

Broadband Wireless for Rural Areas --

WiFiRe: Medium Access Control (MAC) and Physical Layer (PHY) Specifications

Release 2006 – Ver. 1.1

(This document is – Revised version of the Aug 2006 draft)

Center of Excellence in Wireless Technology (CEWiT)



About CEWiT

The Centre of Excellence in Wireless Technology (CEWiT), India, has been set up under a public-private initiative with the mission of making India a leader in the research, development and deployment of wireless technology. It is an autonomous institution temporarily headquartered at IIT Madras.

Broadband wireless technology has great potential in the coming years. Emerging standards can be leveraged to build a system that specifically meets India's broadband access needs. CEWiT will play a pro-active role in engaging with academic and industry research groups in India to focus research on areas with strong potential. CEWiT will also foster collaboration with similar efforts worldwide. CEWiT seeks to actively participate in International standards bodies, and to assist government and public institutions in policy-making, spectrum management and regulation.

CEWiT Std, WiFiRe, 2006 Edition

CEWiT standards are developed within the Technical Committees of CEWiT. Members of the committees serve voluntarily and without any compensation. The standards developed within CEWiT represent a consensus of the broad expertise of the subject. The existence of a CEWiT standard does not imply that there are no other ways to provide services related to the scope of the standard. Furthermore, a standard is subject to change brought about through developments in the state of the art and comments received from the users of the standard. Users should check that they have the latest edition of any CEWiT standard. These may be obtained from <http://www.cewit.org.in/>

Comments on standards and requests for interpretations relating to specific applications should be addressed to:

Secretary, Center of Excellence in Wireless Technology
CSD152 (ESB)
Indian Institute of Technology Madras
Chennai – 600 036, INDIA
email: feedback@cewit.org.in

The distribution and usage of this Standard are as per the Creative Commons license – Attribution-Share Alike. See <http://creativecommons.org/licenses/by-sa/2.5/> for details. A brief excerpt from the license is given below.

You are free:

- to copy, distribute, display, and perform the work
- to make derivative works
- to make commercial use of the work

Under the following conditions:



Attribution. You must attribute the work in the manner specified by the author or licensor.



Share Alike. If you alter, transform, or build upon this work, you may distribute the resulting work only under a license identical to this one.

- Any of these conditions can be waived if you get permission from the copyright holder.

CEWiT Std, WiFiRe, 2006 Edition

Abstract:

WiFiRe stands for **WiFi – Rural extension**. It seeks to leverage the license free nature of the WiFi spectrum (IEEE 802.11b, 2.4 GHz Band) and the easy availability of WiFi RF chipsets, in order to provide long-range communications (15-20 Kms) for rural areas. The key idea in WiFiRe is to replace the 802.11b MAC mechanisms (DCF/PCF), with something more suitable for long-range communication, while continuing to use the 802.11b PHY support. WiFiRe is meant for a star topology - a Base Station (BS) at the fiber Point of Presence (PoP) and Subscriber Terminals (ST) in the surrounding villages – with sectorized antennas at the BS and a directional antenna at each ST. The WiFiRe MAC is time-division duplex (TDD) over a single 802.11b channel along with a multi-sector TDM mechanism.

This document specifies the details of WiFiRe, including services provided to the higher layers, the message formats and sequences, the protocol description and various timings involved. WiFiRe capacity analysis, scheduler design and simulation analysis are also provided as annexure.

Authors:

Sridhar Iyer (IIT Bombay), **Krishna Paul** (IIT Bombay)¹, **Anurag Kumar** (IISc Bangalore) and **Bhaskar Ramamurthi** (IIT Madras).

Contributors:

Person	Contribution
Ashok Jhunjhunwala, IIT Madras	Conceptualization
Bhaskaran Raman, IIT Kanpur	Management sub-procedures
Anirudha Sahoo, IIT Bombay	Data transport sub-procedures
Om Damani, IIT Bombay	Security sub-procedures
Anitha Varghese, IISc Bangalore	Capacity analysis and scheduler design
Anirudha Bodhankar, IIT Bombay	Simulation model and analysis
Alok Madhukar, IIT Bombay	Data flow and state transition diagrams
Anand Kannan, CEWiT ²	Initial concept document

¹ Krishna Paul was with IIT Bombay when this work was initiated. She joined Intel, Bangalore, towards the end of this work.

² Anand Kannan is now with Valued Epistemics (Pvt) Ltd., Chennai

Reviewers:

The following is a list of reviewers who reviewed the draft version of August 2006:

(The final list of reviewers will include all those who review Ver. 1.1 and turn detailed comments)

Varsha Apte, IIT Bombay
Srinath Perur, IIT Bombay
Raghuraman Rangarajan, IIT Bombay
Vijay Raisinghani, TCS
Kameshwari Chebrolu, IIT Kanpur

Acknowledgements:

The following persons contributed to the discussions and/or other supporting activities:

Pavan Kumar, IIT Kanpur	Pratik Sinha, Zazu Networks
Narasimha Puli Reddy, IIT Kanpur	
Klutto Milleth, CEWiT	K. Giridhar, IIT Madras

Contents

1	OVERVIEW.....	7
1.1	BACKGROUND.....	7
1.2	DEPLOYMENT SCENARIO.....	8
1.3	TECHNOLOGY ALTERNATIVES.....	8
1.4	WiFiRE APPROACH.....	10
1.5	SCOPE.....	13
2	GENERAL DESCRIPTION.....	15
2.1	DEFINITIONS AND ABBREVIATIONS.....	15
2.2	DESIGN DRIVERS AND ASSUMPTIONS.....	15
2.3	WiFiRE SYSTEM ARCHITECTURE.....	16
2.4	NETWORK INITIALIZATION.....	18
2.5	IMPACT OF SECTORIZATION.....	19
2.6	MAC PROTOCOL OVERVIEW.....	21
2.7	MAC SERVICES.....	23
2.8	MAC SERVICE INTERFACES.....	24
2.9	TYPICAL FRAME AND SLOT TIMINGS.....	24
2.10	RANGING AND POWER CONTROL.....	28
2.11	PDU FORMATS.....	29
2.12	BS SCHEDULER FUNCTIONS.....	30
2.13	SUPPORT FOR MULTIPLE OPERATORS.....	30
2.14	SUMMARY OF PROTOCOL STEPS.....	31
2.15	DIFFERENCE BETWEEN WiMAX AND WiFiRE.....	32
3	MAC SERVICE DEFINITION.....	33
3.1	SERVICE SPECIFIC SUB-LAYER (SSS).....	33
3.2	LINK SPECIFIC SUB-LAYER (LCS).....	37
3.3	DETAILED DESCRIPTION OF SERVICE PRIMITIVES.....	41
4	MAC DETAILED DESCRIPTION.....	49
4.1	ADDRESSING AND CONNECTION IDENTIFICATION.....	49
4.2	BANDWIDTH REQUEST GRANT SERVICE.....	50
4.3	MAC PDU FORMAT.....	52
4.4	MAC HEADER FORMAT.....	52

4.5	MAC MANAGEMENT PDU(S)	54
4.6	MAC DATA PDU(S).....	55
4.7	NETWORK INITIALIZATION SUB-PROCEDURES	56
4.8	CONNECTION MANAGEMENT SUB-PROCEDURES	63
4.9	DATA TRANSPORT SUB-PROCEDURES	66
4.10	PROTOCOL SUMMARY: STATE-TRANSITION DIAGRAMS	66
5	AUTHENTICATION AND PRIVACY	72
6	MAC MANAGEMENT	72
7	PHY SERVICE SPECIFICATION AND MANAGEMENT.....	72
8	GLOSSARY OF TERMS	73
8.1	ABBREVIATIONS AND ACRONYMS	73
8.2	DEFINITIONS	75
A.	ANNEX A: DESIGN DRIVERS	75
B.	ANNEX B: CAPACITY ANALYSIS	82
C.	ANNEX C: SCHEDULER DESIGN	91
D.	ANNEX D: SIMULATION ANALYSIS	100
E.	BIBLIOGRAPHY	110

WiFiRe: Medium Access Control (MAC) and Physical Layer (PHY) Specifications

1 OVERVIEW

1.1 Background

About 70% of India's population, or 750 million, live in its 600,000 villages, and around 85% of these villages are in the plains. The average village has 250-300 households, and occupies an area of 5 sq. km. Most of this is farmland, and typically the houses are in one or two clusters. Villages are thus spaced 3-4 km apart, and spread out in all directions from the market centers. The market centers are typically spaced 30-40 km apart. Each such center serves around 250-300 villages, in a radius of about 20 km [1], as shown in Figure 1.




Figure 1: Background

The telecommunication backbone network, passing through all these centers, is new and of high quality optical fiber. The base stations of the mobile (cellular) operations are also networked using optical fiber.

However, the solid telecom backbone ends abruptly at the towns and larger villages. Beyond that, cellular coverage extends mobile telephone connectivity only up to a radius of 5 km, and then telecommunications services peter out. Fixed wireless telephones have been provided in tens of thousands of villages, but the telecommunications challenge in rural India remains the “last ten miles”. This is particularly true if the scope includes broadband Internet access.

The Telecom Regulatory Authority of India has defined broadband services as those provided with a minimum data rate of 256 kbps [2]. Assuming a single kiosk (end-point) in each village, generating sustained 256 kbps flows, 300 kiosks will generate traffic of the order of 75 Mbps. This is a non-trivial amount of traffic to be carried over the air, per base station, even with a spectrum allocation of 20 MHz.

1.2 Deployment Scenario

Given the need to cover a radius of 15-20 km from the fiber point-of-presence (PoP), a broadband wireless system will require a system gain of at least 150 dB. The *system gain* is a measure of the link budget available for overcoming propagation and penetration (through foliage and buildings) losses while still guaranteeing system performance. This may be achieved using Base Station towers of 40 m height, at the PoP, and a roof-top antenna of 10 m height at each Subscriber end (kiosk), with line-of-sight deployment. A subscriber kiosk may also be installed in a vehicle, which may be stationed at different villages over a period of time.

A more detailed discussion on the background and deployment considerations is given in Annex A.

1.3 Technology Alternatives

Technical reviews of current wireless broadband technologies and their evaluations are given in [1,3]. A summary is as follows:

- Present day mobile cellular technologies (such as GSM [4], GPRS [5], CDMA [6]) may meet the cost targets but are unlikely to be able to provide broadband services as defined above.
- Proprietary broadband technologies (such as iBurst [7], Flash-OFDM [8]) meet many of the performance requirements, but typically have low volumes and high costs. The indigenously-developed Broadband corDECT technology [9] of Midas Communication Technologies, Chennai is a fixed-access wireless broadband system that meets the performance and cost requirements.
- WiMAX-d (IEEE 802.16d) [10], is a new standards-based technology for fixed wireless-access that meets the performance requirements, but not the cost targets. Beginning 2007, 16.d products are likely to be deployed in urban areas; initial per user cost projection and spectral efficiency suggest

suggests that it will take years to bring down the cost to a level that is acceptable by rural income group. (See Annex A, section 5 for a detailed discussion). WiMAX-e (IEEE 802.16e) [10] is a mobile evolution of the standard, which may see sufficient and faster (than 16.d) adoption beginning 2008 to generate high volume and drive down the cost. If this happens, it is likely to be a technology that meets both performance and cost requirements. (Also, in urban scenario, WiMAX-e may cannibalize 802.16 d for the lack of mobility factor in 16d). However, it will take some years for the costs of WiMAX based technologies to drop to levels viable for rural deployment.

