is active and that the device may contend for access to that lane, called ContentionEligibleOutcome. Some lane definitions may provide a signaling combination that indicates that the lane is active but that the device hearing that sequence may not contend for the lane itself, called ActiveOnlyOutcome. Of note in the table are the number of ContentionEligibleOutcomes for Lane 4. This is because there are many lanes that can be active that do not overlap Lane 4: specifically, Lanes 1, 2, 3, 5, and 7. Lane 7 and Lane 8 have the fewest parameters because there is one and only one signaling outcome that permits contention for these lanes and no signaling that will override their selection. Given this collection of signaling outcomes for all lanes, devices can use their observations to determine all the lanes that are active in their location.

C.5 Priority Phase Set

The design of the priority phase set can also be the same across all types of signaling, however in the case of serial contention there designer has some additional options. In all other signaling approaches the number of priority phases must be identical. In the case of serial signaling, the priority phases may have different numbers of phases lane to lane (except when the same signaling order number is used, for example Lanes 1, 5, 7, and 8, and Lanes 3 and 6 in this design.

Table C-6 provides a simple serial out-of-band signaling design where there are no contentions for frames or epochs, nor are there any reservation on any of the lanes. Much of the description below for out-of-band serial signaling is the same as that used for in-band signaling with two exceptions: the signaling is all done on the same channel, and there are more precedence phases in contention for Lanes 1, 5, 7, and 8. Referring back to Table C-2, the signaling for these lanes all occurs in the same contention period in the series and so contentions for all lanes must have the same number of phases listed in their design. Only one of these types of lanes – 1, 5, 7, or 8 – can exist locally.

Table C-6. PorPPParameters Defining the Priority Phase Set (serial signaling)

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	7	6	6	6	7	6	7	7
Parameter PP2								
PPPName	PriorityPhase	PriorityPhase	PriorityPhase	PriorityPhase	PriorityPhase	PriorityPhase	PriorityPhase	PriorityPhase
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP3								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP4								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	4	3	3	3	4	3	4	4
Parameter PP5								
PPPName	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

Table C-7 provides a simple concurrent out-of-band signaling design. It is identical to the serial in-band signal design described by Table B-7. Four different channels are used in the priority signaling, but in this case the definition of those channels, which would follow in the contention channels, would be on out-of-band channels.

Table C-7. PorPPParameters Defining the Priority Phase Set (concurrent signaling)

Sub - elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	6	6	6	6	6	6
Parameter PP2								
PPPName	PriorityPauses	PriorityPauses	PriorityPauses	PriorityPauses	PriorityPauses	PriorityPauses	PriorityPauses	PriorityPauses
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP3								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP4								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP5								
PPPName	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel	PChannel
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C2	C3	C4	C1	C3	C1	C1

Similarly, Table C-8 provided a design that allowed frame contention on Lanes 3, 4, and 6 when using per-slot contention. In order to use the same number of contention phases as in the other contentions, it was necessary to reduce the number of QoS phases available for these lanes. In serial signaling, as shown in Table E-8, there is no requirement to have the same number of phases for signaling across all contention designs and therefore an additional signaling phase was added to those signaling designs.

Table C-8. PorPPParameters Defining the Priority Phase Set (serial signaling with long-duration contentions)

Sub - elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter PP1								
PPPName	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases	PriorityPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	6	6	7	7	6	7	6	6
Parameter PP1								
PPPName	PriorityPa use	PriorityPa use	PriorityPa use	PriorityPa use	PriorityPa use	PriorityPa use	PriorityPa use	PriorityPa use
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter PP1								
PPPName	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases	QoSPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	3	3	3	3	3	3
Parameter PP1								
PPPName	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases	UserUsePhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	3	3	4	4	3	4	3	3
Parameter PP1								
PPPName	PChannel	PChannel	FrameContention	FrameContention	PChannel	FrameContention	PChannel	PChannel
PPPType	STRING	STRING			STRING		STRING	STRING
PPPValue	C1	C1			C1		C1	C1
Parameter PP1								
PPPName			PChannel	PChannel		PChannel		
PPPType			STRING	STRING		STRING		
PPPValue			C1	C1		C1		

Recall that a concern with long-duration access with in-band signaling was that all the SDSs that would want to use a lane that overlapped in frequency with one of the lanes with a long reservation, would have to listen to the contentions where these reservation are made and delay access until one of these contentions passed. No in-band signaling was to take place during the access of a frame or epoch. In out-of-band signaling, a device that has this type of long-duration access can continue to signal out-of-band to inform neighbors that the frame or epoch is in use.