- WiFi (IEEE 802.11b) [11], is an inexpensive local-area broadband technology. It can provide 256 kbps or more to tens of subscribers simultaneously, but can normally do so only over short distances (less than 50 m indoors). One attraction of WiFi technology is the de-licensing of its spectrum in many countries, including India. Another is the availability of low-cost WiFi chipsets. In rural areas, where the spectrum is hardly used, WiFi is an attractive option, provided its limitations when used over a wide-area are overcome. Various experiments with off-the-shelf equipment have demonstrated the feasibility of using WiFi for long-distance rural point-to-point links [12]. The main issue is that WiFi typically uses a Carrier Sense Multiple Access (CSMA) protocol, which is suited only for a LAN deployment. Further, the Distributed Coordination Function (DCF) mechanism does not provide any delay guarantees, while the Point Coordination Function (PCF) mechanism becomes inefficient with increase in number of stations [13]. When off-the-shelf WiFi equipment is used to set up a wide-area network, medium access (MAC) efficiency becomes very poor, and spectrum cannot be re-used efficiently even in opposite sectors, of a base station. One solution for this problem is to replace the MAC protocol with one more suited to wide-area deployment. This will have to be crafted carefully such that a low-cost WiFi chipset can still be used, while bypassing the in-built WiFi MAC. The alternative MAC can be implemented on a separate general-purpose processor with only a modest increase in cost.

WiFiRe, as defined herewith, is one such alternative MAC designed to leverage the low cost of WiFi technology for providing fixed wireless access. It is a Time Division Duplex (TDD) communication protocol over a single WiFi channel, along with a multi-sector Time Division Multiplex (TDM) mechanism. This is explained in the next section.

There are existing commercial products which support long-distance WiFi links [14,15,16,17]. Some of these products are for point-to-point links and others are for point-to-multipoint links. While the protocol used by such products is proprietary, they are likely to be based on some kind of Time Division Multiple Access (TDMA) mechanism. This is supported by the fact that some of these products allow a network operator to flexibly split the available bandwidth among various clients in a point-to-multipoint setting.

WiFiRe has the following advantages over such products:

- WiFiRe is an *open standard*, whereas the above products involve proprietary protocols which are non-interoperable. The non-interoperability also implies that the *cost* of such products is likely to be higher than standards based products.
- Related to the above, the performance of WiFiRe is more predictable and understandable than that of the proprietary commercial products. This is especially important for large scale deployments.
- All of the commercial products above consider only a single sector operation (single point-to-multipoint link). WiFiRe is designed for higher spectral reuse through multiple carefully planned sectors of operation. Such reuse is estimated to achieve 3-4 times higher throughput performance. With WiFiRe, it is estimated that one can support about 25 Mbps (uplink + downlink) per cell, using a single WiFi carrier at 11 Mbps service. This would be sufficient for about 100 villages in a 15 km radius.

1.4 WiFiRe Approach

WiFiRe stands for **WiFi – Rural extension**. The main design goal of WiFiRe is to enable the development of low-cost hardware and network operations for outdoor communications in a rural scenario. This has two implications: (i) a WiFiRe system avoids frequency licensing costs by operating in the unlicensed 2.4 GHz frequency band, and (ii) WiFiRe uses the WiFi (IEEE 802.11b) physical layer (PHY), due to the low cost and easy availability of WiFi chipsets.

WiFiRe requires a 40 m tower at the base station (BS) near the fiber PoP (point-of-presence) and 10-12 m poles at the subscriber terminals (ST), in order to maintain the desired system gain of about 150 dB. In case of taller obstacles, the height of BS tower and ST tower may require adjustments to maintain the Line of Sight (LOS). The network configuration is a star topology, as shown in Figure 2.

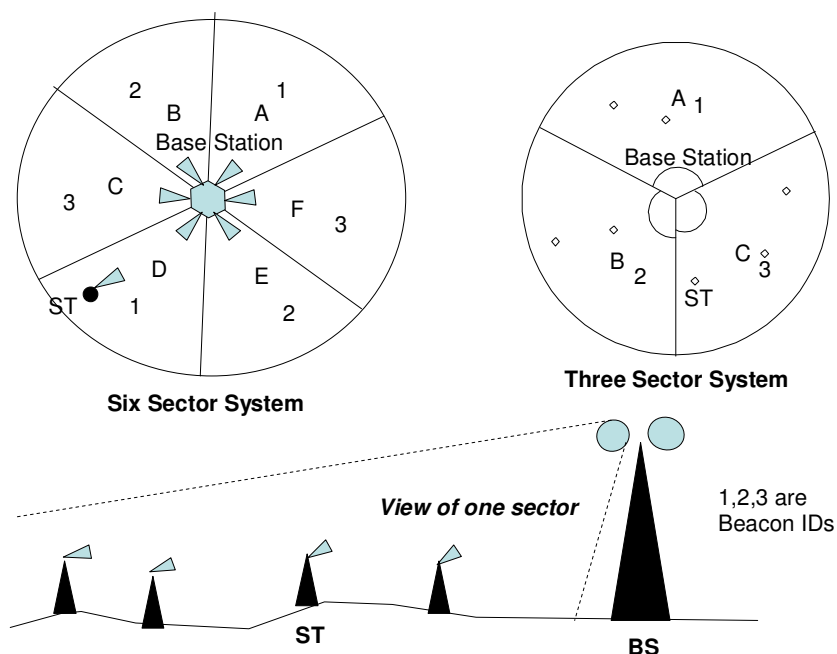


Figure 2: WiFiRe Network Configuration

One System (S), using a single IEEE 802.11b channel, will serve a cell with about 100-120 villages spread over a 15 Km radius. The cell will be sectored, with each sector containing a sectorized BS antenna. Two example configurations: (i) six sectors of 60 degrees each and (ii) three sectors of 120 degrees each, are shown in Figure 2. There will be one fixed subscriber terminal (ST) in each village, which could be connected to voice and data terminals in the village by a local area network, serving the VoIP and Internet connectivity need of the entire village. All ST(s) in a sector will associate with the BS antenna serving that sector. The ST antennas will be directional. While permitting reliable communication with the serving BS, this limits interference to/from other co-located BSs, and more importantly, to/from BSs belonging to adjacent cells.

However, because of antenna side-lobes, transmitters in each sector may interfere with receivers in other sectors. Thus, depending on the attenuation levels, a scheduled transmission in one sector may exclude the simultaneous scheduling of certain transmitter-receiver pairs in other sectors. Further, simultaneous transmissions will interfere, necessitating a limit on the number of simultaneous transmissions possible. This is explained further in section 2.5.

As a result, WiFiRe has one medium access (MAC) controller for all the sectors in a BS, to co-ordinate the medium access among them. The multiple access mechanism is time division duplexed, multi-sector TDM (TDD-MSTDM) scheduling of slots. As shown in Figure 3, time is divided into frames. Each frame is further partitioned into a downlink (DL) and an uplink (UL) segment, which need not be of equal durations. Within

each segment there are multiple slots, of equal duration each. In each DL slot, one or zero transmissions can take place in each sector. Multiple BS antennas (for different sectors) may simultaneously transmit a packet to their respective ST(s), provided they do so in a non-interfering manner. Similarly, in each UL slot, multiple ST(s) (from different sectors) may simultaneously transmit a packet to the BS, provided they do so in a non-interfering manner.

Beacons are transmitted at the start of each DL segment. The beacon for each sector contains information for time synchronization of the ST(s) in that sector, information regarding the DL and UL slot allocations (DL-MAP, UL-MAP) for that frame, and other control information. Due to site and installation dependent path loss patterns, and time varying traffic requirements, the MAP(s) need to be computed in real-time.

In order to ensure that the beacons get through to the ST(s) even under poor channel conditions, the beacons are transmitted at a lower rate (2 Mbps) than the data packets. In case of a three-sector system, the beacon for each sector is transmitted one after another, to ensure that they do not interfere. In case of a six-sector system, opposite sectors may transmit their beacons simultaneously. The order of transmission of the beacons is indicated by the numbers in Figure 2.

Note that having a 3-sector separation between beacons that are transmitted simultaneously is a conservative action. However, this is recommended since the front-to-back attenuation ratio of antenna lobes is more reliable than that of side-lobes. For subsequent data transmission, alternate sectors may transmit simultaneously, based on the interference matrix. This is explained in detail, along with a capacity analysis, in Annex B. A further general description of WiFiRe is given in section 2.

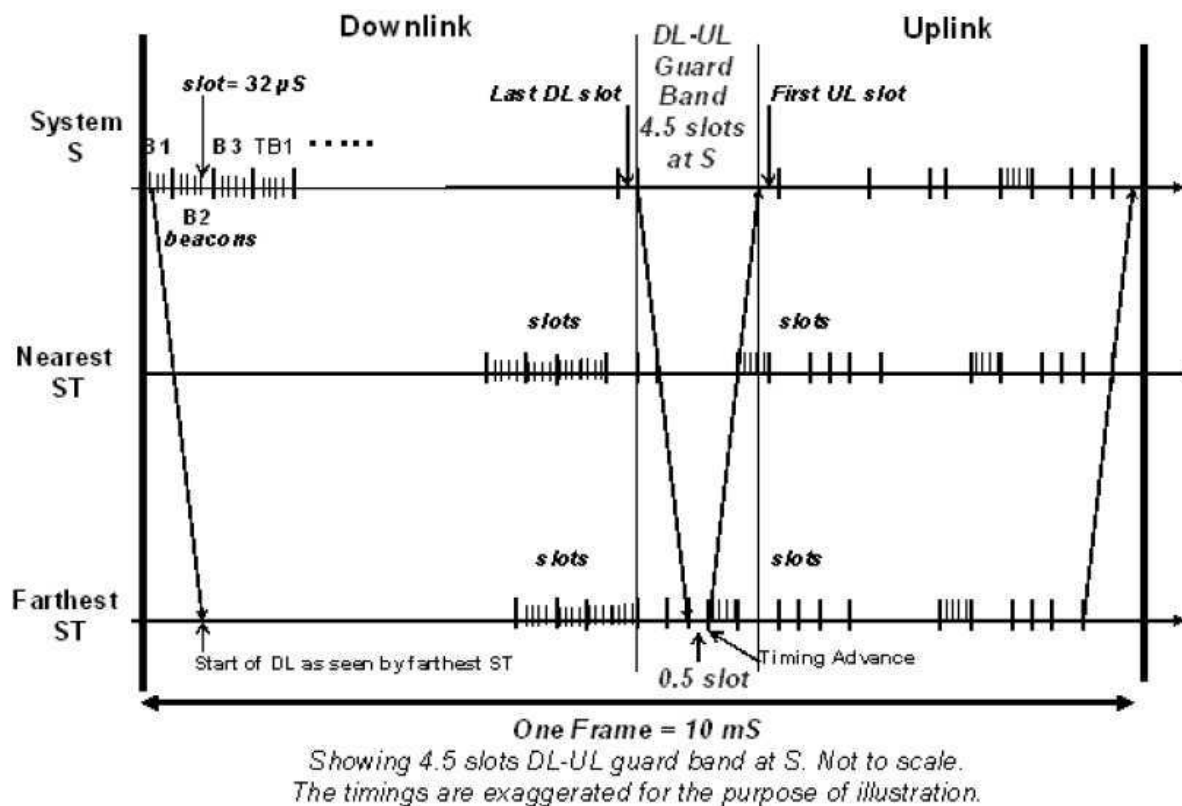


Figure 3: WiFiRe Multiple Access Mechanism

1.5 Scope

The scope of this standard is to develop a medium access control (MAC) and Physical layer (PHY) specification for WiFiRe broadband wireless connectivity for fixed stations within a rural area. In this context, a rural area is characterized by the presence of optical-fiber point-of-presence (PoP) within 15-20 km of most villages and fairly homogenous distribution of about 100-120 villages around each PoP, in the plains. The network configuration is a star topology with sectorized Base Station (BS) antennas on a tower at the PoP and a directional Subscriber Terminal (ST) antenna at each village kiosk.

Specifically, this standard

- Describes functions and services required for a WiFiRe compliant device to operate in the network.
- Defines the MAC procedures and protocols to support the data delivery services.
- Specifies the various aspects of the WiFi PHY being used.

The reference model for the layers and sub-layers of this standard are shown in Figure 4.

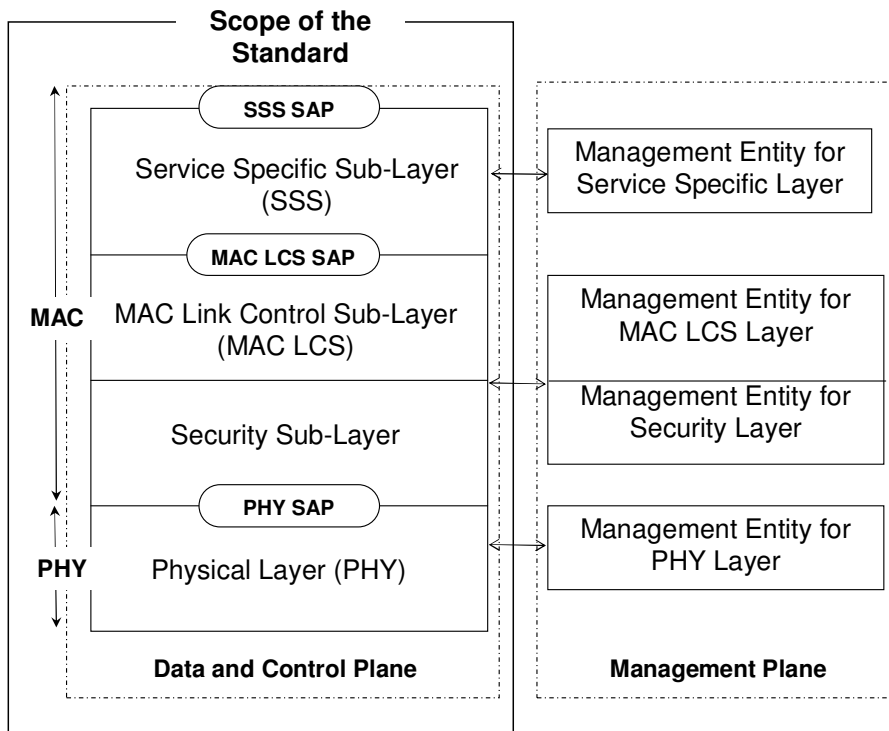


Figure 4: WiFiRe Reference Model showing the service interfaces and the scope of the standard

2 GENERAL DESCRIPTION

A WiFiRe system is one approach to design a *long-range and low-cost* fixed wireless communication network. The WiFiRe physical layer (PHY) directly employs the low-cost WiFi PHY (IEEE 802.11b, Direct Sequence Spread Spectrum). The WiFi PHY is for operation in the 2.4 GHz band and designed for a wireless local area network (LAN) with 1 Mbps, 2 Mbps and 11 Mbps data payload communication capability. It has a processing gain of at least 10 dB and uses different base-band modulations to provide the various data rates, with a typical reach of about 100 meters. WiFiRe extends the transmission range of the WiFi PHY to 15-20 Kilometers, by using a deployment strategy based on sectorized and directional antennas and line-of-sight communication.

The WiFiRe medium access control layer (MAC) replaces the WiFi MAC (IEEE 802.11b, Distributed Coordination Function) with a mechanism more suited to wide-area deployment, in terms of providing efficient access and service guarantees. The MAC is time division duplexed, multi-sector TDM (TDD-MSTDM), as described subsequently. The WiFiRe MAC is conceptually similar to the WiMax MAC (IEEE 802.16) in some respects.

2.1 Definitions and abbreviations

The various terms and abbreviations in this document are defined at the first point of their use. They are also provided collectively in the form of a glossary in section 8, for quick reference.

2.2 Design drivers and assumptions

The key design drivers for WiFiRe are as follows:

- The existence of a fiber point of presence (PoP) every 25 km or so in rural India, for backbone connectivity.
- The availability of unlicensed or free spectrum in the 2.4 GHz band.
- The low cost of WiFi chipsets. Most WiFi chipsets are designed so that the PHY and MAC layers are separate. Thus it is possible to change the MAC, or in the least, bypass it, while retaining the same PHY.
- The link margins for WiFi PHY being quite adequate for line-of-sight outdoor communication in flat terrain for 15-20 Km range.
- It being possible to have high efficiency outdoor systems, providing application service guarantees, without significantly changing radio costs. This is done by retaining the same PHY but changing the MAC, sectorization and antenna design choices and tower/site planning.

- Base Station towers of 40 m height and fixed Subscriber Terminal antennas of 10-12 m height being sufficient to cover a radius of 15-20 km from the fiber PoP for 85% of the area in rural India. This configuration can provide the required system gain with line-of-sight deployment.

The key assumptions in WiFiRe are as follows:

- The wireless links in the system are fixed, single hop, with a star topology. *Handling of mobile nodes, multi-hop wireless links and other topologies are deferred to a later release.*
- There is a fixed carrier frequency f_c and the WiFi radios are operating at 11 Mbps, except for PHY synchronization and certain control packets which may be sent at 2 Mbps.
- About 20 MHz (1 carrier) of conditionally licensed spectrum is available for niche/rural areas. The spectrum mask, power level and carrier location exactly match those for WiFi (IEEE 802.11b).
- All nodes in the system are operated by a single operator who also owns the conditional license.
- Multiple operators will use different carriers to avoid interference and frequency usage policy will be decided by co-operation among the operators.
- The PHY overhead is 192 microseconds for 1 Mbps and 96 microseconds for 2 Mbps and 11 Mbps as per WiFi (IEEE 802.11b)
- No meaningful higher layer information can be sent using the PHY overhead.
- There are no multi-path issues due to the deployment topology and the line-of-sight design.
- All the transmissions in a cell (set of co-located BS) are controlled by a single scheduler.
- All systems in adjacent cells belonging to the same operator use the same frequency, and do not interfere significantly with each other. This is made possible by the use of directional antennas at the Subscriber Terminals.
- The various components in the system have unique IP addresses.
- A single voice over IP (VoIP) packet is approximately 40 bytes. For active connections, VoIP packets are generated periodically, once in 20 milliseconds. This implies the use of a codec such as G.729, which has sampling rate of 8 Kbps. A codec such as G.711 has a sampling rate of 64 Kbps as it includes provisions for modems etc. This is not required in WiFiRe.

A more detailed discussion of the design drivers and assumptions is given in Annex A.

2.3 WiFiRe system architecture

The WiFiRe system architecture consists of several components that interact to provide a wireless wide area network (WAN) connectivity. In order to operate outdoors with a reach of 15-20 Kilometers, using the Direct Sequence Spread Spectrum (DSSS) based 802.11b PHY, WiFiRe adopts a star network topology

using directional antennas with (i) appropriate transmission power and (ii) adequate height of transmitter and receiver, for Line of Sight (LoS) connectivity.

As shown in Figure 2, a WiFiRe system consist of a set of sectorized antennas at the base station (BS), mounted on a transmission tower with a height of 40 meters and directional antennas at the subscribers terminals (ST), mounted on poles with height of around 10 meters. Typically a system is designed to cover an approximately circular area with radius of 15-20 Kms, around the tower. This area is called as a *Cell*. WiFiRe supports a link layer providing long-haul reliable connection, with service guarantees to real time and non real time data applications.

As shown in Figure 5, the key components of the WiFiRe architecture are:

- **System (S)** is a set of co-located BS (typically, six) each with a sectorized antenna, mounted atop a tower with elevation of around 40 meters, providing coverage to a cell of radius around 15-20 Km. All the transmissions in a System are coordinated by a single scheduler.
- **Base Station (BS)** of a system 'S' is a radio transceiver having the electronics for WiFi (IEEE 802.11b) physical layer. A WiFiRe BS uses a sectorized antenna, with a triangular coverage area; the exact shape of the coverage area depends on the design of the antenna and transmission power. The impact of sectorization is discussed in section 2.5.
- **Subscriber Terminal (ST)** is the user premise network equipment. An ST has a directional antenna, and it is pointed towards a System 'S'. The system S is determined at the time of deployment and fixed thereafter. In most cases, S becomes the tower for which the ST the receiving the transmission with highest signal quality level. Appropriate initialization, ranging and registration are required to ensure that a ST can communicate with one and only one BS of system S. This is discussed in section 2.6.
- **User Equipment (UE)** is a user device that connect to ST. UE(s) are source and sink of user data. WiFiRe does not specify the nature of the network media between ST and UE. They may be wired or wireless links. The service interfaces at ST provide a list of services to UE(s). This is discussed in section 2.7.

The S is connected to the external world (Internet) through the fiber PoP, while the ST is connected to voice and data terminals, through a local area network.

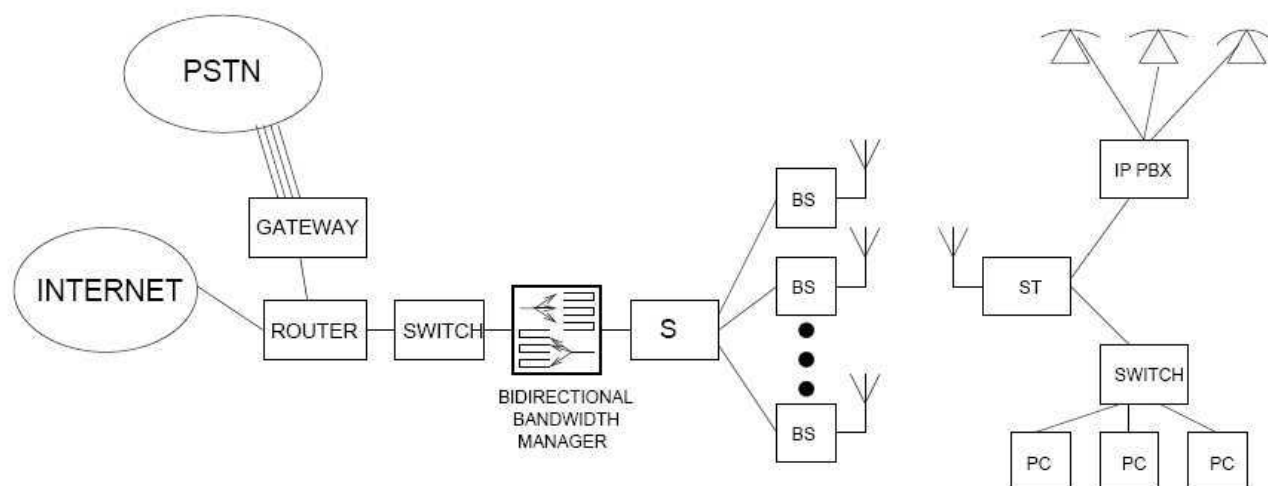


Figure 5: WiFiRe system architecture

2.4 Network initialization

The association between a ST and a System S is static. This is determined by configuration at the ST, during deployment. It is possible that a ST may hear more than one system or more than one BS of a system, depending on spatial planning of the system deployment. Appropriate topological planning and orientation of the ST directional antenna is required to ensure that a ST communicates with one and only one system S.