Conveniently, participating in out-of-band signaling throughout a long frame or epoch access also allows the creation of signaling that permits higher priority users to preempt the long-duration access. This is true for both serial and concurrent signaling.

The designs that are practical for consolidated in-band signaling, such as that defined in Table B-9, are equally executable with either out-of-band signaling technique. Out-of-band signaling readily supports highway designs with multiple types of long-term access: timeslot reservation, frame access, frame reservation, and epoch access. In execution, however, each of the long-duration accesses precludes the others. A device may not interrupt the long-duration access of another device unless it has preemption authority.

The expectation in this design is that those devices that contend and win access to a whole frame will subsequently contend in the follow-on contentions of the frame using the override lane use precedence to preserve the designated lane as active. They would also use the frame priority to preserve their access to the frame and to cause other devices in their vicinity to defer to another frame or epoch or move to another lane for their timeslot contentions.

This design also supports timeslot reservations and frame reservations. Timeslot reservations and any type of frame contention on the same lane are mutually exclusive. Thus, a device wanting to make a timeslot reservation on a lane must wait for a frame in which no a frame contention occurs. When frame reservations are permitted, then the device must wait an epoch number of consecutive frames with no frame contention before making a timeslot reservation. The use of the ReservationPrecedence parameter and the rule it invokes prevents a frame contention from interrupting a timeslot reservation.

However, a device with authority to preempt a reservation may use all signaling phases of the priority phase set to do so. It will still have to contend in the contention phases because their might be a peer that also has an urgent mission.

Similar to in-band signaling and as defined by Table B-10, a highway design without overlapping lanes using out-of-band signaling is much more supportive of long duration contentions for frames and epochs. As is the case with in-band signaling, using the priority phase set design of Table B-10 would eliminate the need for a lane use precedence phase set and a lane selection phase set. The active lanes do not change. However, when using serial signaling, it is possible for a device to contend and win several adjacent lanes and in so doing create access for a wider band channel. This approach to wider lanes does put a burden on a device to be able to rapidly change its spectrum use based on success in contention. It would not be appropriate for a device wanting to have access to adjacent lanes to let a lane go unused if it is unsuccessful in gaining access to all the adjacent lanes. It should have a way to productively use the lane to which it gets access.

C.6 Contention Phase Set

The contention phase set is one of the simpler ones to define and the only difference between in-band and out-of-band signaling is the channel used for the signaling. These channels are almost always the same as those used for the priority phase set. In this lane design, concurrent out-of-band signaling would still use four channels for contention – the same as the design for in-band signaling – as shown in Table B-11. In the case of serial signaling, all lanes would use the same

channel for their contention. Table C-9 specifies a possible design. Again each of the lanes uses different contention designs, but with the same number of contention phases.

Table C-9. PorPPParameters Defining the Contention Phase Set (with serial out-of-band signaling)

Sub-elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter C1								
PPPName	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases	Contention Phases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	7	7	7	7	7	7	7	7
Parameter C2								
PPPName	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause	Contention Pause
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter C3								
PPPName	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity	DesignDensity
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	200	200	200	1000	200	200	50	50
Parameter C4								
PPPName	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel	CChannel
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

C.7 Contender ID Phase Set

Serial in-band contention signaling supports the use of a phase set to identify contention winners. The contender ID allows individual devices to move across lanes on their own initiative and then indicate their success so that other members of the SDS can move there as well. If a device wins access on a lane it will preempt other devices from contending for subsequent lanes in the series of contentions.

There are three approaches to signaling the SDS ID. The first is through the use of the same type of signals as used in the rest of the contention but in this case the signals convey a binary number. The contention winner uses both the signal and echo, so a single phase provides two bits. The contention winner can only alert the devices that are within one hop of itself of its SDS ID. Table C-10 provides an example of this type of design. Each lane is designed to use the same number of phases. In this design, there are four phases and therefore eight bits and 256 possible identities for SDSs on the highway. The entity that manages the highway will have to give SDSs the ID that devices of the SDS may use on the highway.