The association of a ST with a BS of a particular system S, is dynamic and can change during each 'power-on' scenario of the ST. Appropriate initialization, ranging and registration are required to ensure that a ST communicates with one and only one BS of system S. This association depends on antenna gain and other selection factors. Once this association is performed, it is fixed as long as the ST remains in 'power-on' mode. This is described in more detail in section 4.7.

The impact of inter-cell interference caused by neighborhood system at a ST (a common issue in cellular systems) is considered minimal since the ST directional antenna is locked onto one system and BS at the time of deployment and initialization, respectively.

2.5 Impact of sectorization

All BS in a system, use the same WiFi channel (single carrier) for communication with their respective STs. (As discussed in earlier section, operator owning another system ‘S’ shall use different frequency channel of WiFi). This is unlike typical sectorized deployments, in which co-located sectors use separate frequency channels.

In WiFiRe all the sectors in a multiple antenna configuration continue to use the same frequency channel. As a result, several interference scenario may happen: transmission by one BS may interfere withreception at BS in adjacent sectors. An ST may hear transmission from more than one BS of a system S. An ST may or may not be able receive the transmission from its BS, depending on interference caused by the neighboring sector BS antenna. Also, transmitters in one sector may interfere significantly with receivers in other sectors, because of antenna side-lobes. Hence, the MAC layer design at S includes a functionality that coordinates and manages the transmission of different BS.

A situation in which the system coverage area is partitioned into six sectors of 60 degrees each is shown in Figure 6. All ST(s) in a sector will associate with the BS antenna serving that sector. Each antenna's radiation pattern covers an additional 20 degrees on either side. Thus, depending on the attenuation levels, a scheduled transmission in one sector may exclude the simultaneous scheduling of certain transmitter-receiver pairs in other sectors. A detailed discussion on the radiation pattern for a typical BS antenna, the regions of interference and system capacity bounds is given in Annex B.

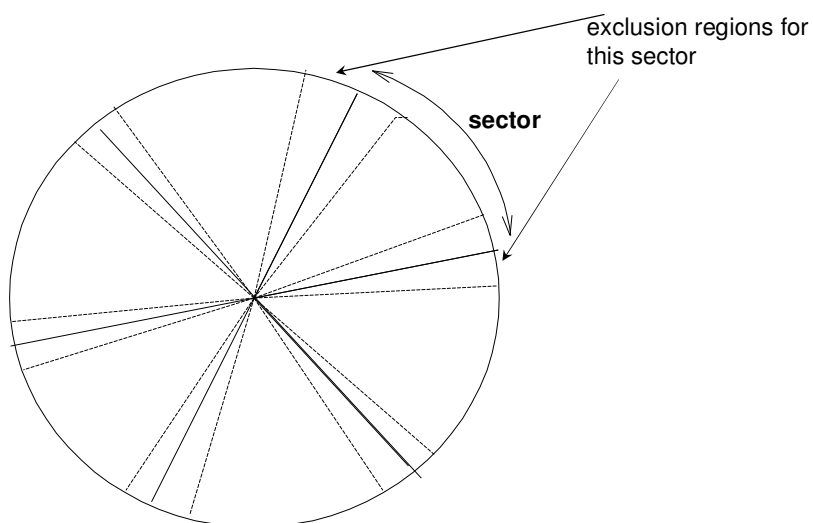


Figure 6: A simple antenna coverage and interference for six sectors

This aspect of sectorization of the coverage area while using the same frequency channel for all the sector antennas, is a key feature in WiFiRe. It not only impacts the design of the MAC protocol between transmitters and receivers, but also the scheduling policies and the system performance. During downlink transmission, a significant amount of power from the transmitting BS reaches the adjacent BS antennas, the distance separating them being very small. Hence, when a downlink transmission is scheduled in any one of the sectors, the other BS(s) cannot be in receiving mode. Hence downlink (DL) and uplink (UL) transmissions must alternate. As a result, the MAC layer avoids conflict between interfering BS antennas by using time division duplex (TDD) between the DL and UL directions (See Figure 3). The MAC scheduler at S further needs to ensure that the adjacent/interfering BS do not transmit simultaneously. Only non-interfering BS(s) may transmit simultaneously and that too in a synchronized manner. This is explained further in section 2.9.

Each BS antenna is controlled by an IEEE 802.11b PHY. The MAC layer at S is on top of all of these PHY(s), as shown in Figure 7. From the perspective of the MAC, each PHY (hence each BS antenna) is addressable and identifiable. Thus a single MAC controls more than one PHY and is responsible for scheduling MAC packets appropriately in one more PHY(s), while resolving possible transmission conflicts from the perspective of the receivers. The MAC at S can individually address each PHY and can schedule packets for transmission through any of the PHY(s), either sequentially or in parallel.




Figure 7: Single WiFiRe MAC controlling six WiFi PHY(s).

The system S broadcasts a downlink map (DL-MAP) and uplink map (UL-MAP) in specific slots (Beacons) of the downlink. These MAP(s) contain the slot allocations for the various transmissions and convey the link schedule information to the ST(s). Adjacent sectors (for example, sector 1, 2, 6 in Figure 7), resolve

interference issues by employing time-division multiplexing (TDM) within each DL and UL period. The DL and UL are non-overlapping in time. Opposite sectors (for example, sector 1 and 4 in Figure 7), are not expected to interfere with each other in a typical installation and may transmit simultaneously during DL. Opposite sectors may also receive simultaneously during UL. This leads to better resource utilization, as shown in Figure 8.

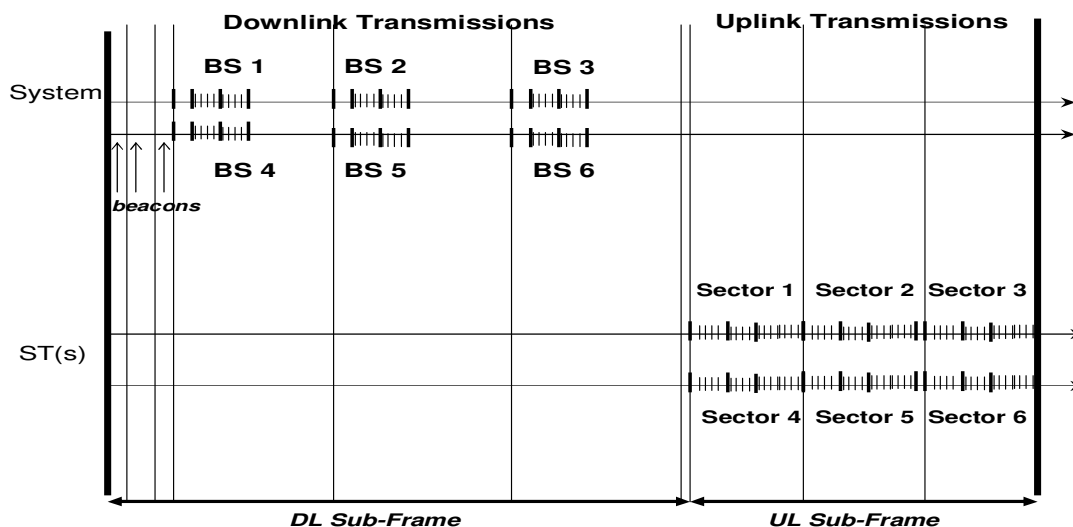


Figure 8: Parallel transmissions in a six sector system (conservative case)

The scheduler may further exploit such a situation to increase throughput by increasing the parallelism to include other non-interfering transmissions also. For example, in a particular deployment of ST(s), it may be possible to schedule parallel transmission of alternate BS(s), such as BS 1, 3 and 5. This depends upon the interference matrix determined by the ST locations and antenna radiation patterns. A discussion on scheduler design is given in Annex C.

2.6 MAC protocol overview

The MAC mechanism is a time division duplexed, multi-sector TDM (TDD-MSTDM) scheduling of slots. Time is divided into frames (See Figure 8). Each frame is further partitioned into a downlink (DL) and an uplink (UL) segment, which need not be of equal durations. The downlink - from the system S to the ST(s) - operates on a point-to-multipoint basis. The uplink - from a ST to system S - operates on a point-to-point basis. Within each segment there are multiple slots, of equal duration each. The slot duration and various timings are discussed in section 2.9.

The DL segment begins with each BS in the system transmitting a Beacon packet, in a non-interfering

manner. For example, in the six sector system shown in Figure 7, the BS(s) for sectors PHY 1 and PHY 4 may transmit their beacons (say B1 and B4) simultaneously, followed by the BS(s) for PHY 2 and PHY 5, followed by PHY 3 and PHY 6. All BS(s) are synchronized with each other; hence transmission of beacon B2 by PHY 2 starts only after completion of transmission of beacon B1 from PHY 1. Note that even though two beacons may get transmitted simultaneously (such as B1 and B4), their contents are not identical.

The beacon for each sector contains information for time synchronization of the ST(s) in that sector, information regarding the DL slot allocations (DL-MAP) and UL slot allocations (UL-MAP) for that frame, and other control information. Informally, a beacon contains <Operator ID, System ID, BS ID, All registered ST(s) scheduled for that frame and their corresponding slot assignments>. The BS ID identifies the BS (or the PHY) through which this beacon is transmitted. The structure of a beacon is given in section 4.5.

The rest of the DL transmissions follow the DL-MAP in the Beacon. In each DL slot, one or zero transmissions can take place in each sector. The DL-MAP may allow multiple non-interfering BS to simultaneously transmit a packet to the ST(s) in their respective sectors, in each slot. The DL segment ends when all the transmissions as given in the DL-MAP have been completed.

To account for propagation delays, there is a guard time of a few slots between the end of the DL segment and the start of UL segment (see section 2.9). In each UL slot, one or zero transmissions can take place in each sector, as governed by the UL-MAP. Each sector receives its own UL-MAP (contained in respective beacon for that sector), and the UL-MAP is constructed in such a way that multiple ST(s) from different sectors, may transmit in the same UL slot, provided these transmissions are non-interfering at the BS. Because of path loss patterns and time varying traffic requirements, the DL-MAP and UL-MAP need to be computed on-line. A discussion on scheduler design is given in Annex C.

The link protocol includes mechanisms that allow a ST to transmit resource (slot) reservation requests to S, for the UL and DL segments. This enables a ST to request for specific delay and bandwidth guarantees. On receipt of such resource reservation requests, the MAC layer at S executes a scheduling functionality that tries to meet the demands of the ST(s), for the next time frame. This link schedule information is captured as the DL-MAP and the UL-MAP and transmitted with the corresponding beacon. An ST listens to all the beacons from its associated BS. From the DL-MAP, the ST determines the DL slots to be monitored for its downlink data packets. From the UL-MAP, the ST determines the UL slots in which to send its data (or control) packets to the BS. Depending on the class of service, a ST may have regular slot(s) allocated in each time frame, or may be granted slot(s) by the S, after explicit resource requests. The protocol details, including the request-grant mechanism, packing, data transmission etc. are given in section 4.8 onwards.

2.7 MAC services

The WiFiRe MAC is connection-oriented. A connection defines both the mapping between peer data link processes that utilize the MAC and a service flow category. The service flow category defines the quality of service (QoS) parameters for the PDU(s) (protocol data units) that are exchanged on the connection. Each connection has a unique identifier (CID). Service flow categories provide a mechanism for uplink and downlink QoS management. Each ST adheres to a transmission protocol that controls contention and enables the service to be tailored to the delay and bandwidth requirements of each user application. This is accomplished through different types of uplink scheduling mechanisms. An ST requests uplink bandwidth (slots) on a per connection basis (implicitly identifying the service flow category).

A system S may grant bandwidth to an ST in one or more of the following ways: (i) Unsolicited bandwidth grants, (ii) Polling, and (iii) Contention Procedures. For example, real-time applications like voice and video require service on a more uniform basis and would fall in the Unsolicited bandwidth grant category, data applications that are delay-tolerant may be serviced by using the Polling mechanism and the Contention mechanism may be used when an ST has been inactive for a long period of time. These are described in more detail in section 4.8.

A default set of service flows may be provisioned when a ST is initialized. Subsequently, connections (identified with connection-ids, allocated at the time of connection creation) are associated with these service flows. Connection-id and service flow type together provide a reference against which ST request QOS for any data pipe (data being generated by a application or a set of applications). New connections may also be established when required. Connections once established may require active maintenance, depending on the type of service. For example, VoIP services are fixed demand and would require virtually no connection maintenance. On the other hand, Internet access services may require a substantial amount of ongoing maintenance due to their bursty nature and due to the high possibility of fragmentation. Finally, connections may be terminated. All connection management functions are supported through the use of static configuration and dynamic addition, modification, deletion of connections.

Also, within a scheduling interval, bandwidth may be granted by S on a per connection basis (Grant Per Connection) or as an aggregate of grants for each service flow category (Grant Per Service Flow) or as an aggregate of all grants for a ST (Grant Per Subscriber Terminal). The grant per connection would be typically used for VoIP, while the grant per service flow would be used for TCP traffic. These are described in more detail in section 4.2.

Mechanisms are defined to allow vendors to optimize system performance using different combinations of

these bandwidth allocation techniques while maintaining consistent inter-operability definitions.

2.8 MAC service interfaces

The service interfaces include the Service Specific Sub-layer (SSS), MAC Link Control Sub-layer (LCS) and the MAC Security Sub-layer. The reference model for service access points (SAP) is shown in Figure 4. A brief mention of the main services is given below. The detailed service specification is given in section 3.

The SSS should provide protocol-specific services to UE(s) for protocols such as IP, ATM, Ethernet, etc. The MAC being connection-oriented provides for higher layer peer-to-peer connection(s) between a ST and BS, with associated QoS parameters for data transport. The SSS should provide connection management and packet classification services, for mapping higher layer PDU(s) (protocol data units) to connections provided by the MAC LCS sub-layer. These services should provide following capabilities:

- *Connection Management:* The SSS should provide SAP(s) to higher layers to create and maintain higher layer peer-to-peer connection(s) between a ST and BS, with associated QoS parameters.
- *Packet Classification:* The SSS should provide SAP(s) to carry out the task of classifying higher layer PDU(s) into appropriate connections (based on some policy database), and mapping the higher layer PDU(s) to MAC PDU(s).

The SSS in turn uses the following LCS services to communicate with the peer SSS:

- *Connection Provisioning:* This includes primitives for creating and terminating MAC connections. Each connection has a unique identifier (CID).
- *Data Transport:* This includes primitives for delivery of the MAC SDU(s) (service data units) to the peer MAC entity, in accordance with the QoS associated with a connection's service flow category.
- *Security:* This includes primitives for the security sub-layer, for authentication of the end-points, and for secure transmission of the connection's PDU(s).

2.9 Typical Frame and Slot Timings

The MAC assumes that a single voice over IP (VoIP) packet, approximately 40 bytes long, will fit into one time slot of the frame. Since a single voice call may be the only traffic to/from an ST in several instances, it is important to design the MAC so that system capacity utilization is as efficient as possible even with a large number of STs with single VoIP calls. Also, VoIP packets are generated periodically, once in 20 or 30 milliseconds, for active connections. As a result, the duration of a frame is chosen as 10 milliseconds and a slot is defined as 32 microseconds. At 11Mbps, one slot corresponds to 44 bytes (which can carry a VoIP

packet); at 2 Mbps, this is 8 bytes. The PHY overhead at 1 Mbps is 6 slots (192 microseconds) and 3 slots at 2 Mbps and 11Mbps (96 microseconds). In case the VoIP packet is longer, the slot duration will need to be increased and the number of slots per frame correspondingly reduced. The Beacon carries system information using which the ST(s) can appropriately interpret the frames.

A frame corresponds to $10 * 1000 / (32) = 312.5$ slots. This is partitioned between the downlink (DL) and uplink (UL). The DL to UL ratio is to be fixed at the time of system initialization. A roughly 2:1 ratio is the default value. Hence there are 208 slots for the DL and 100 slots for UL, including overheads. As shown in Figure 9, 4.5 slots are used as guard time between the DL and UL, to account for propagation delays and to provide for transmitter-receiver turn-around at the BS radio. This gives a maximum possible range of about 24 Kms. Varying the DL to UL ratio dynamically, on a periodic or per frame basis, is optional. In this case, care needs to be taken to ensure synchronization of the BS antennas, to prevent UL, DL interference.

Beacons

Beacons are sent consecutively (for 3 adjacent sectors) at the beginning of each frame. These beacons are broadcast from the system, each by a different BS. The beacons are several slots long. Opposite sectors may transmit beacons simultaneously, when number of sectors is greater than 3.

A beacon is sent at 2 Mbps. Since a slot is defined as 32 microsecond, at 2 mbps, 8 bytes can be transmitted in one slot within a beacon. Thus a beacon is 3 slots (PHY Overhead) + 1 slot (Control Overhead) + 1 slot (DL-MAP) + 1 slot (UL-MAP), assuming 1 slot each for DL-MAP and UL-MAP is sufficient. The control overhead includes the beacon header, operator ID, etc.

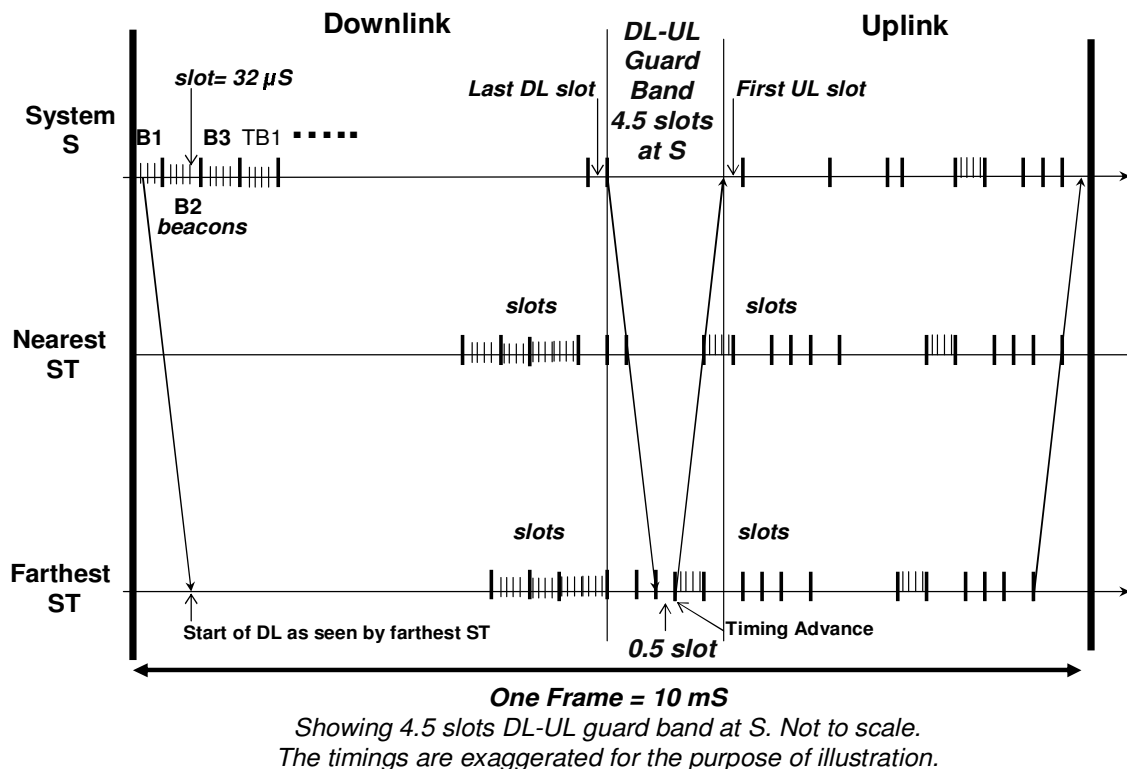


Figure 9: WiFiRe Timing Diagram

Downlink Transport Block

All downlinks, excluding the beacon, are at 11 Mbps. Transmitting a single DL slot consumes at least three additional slots as PHY overhead making it 4 slots long. Since DL is point-to-multipoint within each sector, (i) multiple MAC PDU(s) for the same ST can be combined and (ii) MAC PDU(s) for different ST can be combined, and transmitted using a single PHY overhead. This is termed as a *Downlink Transport Block* (DL-TB). The DL-TB should always begin at slot boundary and may be of variable size. However, it should fit in an integral number of slots (minimum 4) and should not exceed the maximum payload size defined by the chosen WiFi PHY; $ampDUMaxLength = 2312$ bytes, for DSSS as specified in IEEE 802.11b [11]. The DL-MAP specifies the <ST-ID> of the ST(s) for which there are packets in the current DL sub-frame. The MAC header specifies how one or more ST(s) extract one or more IP packets (including VoIP) from the DL-TB payload.

Uplink Transport Block

All uplinks are at 11Mbps. A UL slot is at least of 4 slot-duration (3 PHY overhead + 1). Since UL is point-to-point within each sector, multiple MAC PDU(s) at a given ST can be combined and transmitted using a

single PHY overhead. This is termed as a *Uplink Transport Block* (UL-TB). The UL-TB should always begin at slot boundary and may be of variable size. However, it should fit in an integral number of slots (minimum 4) and should not exceed the maximum payload size defined by the WiFi PHY. The UL-MAP specifies the <ST-ID, Slot No> mapping for which ST is to transmit in which slot. The MAC header specifies how the BS extracts one or more IP packets from the UL-TB payload. The key difference between UL-TB and DL-TB is that the UL-TB is always for one ST whereas DL-TB can be for multiple ST(s) in the same sector.