Table C-10. PorPPParameters Defining the Contender ID Phase Set (only used with serial out-of-band signaling)

Sub - elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter CID1								
PPPName	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	Binary	Binary	Binary	Binary	Binary	Binary	Binary	Binary
Parameter CID2								
PPPName	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases
PPPType	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	4	4	4	4	4	4	4	4
Parameter CID3								
PPPName	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter CID4								
PPPName	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel
PPPType	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

The second alternative is to use a frequency in signaling to indicate the ID. After a contention, the devices of each SDS participate in the contention phase on the frequency assigned to that SDS. The phase can have the typical signal and echo and so reach two hops. Table C-11 provides an example definition of this type of design. The entity managing the highway would give guidance on how to select the frequency for this signal. It is desirable that signaling approaches provide a means for the signal frequency to change from timeslot to timeslot.

Table C-11. PorPPParameters Defining the Contender ID Phase Set with Frequency Identification (only used with serial out-of-band signaling)

Sub - elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter CID1								
PPPName	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod	Contender IDMethod
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	Frequency	Frequency	Frequency	Frequency	Frequency	Frequency	Frequency	Frequency
Parameter CID2								
PPPName	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases	Contender IDPhases
PPPTYPE	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER	INTEGER
PPPValue	1	1	1	1	1	1	1	1
Parameter CID3								
PPPName	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter CID4								
PPPName	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C1	C1	C1	C1	C1	C1	C1	C1

The third alternative sets aside a window of time for a contending devices to send a modulated signal that contains the identification of the SDS. Table C-12 provides an example of how to specify this approach to convey contender ID. In this design, 750 microseconds are set aside for this transmission. The IDChannel for all types of lanes would be the same. Here, the definition specifies the use of C2. C2 would be a different class of signal as compared to C1. In all other cases just pulses of energy are sent. Here, the transmission contains more than one bit of information, a modulated version of the ID number and so the modulation and not a signal is specified by the channel definition. It is reasonable that with the development of a highway ecosystem there would be multiple signaling channels and modulated data channel definitions that could be accessed by an out-of-band signaling device.

Table C-12. PorPPParameters Defining the Contender ID Phase Set with Modulated Signal Identification
(only used with serial out-of-band signaling)

Sub - elements	1	2	3	4	5	6	7	8
Parameter CID1								
PPPName	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod	ContenderIDMethod
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	Modulated	Modulated	Modulated	Modulated	Modulated	Modulated	Modulated	Modulated
Parameter CID2								
PPPName	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause	IDPause
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5000	5000	5000	5000	5000	5000	5000	5000
Parameter CID3								
PPPName	IDWindow	IDWindow	IDWindow	IDWindow	IDWindow	IDWindow	IDWindow	IDWindow
PPPTYPE	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	750000	750000	750000	750000	750000	750000	750000	750000
Parameter CID4								
PPPName	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel	IDChannel
PPPTYPE	STRING	STRING	STRING	STRING	STRING	STRING	STRING	STRING
PPPValue	C2	C2	C2	C2	C2	C2	C2	C2

C.8 Contention Channels

The contention channels definition of Table B-14 could be a definition for concurrent out-of-band signaling as well. When serial out-of-band signaling is used, the channel definition would be the exact same for all lanes. If contender ID is used and the method of conveying that ID is a modulated signal then the contention channel definition would have to define that modulated channel.

C.9 Range Symmetry

The range symmetry provides important guidance when out-of-band signaling is used. It is possible that the signaling is done on a completely different band where propagation characteristics are very different. For example, signaling may use VHF channels while the lanes are defined for UHF or SHF channels. It is less meaningful to specify power symmetry. All signaling devices are likely to use the same power. Directional antenna effects would not exactly match those of the SDS devices. The designers of the highway would address symmetry issues by specifying the power used by signals and by placing restrictions on the range of the SDSs that use the highway. The lanes would also be managed by putting a restriction on the power transmitted on the lanes as defined by the spectrum masks.

Table C-13 provides the definition of the range symmetry. It describes the planning ranges of the signaling based on the signal power that is used and then restricts the range of the SDSs operating on the highway to some fraction of the range of the signals. In this case, that fraction, UseAsymmetry is 1.1 of the signal range. This design anticipates that users will operate on a frequency that will have a greater rate of propagation attenuation and that some SDSs can take advantage of processing gain to mitigate interference and so operating at these greater ranges will not have a significant effect on peers.

Table C-13. PorPPParameters Defining the Range Symmetry

Sub - elements	Value							
	1	2	3	4	5	6	7	8
PorPPParameters								
Parameter RS1								
PPPName	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower	NominalPower
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	5	5	5	5	5	5	5	5
Parameter RS1								
PPPName	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance	MaxPowerVariance
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	0.5	2	2	2	2	2	2	2
Parameter RS1								
PPPName	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain	MaxAntennaGain
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	3	3	3	3	3	3	3	3
Parameter RS1								
PPPName	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange	TimingRange
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	15000	15000	15000	15000	15000	15000	15000	15000
Parameter RS1								
PPPName	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange	SquareLawRange
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	40000	40000	40000	40000	40000	40000	40000	40000

Sub - elements	Value							
	1	2	3	4	5	6	7	8
Parameter RS1								
PPPName	Terrestrial Range	Terrestrial Range	Terrestrial Range	Terrestrial Range	Terrestrial Range	Terrestrial Range	Terrestrial Range	Terrestrial Range
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	9000	9000	9000	9000	9000	9000	9000	9000
Parameter RS1								
PPPName	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry	UseAsymmetry
PPPType	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER	NUMBER
PPPValue	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1

C.10 Design Metrics

The design metrics for out-of-band signaling highways will be different from those used for in-band signaling highways. The issue when using out-of-band signaling when compared to in-band signaling is the additional spectrum consumed by the signaling rather than the time consumed by the signaling. Since signaling can be designed with rather narrowband channels, this may not be a significant cost. If frequency hopping is desirable to harden the signaling against adversaries wanting to deny service, then the cost may be higher, depending on how that signaling precludes other uses of the same spectrum. The signaling of multiple highways could all occur in the same spectrum, in which case larger bands for frequency hopping would be less of an issue.

A benefit of out-of-band signaling is that the duration of transmission opportunities can be shorter with out-of-band signaling without sacrificing efficiency. The duration of transmission opportunities can be no shorter than the time it takes to do the signaling so concurrent signaling designs support shorter transmission opportunities than serial designs.

As before, some metrics require estimates of the characteristics of the SDSs that will operate on the highway, e.g. use efficiency. The subsections below provide metrics that are applicable to the out-of-band signaling designs defined above.

C.10.1 Temporal Overhead

In out-of-band signaling designs, the only time that must be set aside is to end the transmission 50,000 nanoseconds before the next timeslot. Thus in the designs with 40,000,000-nanosecond timeslots the temporal overhead is

$$\frac{50,000}{40,000,000} = 0.0013. \quad (C-1)$$

In the designs with 20,000,000-nanosecond timeslots the temporal overhead is

$$\frac{50,000}{20,000,000} = 0.0025. \quad (C-2)$$

Highways that allow frame and epoch contentions will have even better temporal overhead measures.

C.10.2 Spectral Overhead

The spectral overhead is the ratio of the spectrum used for access to the bandwidth of the lane. The combined bandwidth of the base lanes, Lanes 1, 2, 3, and 4, is 4 MHz and the bandwidth of a signal, as defined in Section B.7, is 25 kHz. The spectral overhead of serial signaling where only one channel is required is

$$\frac{25,000}{4,000,000} = 0.0063 . \quad (C-3)$$

The spectral overhead of concurrent signaling where four channels are required is

$$\frac{100,000}{4,000,000} = 0.0250 . \quad (C-4)$$

The signal design assumed frequency hopping. In this case the Spectral Overhead may be larger. The method to compute this measure would depend on the number of highways and lanes served by the band of spectrum used for signaling. The overhead becomes the bandwidth of the signal itself and a proportional part of the hopping band that is not supporting any signal. At best the measure is the efficiency computed above.

Additional spectral loss may be computed as the proportion of a lane that is lost to guard bands between lanes. The bandwidth of the guard band would be measured from some point where the spectrum mask of the lane is some power level below the maximum power level to the point that divides adjacent lanes. Using the example of Section B.9.2 and the spectral overhead computed in Equation (B-4), the total temporal overhead is the sum of the overhead associated with the spectral occupancy of the lane and the use of spectrum by contention resolution signaling, so 0.1913 for serial signaling and 0.2100 for concurrent signaling.

C.10.3 Use Efficiency

As described in Section B.9.3 this measure is best made when a spectrum highway is in operation. Use efficiency is achieved by matching the highway designs to the ways the SDSs that use the highway use spectrum.