There should be a few microseconds of silence after every UL-TB to accommodate for estimation errors in Ranging (see section 4.7.1 for details of Ranging phase). This is ensured during slot allocation, depending on the fraction of last slot that is actually occupied by an ST's transmission. The MAC headers for the UL-TB and DL-TB are similar. The MAC header includes information for concatenating fractional IP packets split between the last TB of one frame's DL/UL and the first TB of the next frame's DL/UL. This is described in more detail in section 4.9.

Note that a maximum of $100 / 4 = 25$ simultaneous users can be supported per sector in one UL, when there is no spectrum reuse among the sectors. This means a payload of $25 * 2$ bytes <ST-ID, Starting slot> or 50 bytes, for the UL-MAP (and DL-MAP). For transmitting 50 bytes in UL-MAP (and in DL-MAP), the beacon size requires to be as follows:

3 slots (PHY Overhead) + 1 slot (Control Overhead) + 7 slot (DL-MAP) + 7 slot (UL-MAP),

If there are no allocations for an ST in DL-MAP (and UL-MAP), the ST may go into power-save mode.

The start of the UL may have ranging blocks (see section 2.10 for 'Ranging and Power Control"). Each ranging block is of size 8.5 slots (3 PHY overhead + 1 slot + 4.5 slots guard time). An ST-ID of all 1's in the UL-MAP indicates that the corresponding slot is a ranging block. These slots are used to transmit ranging request messages. The guard time is required to account for the propagation delay(s) between the BS and the ST and for computing the timing advance by the BS (See Figure 11).

The end of the UL may have contention slots. Each contention slot is of size 4 slots (3 PHY overhead + 1 slot). An ST-ID of all 0's in the UL-MAP indicates that the corresponding slot is a contention slot. Contention slots are used to transmit registration request messages, resource reservation messages and data for best-effort connections. There should be at least one contention slot per frame. Also, *polling slots*, specifically for transmission of resource reservation requests by an ST, may occur optionally in each frame. A polling slot should occur at least once every 50 frames. The sequence of slots for one BS is shown in

Figure 10.

Beacon (BS ID DL-MAP UL-MAP)	DL-TB(s) (1...M MAC PDU(s) for ST(s), each of size at least 4 slots)	DL-UL Guard time (4.5 slots)	Ranging Block(s) (8.5 slots) (optional)	UL-TB Polling Slot(s) (optional)	UL-TB(s) (1...N MAC PDU(s) from an ST, each of size at least 4 slots)	Contention Block(s) (Min 4 slots)
---	---	---	--	---	--	--

Figure 10: Frame structure as seen by one BS

2.10 Ranging and power control

New and un-synchronized ST(s) are allowed to Range and Register. During power-on initialization, a ST gets attached to a BS of the system S, depending on the beacons it is able to hear from the system S. On powering up a ST listens for one or more beacons from the Operator and System ID it is programmed for. There are specific time slots defined in the uplink segment for ranging. These are called ranging blocks and ranging request packets are transmitted in them. Informally, ranging request has the following information: <System ID, ST ID, BS IDs that are audible to the ST, Signal strengths of beacons from the various BS>. Based on this, the system S associates the ST with one of the BS. Then S informs the ST about the timing synchronization and BS ID that will service the ST. This is done through a ranging response packet. Upon receipt of a ranging response from a BS, the ST is live and ready to receive from and transmit data to that BS. The ranging process is shown in Figure 11 and is described in more detail in section 4.7.

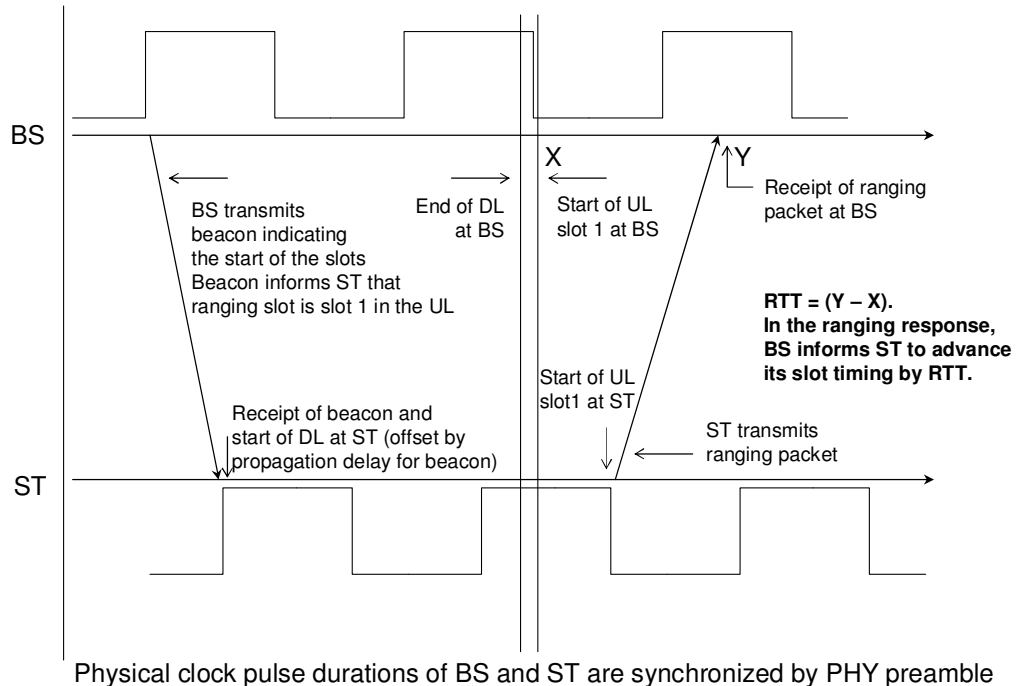


Figure 11: Ranging and Timing Advance Mechanism

The ranging response may optionally recommend the transmitter power level to be used by the ST. This may facilitate power control and better re-use across sectors. It may also contain information to enable the ST to switch to sleep modes to conserve power when needed. *The specification of protocol actions and PDU formats for power control are deferred to a later release.*

2.11 PDU formats

The details of the formats for the various protocol data units (PDU(s)) are given in section 4.3 onwards. A brief description of some of the important PDU(s) is as follows:

- *Beacon*: This contains <Operator ID; System ID; BS ID; DL-MAP; UL-MAP>. It also specifies whether a ranging block is present in the UL sub-frame.
- *Ranging Request*: This contains the <Operator ID; System ID; ST ID> of the ST sending the request. It also contains the < BS ID, Signal Strength> information for each beacon heard by the ST.
- *Ranging Response*: This contains the <BS ID, Basic CID, Primary CID, Timing Advance>. It assigns two connection identifiers - a Basic CID to the ST for periodic ranging and a Primary CID for further exchange of management messages. It also conveys the Timing Advance information to the ST to synchronize the ST transmissions with the BS slot timings.

- *Registration Request and Response*: The ST exchanges capabilities with the BS and gets assigned an IP address using these messages.
- *Dynamic Service Addition Request and Response*: The ST requests for and gets assigned a data CID using these messages. The service flow parameters are sent as Type-Length-Value tuples.
- *Dynamic Service Change Request and Response*: The ST uses these messages to either change the properties of a connection or to send a resource reservation request to the BS.
- *Data*: This contains the higher-layer data (MSDU) to be transferred from ST to BS or vice versa.

2.12 BS Scheduler functions

As mentioned earlier, a single MAC at S controls more than one PHY and is responsible for scheduling. The MAC at S can individually address each PHY and can schedule packets for transmission through any of the PHY(s) either sequentially or in parallel. The scheduler should optimally do the following:

- Simultaneously schedule multiple pairs of transmissions to/from BS(s) from/to ST(s), in a non-interfering manner.
- Appropriately combine traffic to one or more ST(s) in a sector into one DL-TB, without affecting the scheduling of other sectors.
- Assign uplink capacity keeping QoS requirements in consideration, especially the periodic nature of VoIP packets and TCP ACK(s).
- Adapt to new additions or dropout of ST(s) in the system, within a frame.

The specification of the scheduler is beyond the scope of this document. However, a detailed discussion on scheduler design is given in Annex C.

2.13 Support for multiple operators

The WiFiRe channel model requires about 20 MHz (1 WiFi Carrier) spectrum in order to provide VoIP and broadband Internet services to the users in a cell. In order to support multiple operators in an outdoor environment, a WiFiRe system operator may require conditional licensing of one channel (frequency band of 20 MHz) within the unlicensed 2.4 GHz band. The charges/fees for this channel licensing are expected to be negligible. A single operator is expected to own the conditional license and operate the site towers, in any given area. All the components (transmitters, receivers and directional antennas) belonging to an operator should use the same channel, while another operator should use a different channel. In case multiple operators are to be permitted in the same area, each operator would need to conditionally license one channel, in a non-overlapping manner. Receivers located in a coverage area of multiple antenna(s) should point towards a designated antenna during deployment time and remain locked to this tower.

2.14 Summary of protocol steps

The main steps involved in the protocol are as follows (see figure 23-25):

1. ST powers On and determines the Operator ID and System ID from its configuration.
2. ST listens for beacon messages – format is defined in section 4.5.
3. For each beacon received, ST notes the BS ID, the signal strength of the beacon and the ranging blocks as allocated in the UL-MAP.
4. ST constructs a ranging request message – format is defined in section 4.7.
5. ST determines the BS to transmit ranging request to – beacon received with highest signal strength.
6. ST waits for start of ranging block in the corresponding UL sub-frame.
7. ST transmits the ranging request message in the ranging block.
8. ST waits for ranging response – monitors DL-MAP in all beacons of the subsequent frames.
9. If no response is received within a timeout period, ST waits for a random backoff time and repeats the actions from step 6.
10. S receives the ranging request message and selects an appropriate BS and determines the timing advance to be used by the ST for being in slot synchronization with the BS.
11. S constructs a ranging response message – format is defined in section 4.7
12. S puts the ranging response in transmit queue of the corresponding BS and invokes the scheduler.
13. Scheduler (asynchronously) constructs the DL-MAP for the next frame. Transmission of the ranging response may get scheduled in the next or some other subsequent frame.
14. Scheduler may (optionally) provide a UL slot allocation (in the UL-MAP) for the registration request transmission by the ST.
15. S transmits the DL-MAP and UL-MAP in the next beacon.
16. S transmits the ranging response in appropriate DL slot.
17. ST finds its id in DL-MAP and receives the ranging response message in the corresponding slot. ST determines the basic CID and primary CID to be used for further exchanges.
18. ST constructs a registration request message.
19. ST transmits registration request in the allocated UL slot (if any) or in one of the UL contention slots and waits for a registration response. If no response is received within a timeout period, ST waits for a random backoff time and retransmits the registration request.
20. S receives the registration request and assigns an IP address to the ST, after authentication.
21. S constructs a registration response message and transmits it in the appropriate DL slot.
22. Registration is complete when the ST receives and is able to process the registration response. Now the ST has an IP address and is ready to setup data connections.
23. When the higher layer at ST has a data packet to send, it sends a dynamic service addition request message to S, in one of the polling slots or contention slots. If no response is received within a

timeout period, ST waits for a random backoff time and retransmits the request.