C.10.4 Access Delay

The access delay for the serial contention is the same for all lanes when the individual lane contentions are preceded by the lane selection phases set, since there is only one. It is the start offset plus the duration of the timeslot. In this case, with a 0 duration OutOfBandOffset the access delay is 40,000,000 nanoseconds. In the case of the concurrent signaling design it is 20,000,000 nanoseconds.

C.10.5 Slot Delay

Since in both out-of-band signaling designs the contention for the next timeslot starts concurrently with the current timeslot, the soonest a contention can change the use of spectrum is during the contention for the subsequent slot and so the slot delay in both cases is the duration of a timeslot. It is therefore 40,000,000 nanoseconds in the serial signaling design and 20,000,000 nanoseconds in the concurrent signaling design.

C.10.6 Responsiveness

The responsiveness measure is the sum of the access delay and the slot delay. Therefore it is 80,000,000 nanoseconds for the serial signaling design and 40,000,000 nanoseconds for the concurrent design.

C.10.7 Spatial Capacity

The signal design used for these out-of-band contention designs is the same as that used for the in-band contentions and so the spatial capacity will be the same, one lane user per 509 square kilometers. See Section B.9.7 for an explanation of spatial capacity.

C.10.8 Spectral Agility

A unique benefit of serial out-of-band signaling is that it supports spectrum agility. An individual device can choose to change lanes and can inform peers of the same SDS of the change using the SDS ID portion of the contention. SDSs can move their use to different lanes from timeslot to timeslot. The time to change lanes is the same as the access delay, 40 milliseconds.

Concurrent out-of-band signaling does not have this advantage and spectral agility is not a function of either the highway design or its operation. It is a function of protocols internal to the SDS.

C.10.9 Asymmetry Risk

The asymmetry risk is a measure of the extent to which a particular highway instantiation with its collection of participants will result in contentions with a non-ideal outcome where a lower priority contender can preempt a higher priority contender because of asymmetry. This is entirely a function of the technologies that use a spectrum highway and thus is not a function of the highway design, but of deciding which users can operate on the same highway.

C.10.10 Contention versus Device Asymmetry

This metric is usually a measure of the operation of a highway an requires and understanding of the devices that operate on the highway. The parameter UseAsymmetry in Table C-13 restricts the device range to 1.1 of the signal range. Assuming some devices operate beyond the signal range there will be some risk of degradation. In this case, that degradation has to be measured by observing the highway in operation. Simulation may be used to estimate this measure.

C.10.11 Multifunction Count

The multifunction count captures the different varieties of uses of the highway resources. Each lane is a category of use and then each duration and persistence of the duration identifies the variety of use. Tables C-14 through Table C-16 list the different functions for each lane of the highways whose priority phase sets are described in Section C.5 and tallies the count. The serial signaling design of Table C-6 has a multifunction count of 8. The concurrent signaling design of Table C-7 has a multifunction count of 8. The serial signaling design of Table C-8 has a multifunction count of 11. These counts are relatively low and only the design of Table C-8 supports long duration contentions. Ironically, because signaling is done on a separate out-of-band channel, these signaling approaches are best equipped to support long duration contentions.

Table C-14. Functions and Their Count for the Serial Signaling Design of Table C-6

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation								
Frame Access								
Frame Reservation								
Epoch Access								
Total	1	1	1	1	1	1	1	1

Table E-15. Functions and Their Count for the Concurrent Signaling Design of Table E-7

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation								
Frame Access								
Frame Reservation								
Epoch Access								
Total	1	1	1	1	1	1	1	1

Table C-16. Functions and Their Count for the Serial Signaling Design of Table C-8

Function	Value							
	1	2	3	4	5	6	7	8
Timeslot	✓	✓	✓	✓	✓	✓	✓	✓
Timeslot Reservation								
Frame Access			✓	✓		✓		
Frame Reservation								
Epoch Access								
Total	1	1	2	2	1	2	1	1

C.10.12 Collision Resolution

When using CRS, the collision resolution metric follows from the contention design chosen. The contention design for all lanes use the same number of phases sets, seven, but they assume different design densities. Only the Lane 4 design uses a design density of 1000. The rest of the

lanes use a design density of 200, with the exception of the very broadband Lanes 7 and 8 which use a design density of 50. The designs based on smaller design densities will have greater success at contender densities less than or equal to those design densities. The comparison of the collision resolution performance of these design can be found in Table 3-2, Figure 3-12, and Figure 3-13.