24. Upon receipt of the message, S assigns a data CID and responds with a dynamic service response message. The service flow and QoS parameters associated with the CID are now known to both.
25. If it is a UGS flow, the scheduler at S assigns periodic bandwidth grant in the UL sub-frame to ST.
26. If it is a rtPS or nrtPS flow, the ST requests bandwidth whenever required by sending an appropriate dynamic service change request message. Subsequently it transmits the data in the assigned slots.
27. Finally, ST transmits a dynamic service deletion message to terminate the connection.

2.15 Difference between WiMAX and WiFiRe

WiFiRe has similarity with WiMAX in terms of MAC level message exchange /handshake between System S and ST, MAC services provided to higher layers. In WiFiRe, the protocol initialization phases of Ranging, Authentication, Registration, Connection/Service, handshake mechanism for acquiring QOS follows similar procedures like WiMAX. It also attempts to use similar terminologies (e.g. TLV) from WiMAX. Above has been a choice of design in order to exploit existing implementation knowledge, if any, of high level WiMAX or similar MAC by implementer community.

The difference between WiMAX and WiFiRe is the use of available WiFi PHY with directional antennas in order to reach longer distances. WiFiRe system can use multiple antennas controlled by distinct PHY layer which in turn is governed by a single system level common MAC. Therefore, unlike WiMAX, WiFiRe MAC is a multi-sector MAC which is designed to mitigate the interferences issue caused by antenna side-lobes. Management of multiple antenna related issues and flexibility to select the number of antennas by system designer is a built-in feature of WiFiRe.

3 MAC SERVICE DEFINITION

The MAC provides a connection-oriented wireless link with provisioning to meet the QoS requirements of higher layer data streams. The information flow across the boundaries between the layers can be defined in terms of primitives that represent different items of information and cause actions to take place. These are called service access point (SAP) primitives (See Figure 4). These primitives describe the information that must necessarily be exchanged between the MAC and the higher layer to enable correct functioning of each. These primitives do not appear on the air interface but serve to define the relations of the different layers. The semantics are expressed in the parameters that are conveyed with the primitives.

The WiFiRe MAC being connection-oriented provides for higher layer peer-to-peer connection(s) between a ST and BS, with associated QoS parameters for data transport. This section defines the services provided by the MAC sub-layer(s). It does not impose message formats or state machines for these primitives.

3.1 Service Specific Sub-Layer (SSS)

The Service Specific Sub-layer (SSS) resides on top of the MAC Link Control Sub-layer (LCS). It utilizes the services provided by the LCS and in turn provides services to external higher layers. The SSS provides protocol-specific services to UE(s). It provides connection management and packet classification services for mapping higher layer PDU(s) to connections provided by the MAC LCS sub-layer.

The SSS is used for transport for all packet-based protocols such as Internet protocol (IP), point-to-point protocol (PPP), and IEEE 802.3 (Ethernet). The SSS should perform the following functions:

1. providing SAP(s) to higher layers for creating and maintaining higher layer peer-to-peer connection(s) between a ST and BS, along with associated QoS parameters.
2. accepting higher-layer PDU(s) from the higher layer protocol.
3. classification of the higher-layer PDU(s) into the appropriate MAC layer connection(s).
4. processing (if required) the higher-layer PDU(s) based on the classification.
5. mapping the higher layer PDU(s) to MAC SDU(s) (service data units).
6. delivering the MSDU(s) to the appropriate LCS SAP.
7. accepting the MSDU(s) from the peer LCS entity.
8. mapping the MSDU(s) received from peer entity into appropriate higher layer PDU(s) and delivering them to the higher layer.

For each MSDU, the sending SSS is responsible for delivering the MSDU to the LCS SAP. The LCS is responsible for delivery of the MSDU to peer LCS SAP. This is done in accordance with the QoS, fragmentation, concatenation and other functions associated with a particular connection's service flow characteristics. The LCS uses appropriate MAC management PDU(s) to get the resources (slots) required for sending the MSDU to its peer LCS. The MSDU maps to the payload part in a MAC data PDU. Finally, the receiving SSS is responsible for accepting the MSDU from the peer LCS SAP and delivering it to a higher-layer entity.

3.1.1 Classification

Classification is used to map MSDU(s) to a particular connection for transmission to its MAC peer. Different classifiers are to be used, depending on the upper layer protocol.

A classifier is a set of matching criteria applied to each packet entering the WiFiRe MAC. It consists of protocol-specific matching criteria, a classifier priority and a reference to a CID. When a packet matches a criterion, then the packet is associated with the corresponding CID and is delivered to the corresponding SAP. The service flow characteristics associated with the CID decides the QoS offered to the packet. Classifier priority is provided because overlapping matching criteria may be used by multiple classifiers. Hence priority is used to determine the order in which the classifiers will be applied to a packet.

IP classifier: IP classifiers can be based on one or more of the following fields

- ToS/DSCP bits
- Source address
- Destination address
- Protocol id
- Source port
- Destination port

Ethernet classifier: This classifier may consists of one or more of

- Source MAC address
- Destination MAC address
- Ether type

The packet classification may be done implicitly by the SSS, by examining the higher layer headers. For example, TCP ACK(s) may be identified by parsing the higher layer headers and are assigned a separate data CID. Similarly VoIP packets are identified by their size (VoIP payload is about 20 bytes; VoIP header + tolerance for power amplifier and estimation errors is less than 20 bytes). The packet classification may also be done explicitly. The explicit classification rules follow the TLV encoding (Type, Length, Value) similar to IEEE 802.16.

Further detailed specifications of the SSS primitives are deferred to a later release.

3.1.2 Concatenation

Multiple MPDU(s) may be concatenated as a single PHY PDU. Maximum size of MPDU can be 2312 bytes, as defined in IEEE 802.11 (when DSSS PHY is used). When individual MPDU(s) are smaller than maximum MPDU size, they should be concatenated as shown in Figure 12.

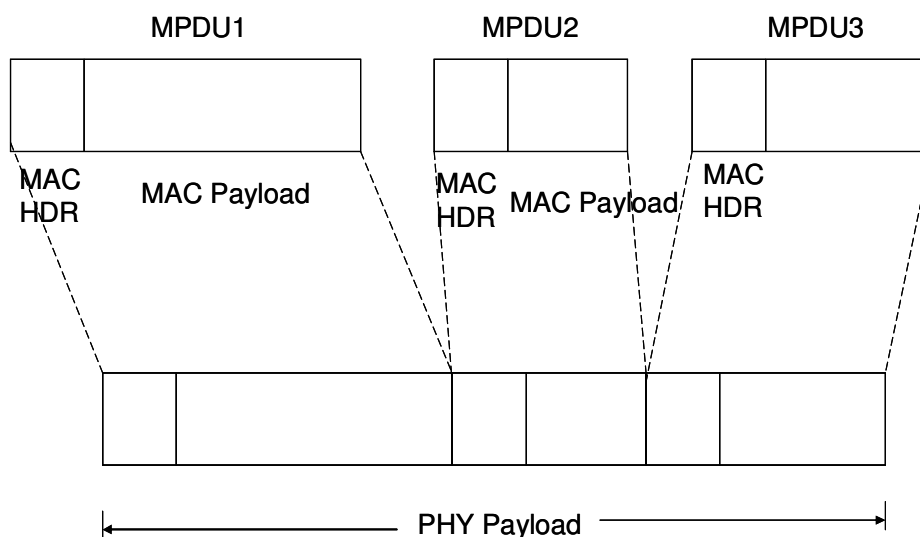


Figure 12: Concatenation of Multiple MPDU(s)

In IEEE 802.16, both concatenation and packing is recommended. In *packing* function, multiple small MAC SDU(s) (MSDU) are packed into one MAC PDU (MPDU). To keep the design simple, WiFiRe does not recommend *packing*. Instead, multiple MSDU(s) can be concatenated in a similar manner as shown above. MPDU(s) should be constructed out of each individual MSDU and concatenation operation may be performed on the MPDU(s).

ARQ is an optional mechanism in WiFiRe (see section 3.1.4). When ARQ is enabled, each MPDU will carry fragmentation sub-header which will have Fragment Control (FC) value set to “unfragmented” and will carry appropriate Fragment Sequence Number FSN value. The fragmentation procedure is described in the next sub-section (3.1.3)

3.1.3 Fragmentation

Fragmentation is the process by which a MSDU may be divided into multiple MPDU(s). The constituent fragments are then reassembled at the receiver to construct the original MSDU. Capabilities of

fragmentation and reassembly are mandatory.

The IP layer will have a MTU of 2312 bytes. For a given MPDU, if the MAC scheduler can assign enough slots, then the MPDU can be transmitted without fragmentation. But, when the scheduler assigns fewer numbers of slots, then the MPDU has to be fragmented. Each fragment will carry a *fragmentation sub-header* which will carry information required for reassembly at the receiver. The fragmentation sub-header is shown below in Figure 13 and the fragmentation of an IP packet is shown in Figure 14.

FC (Fragment Control) 2 bits	FSN (Fragment Sequence Number) 11 bits	Reserved 3 bits
---------------------------------	---	--------------------

Fragment	Fragment Control (FC) value
Unfragmented	00
First Fragment	01
Continue Fragment	10
Last Fragment	11

Figure 13: Fragmentation Sub-Header

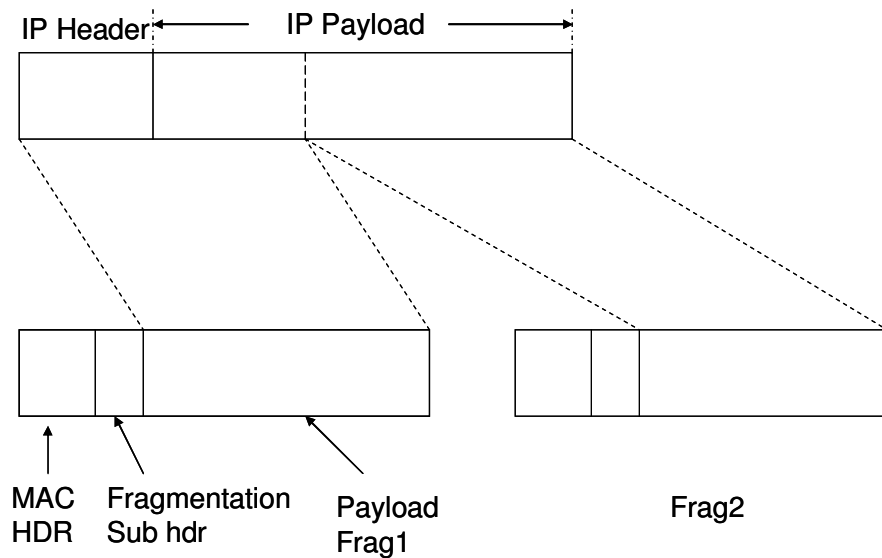


Figure 14: Fragmentation of an IP packet

3.1.4 Fragmentation and ARQ

ARQ is an optional mechanism defined under the MAC to provide link layer reliability. When an MPDU is not fragmented, but the connection has ARQ enabled, ARQ is applied as though the entire MPDU is a single fragment. The MPDU should carry fragmentation sub header with FC field set to “unfragmented” and the FSN should be set to current FSN.

For fragmented MPDU(s), ARQ is applied on each fragment. FSN is increased sequentially for the fragments. The worst case of fragmentation happens when a max size frame (2312 bytes) is allocated only one slot (40bytes) at a time. Hence the number of fragments is $2312/40 = 57.8$. Hence FSN should be at least 6 bits long.