C.10.13 Precedence Count

The precedence count captures the quantity of differentiation levels. This is driven by the number of phases allocated to the priority phase set, their distribution to user/use priority or QoS, and then whether long duration accesses are permitted and whether there are timeslot reservations. Table C-17 and Table C-18 lists the count of precedence levels for each lane for the priority phase set designs of Table C-6 and Table C-8. The priority phase set design of Table C-7 is the same as that of Table B-7, and Table B-20 provides that count. These illustrate that with out-of-band signaling it is possible to have signaling designs that vary from lane to lane. Additional phases are added to the priority phase sets that support access to long duration lanes in the case of the serial signaling design of Table C-6. Additional phases are added to the lanes that support frame access in the case of the serial signaling design of Table C-8.

Table C-17. Precedence Count for the Serial Signaling Design of Table C-6

Function	Value							
	1	2	3	4	5	6	7	8
Frame Access								
Epoch Access								
User/Use Levels	16	8	8	8	16	8	16	16
QoS Levels	8	8	8	8	8	8	8	8
Timeslot Reservation								
Frame Reservation								
Total	24	16	16	16	24	16	24	24

Table C-18. Precedence Count for the Serial Signaling Design of Table C-8

Function	Value							
	1	2	3	4	5	6	7	8
Frame Access			1	1		1		
Epoch Access								
User/Use Levels	8	8	14	14	8	14	8	8
QoS Levels	8	8	8	8	8	8	8	8
Timeslot Reservation								
Frame Reservation								
Total	16	16	23	23	16	23	16	16

C.10.14 Reservation Count

None of the priority phase sets support reservations so their reservation counts are all 0.

C.10.15 Resilience

This measure requires knowledge of the ability of a receiver to detect a signal in the presence of jamming. This typically requires some level of experimentation to discern how well a receiver can detect a signal and so is a function of both the signaling design and the receiver used in the execution of the contention. This information is not available to provide a metric.

C.10.16 Mobility

The measure used for mobility is the duration of time from the end of contention to the end of use. Shorter times are more tolerant to the mobility of the devices. These durations are comparatively small for out-of-band signaling. In the case of the concurrent out of-band signaling specified by Table C-7, the mobility metric is 29,695,000 nanoseconds. The mobility metric in the case of serial signaling design will vary by lane since lane contentions end at different times. In all cases the mobility metric is more than 40,000,000 nanoseconds. In the case of the design of Table C-6 where the first contention has an additional phase the mobility metric of Lanes 1 to 4 are

$$mobility = 80000000 - 50000 - 10876667 - (n-1) \cdot 7441667 \text{ nanoseconds.} \quad (C-5)$$

The mobility metric of Lanes 4, 7, and 8 are the same as Lane 1 and the mobility metric for Lane 6 is the same as that of Lane 3.

C.11 Conclusion

This appendix has provided examples of the timing and signaling design of spectrum highways that use out-of-band signaling. As with in-band signaling, a separate signaling design is provided for each lane. There are marked differences between the definitions of serial and concurrent signaling. Concurrent signaling designs are very similar to those of per timeslot in-band signaling designs. All signaling contentions use the same number of phases, there is a requirement to operate on multiple signaling channels for some lanes, and devices focus on the contention for just one lane after the active lanes are established. The only difference is that the contentions are executed on different channels and are simultaneous with lane use. In serial signaling, all contentions use the same channel and all devices, regardless of what lane they are contending to use, can see the results of all contentions on the highway. This means there is a common design for the phase set used in lane selection, but the priority and contention phase sets can be unique for each lane. Additionally, SDSs on spectrum highways using serial signaling can be much more agile in their movement across the lanes of the highway thanks to the combination of their ability to observe all contentions with the use of the contender ID phase set. At the end of the serial signaling, all participating devices know not only which lanes are active but also which SDSs are operating on each lane.

An advantage of out-of-band signaling is that none of the capacity of the lanes is lost to the execution of signaling. This avoids the spectral waste that occurs in highways using in-band signaling when the lanes of the highway are relatively broadband and signals are narrowband. Further, since signaling is protected through the use of frequency hopping, it is better and more efficient if the signaling of multiple lanes and even multiple highways were executed in the same band. Further, the separation of the signaling in frequency can also protect the operation of the

highway. It would be harder for adversaries that can observe the signaling to know where in spectrum that signaling indicates activity that should be monitored or attacked.

Another advantage of out-of-band signaling is the ability to create shorter duration transmission opportunities. The duration is limited to the time it takes to execute the signaling for each timeslot. Thus the duration of timeslots depends on the duration of a signaling phase and the number of signaling phases used to resolve the active lanes and the lane users. The shortest transmission opportunities can be achieved with concurrent signaling, since all lane contentions are concurrent as opposed to serially executed as done in serial signaling.

Appendix D Glossary

Term	Definition
Goodput	The application-level throughput (i.e. the number of useful information bits delivered by the network to a certain destination per unit of time). The amount of data considered excludes protocol overhead bits as well as retransmitted data packets.
Lane Selection Precedence	A precedence level that distinguishes which SDS devices can chose the active lanes of a highway.
Phase Set	A set of signaling phases used to achieve a specified highway access objective to include establishing the active highway lanes, differentiating priority of access based either on the precedence of users and uses or the required quality of service, and resolving contention.
Quality of Service Precedence	A precedence level used to distinguish access precedence based on the internal SDS activity (e.g. distinguishing access of network devices based on message priority).
Signaling Phase	A period where one or two signals are conditionally used by contending devices to indicate a choice, precedence, or random outcome to peer neighboring signalers. In two-signal phases, contending devices transmit on the first signal and immediate neighbors repeat the signal in the second signal of the phase.
Spectrum Highway	A dynamic spectrum access approach that follows the highway metaphor in its design. It consists of bands of spectrum assigned to a geographical area, which are further subdivided into smaller bands called lanes and time periods called transmission opportunities. Access to lanes and transmission opportunities are arbitrated by signaling. There are defined rules for operating on the highways. Each of a plurality of SDSs and their devices that operate on the highway are configured with their authorized precedence for their operation and then collaborate in sharing the spectrum of the highway by following the highway rules.
Spectrum Highway Access Device (SHAD)	A device designed to execute the signaling required for access to a spectrum highway that serves as an SDS device's proxy in negotiating access. SDS devices provide the SHAD with information on their access needs and the SHAD performs the signaling necessary to gain

Term**Definition**

	access. SHADs inform SDSs when they may use or should listen to a transmission opportunity.
Spectrum Lane	A sub-band of a spectrum highway defined by a transmitter spectrum consumption model that indicates the allowed spectral flux density of transmissions on that lane. Authorized SDSs may access the lane for their use by using signaling for that lane and by following the rules of the highway.
Transmission Opportunity	A quantum of spectrum that is accessed on a spectrum highway defined as a bandwidth for a duration of time. A transmission opportunity may be used for transmitting or receiving.
User/Use Precedence	A precedence level used to distinguish access precedence based either on the user of the SDS or the particular use of the SDS (e.g. function in a multifunction system).

Appendix E Abbreviations and Acronyms

ACK	Acknowledgement
AGL	Above Ground Level
BCT	Brigade Combat Team
BN	Battalion
CA	Collision Avoidance
CDR	Commander
CR	Collision Resolution
CRS	Collision Resolution Signaling
CSMA	Carrier Sense Multiple Access
CTS	Clear to Send
DARPA	Defense Advanced Research Projects Agency
DSA	Dynamic Spectrum Access
EA	Electronic Attack
EIRP	Effective Isotropic Radiated Power
EME	Electromagnetic Environment
EMS	Electromagnetic Spectrum
ES	Electronic Support
HAAT	Height Above Average Terrain
ISM	Industrial, Scientific, and Medical
LBT	Listen Before Talk
LPD	Low Probability of Detection
LPI	Low Probability of Interception
MAC	Media Access Control
NAV	Network Allocation Vector
PDU	Protocol Data Unit
PSMA	Packet Sense Multiple Access
QoS	Quality of Service
RBW	Resolution Bandwidth
RTS	Request to Send

SAS	Spectrum Access System
SC2	Spectrum Collaboration Challenge
SCR	Synchronous Collision Resolution
SDS	Spectrum Dependent System
SHAD	Spectrum Highway Access Device
WISP	Wireless Internet Service Provider
WNW	Wideband Networking Waveform

This page intentionally left blank.