Fragments in ARQ enabled connections are handled according to a standard ARQ process. If a fragment is not received, then the sender retransmits the fragment as per the ARQ functionality. Note that when ARQ is enabled, each fragment acts as an ARQ block. FSN across MPDU(s) of a connection increases sequentially according to ARQ rules. The ARQ rules and parameters are similar that used in IEEE 802.16.

For non-ARQ connections, fragments are sent only once and in the increasing order of FSN. If all the fragments were received correctly by the receiver, then the receiver should build the original MPDU from the fragments. FSN assigned to each fragment enables the receiver in constructing the original MPDU. If there was any loss in fragment, then the receiver should discard all the fragments of the MPDU, including the ones received subsequently until it finds a new first fragment or an un-fragmented MPDU.

3.2 Link Specific Sub-Layer (LCS)

The MAC-LCS services are used by Service Specific Sub-layer (SSS) to access the connection-oriented wireless link for data packet transport. The LCS layer provides the following categories of services:

1. Connection provision services, including creation, termination and change.
2. Data delivery services, from/to the higher layer SSS to/from the peer LCS entity.
3. Security services, including authentication and privacy.
4. Management services, for configuration of various default and power-on values.

The initial request for service from the LCS is provided by the “request” primitive. When this request is made by the initiating SSS, the initiating-side LCS constructs the appropriate Dynamic Service Request message (addition, change, or deletion; see section 3.2.1) and sends it across the wireless link to the peer (receiver-side) LCS. This peer LCS generates an “indicate” primitive to inform its SSS of the request. The peer (receiver-side) SSS entity responds with a “response” primitive to its LCS. This causes the receiver-side LCS to send an appropriate Dynamic Service Response message to its peer (initiating-side) LCS. This LCS generates a “confirm” primitive to the original requesting SSS entity. The LCS may also send a

Dynamic Service Acknowledge message to its peer, if appropriate. At any point along the way, the request may be rejected (such as due to lack of resources), terminating the protocol.

In some cases, for example when the MAC LCS on the initiating-side itself rejects the request, it is not necessary to send information to the peer entity and the “reject” primitive is issued directly by the LCS.

3.2.1 Connection Provision Services

The use of these services is to provide peer communication between System S and a ST for the purpose of creating a connection with QoS parameters. The traversal of connection request and response messages is as shown in Figure 15.




Figure 15: Primitives to request service of MAC sub-layer and generate response

The following primitives are supported:

1. **MAC_CREATE_CONNECTION.request:** This primitive is issued by a SSS entity in a system S or at a ST, to request S to dynamically set up (add) a connection. If the primitive is generated on the ST side, the receipt of this primitive causes the corresponding LCS to pass the request (in the form of a Dynamic Service Addition Request message) to its peer LCS entity in the S. The originating LCS at the ST maintains the correlation between sequence number and the requesting SSS entity.
2. **MAC_CREATE_CONNECTION.indication:** This primitive is issued by the receiver-side LCS entity to its SSS, to request the dynamic addition of a connection, (typically in response to the receipt of a Dynamic Service Addition Request message). If the LCS entity is at S, a CID is generated and the request is authenticated.
3. **MAC_CREATE_CONNECTION.response:** This primitive is issued by a receiver-side SSS entity to its LCS in response to the request for creation of a new connection. The LCS then passes on the response to its peer (initiating-side) LCS entity, in the form of a Dynamic Service Addition Response message.
4. **MAC_CREATE_CONNECTION.confirmation:** This primitive is issued by initiating-side LCS entity to its SSS entity, upon receipt of a Dynamic Service Addition Response message from its peer LCS. This informs the SSS of the status of its request and provides a CID, if the request was successful.

5. **MAC_CHANGE_CONNECTION.request:** This primitive is issued by a SSS entity in a system S or at a ST, to request S to dynamically change a connection's characteristics. For example, this may be to reflect changing bandwidth requirements. If the primitive is generated on the ST side, the receipt of this primitive causes the LCS to pass the request (in the form of a Dynamic Service Change Request message) to the LCS entity in the S.
6. **MAC_CHANGE_CONNECTION.indication:** This primitive is issued by the receiver-side LCS entity to its SSS, to request the dynamic change of a connection.
7. **MAC_CHANGE_CONNECTION.response:** This primitive is issued by a receiver-side SSS entity to its LCS, in response to the request for changing a connection. The LCS then passes on the response to its peer (initiating-side) LCS entity, in the form of a Dynamic Service Change Response message.
8. **MAC_CHANGE_CONNECTION.confirmation:** This primitive is issued by initiating-side LCS entity to its SSS, upon receipt of a Dynamic Service Change Response message from its peer LCS. This informs the SSS of the status of its connection change request.
9. **MAC_TERMINATE_CONNECTION.request:** This primitive is issued by a SSS entity in S or ST to request the termination of a connection. If the primitive is generated on the ST side, it causes the LCS to pass the request (in the form of a Dynamic Service Deletion Request message) to the LCS entity in the S.
10. **MAC_TERMINATE_CONNECTION.indication:** This primitive is issued by the receiver-side LCS entity to request the termination of a connection, in response to the receipt of a Dynamic Service Deletion Request message from its peer LCS.
11. **MAC_TERMINATE_CONNECTION.response:** This primitive is issued by the receiver-side SSS entity to its LCS, in response to a request for the termination of a connection. The LCS then passes it on to its peer (initiating-side) LCS entity in the form of a Dynamic Service Deletion Response message.
12. **MAC_TERMINATE_CONNECTION.confirmation:** This primitive confirms to an initiating SSS entity that a requested connection has been terminated.

While processing these Connection Provisioning messages, MAC LCS at S and ST also verify the QoS constraints requirement (indicated with the request). The resultant MAC connection then provides QoS guaranteed data transport service to the originator. More detail about the primitives is given in section 3.3.

3.2.2 Data Delivery Services

These services provide peer LCS entities with the ability to exchange MAC service data units (MSDU(s)). At S, the LCS determines the destination SAP, sector and associated BS antenna for a given MSDU, by

looking up a connection table; the specification of the table fields are left open for implementation choices. The LCS then uses the underlying Security sub-layer (if required) and PHY-level services to transport an MSDU to a peer LCS entity. Such asynchronous MSDU transport is performed on a QoS constrained and/or best-effort basis. An acknowledgement procedure ensures reliable delivery of MSDU.

The overview of primitives falling in this class of service by MAC LCS is as follows:

1. **MAC_DATA.request:** This primitive is issued by a SSS entity to transfer data to its LCS SAP. This causes the LCS entity to transfer the data to its peer LCS entity, in the appropriate downlink or uplink slot(s), as governed by the DL-MAP/UL-MAP and the MAC protocol.
2. **MAC_DATA.indication:** This primitive is issued by a LCS entity to transfer data from the MAC to the SSS. The specific SSS to receive the indicate message is implicit in the Connection Identifier (CID) information in the MAC header.

More detail about the primitives is given in section 3.3.

3.2.3 Security services

No SAP needs to be provided for security services. Authentication and encryption are part of only LCS. Using the Management Service primitives given in section 6, the administrator configures the security services for the BS/ST. The LCS then provides the security services without any intervention from the SSS. These services include mechanisms for:

1. Mutual authentication between BS and a ST. An appropriate authentication protocol may be used. In the simplest case, the BS may be configured to know the MAC addresses of all the ST(s) in the system. An ST sends its MAC address to the BS along with the Registration Request. The BS verifies its authenticity and then proceeds with the next steps in registration. Similarly each ST may be configured with the MAC address of the BS, which is verified by the ST upon receipt of the Registration Response.
2. Encryption of the MPDU(s) by the LCS before they are transmitted over the air and decryption by the peer LCS. Appropriate key-exchange and encryption protocols may be used.

The detailed specifications of appropriate security sub-layer primitives are deferred to a later release.

3.2.4 Management services

These services provide mechanisms for configuring various default and power-on values, including:

1. Operator ID, System ID, Time.
2. MAC address of peer entity (if required, for authentication).
3. Various keys (if required, for encryption/decryption).
4. TDD frame duration, DL to UL ratio, slot duration, no of slots for ranging, guard time, max TB size, and other values may also be made into configurable parameters. In this case, the operator must

ensure that all the active entities in a cell (system S and its ST(s)) are configured to have the same value for any given parameter. *As a result, these system parameter values need to be transmitted along with the Beacon, periodically.* Care needs to be taken to ensure one mis-configured device does not lead to inappropriate or incorrect functioning of the entire network.

The detailed specifications of appropriate management service primitives are deferred to a later release.

3.3 Detailed description of service primitives

3.3.1 **MAC_CREATE_CONNECTION.request**

3.3.1.1 Function

This primitive is issued by a SSS entity in a S or ST unit to request the dynamic addition of a connection.

3.3.1.2 Semantics of the service primitive

The parameters of the primitive are as follows:

MAC_CREATE_CONNECTION.request

(
scheduling service type,
service flow parameters,
payload header suppression indicator,
length indicator,
encryption indicator,
Fixed-length or variable-length SDU indicator,
SDU length (only needed for fixed-length SDU connections),
CRC request,
ARQ parameters,
sequence number
)

The scheduling service type (see section 2.7) is one of the following: Unsolicited bandwidth grant (UGS), Polling (PS), and Contention or best effort (BE) service. The service flow parameters include details such as peak and average rate. These parameters are the same as those in the Dynamic Service Change Request message. The payload header suppression indicator specifies whether the SDUs on the service flow are to have their headers suppressed. The fixed-length or variable-length SDU indicator specifies whether the SDUs on the service flow are fixed-length or variable-length. The SDU length specifies the length of the SDU for a fixed-length SDU service flow. The encryption indicator specifies that the data sent over this connection is to be encrypted, if ON. No encryption is used, If OFF. Cyclic redundancy check (CRC) request, if ON, requests that the MAC SDUs delivered over this connection are transported in MAC PDUs with a CRC added to them. The automatic repeat request (ARQ) parameters are: whether or not

ARQ is used for the connection and the maximum retransmission limit. As specified in section 2.6, selective-ARQ is to be used. The sequence number is used to correlate this primitive with its response from the *S* via the MAC.

3.3.1.3 *When Generated*

This primitive is generated by a SSS of an *S* or *ST* unit to request the *S* to set up a new connection.

3.3.1.4 *Effect of Receipt*

If the primitive is generated on the *ST* side, the receipt of this primitive causes the MAC to pass the request (in the form of a Dynamic Service Addition Request message) to the MAC entity in the BS. The *ST* MAC remembers the correlation between sequence number and the requesting SSS entity. If the primitive is generated on the *S* side, the *S* checks the validity of the request and, if valid, chooses a CID and includes it in the Dynamic Service Addition Request message sent to the *ST*. This CID shall be returned to the requesting SSS via the CONFIRM primitive. If the primitive originated at the *ST*, the actions of generating a CID and authenticating the request are deferred to the INDICATION/RESPONSE portion of the protocol.

3.3.2 **MAC_CREATE_CONNECTION.Indication**

3.3.2.1 *Function*

This primitive is sent by the receiver-side (non-initiating) MAC entity to the SSS, to request the dynamic addition of a connection in response to the MAC sublayer receiving a Dynamic Service Addition Request message. If the non-initiating MAC entity is at *S*, a CID is generated and the request is authenticated.

3.3.2.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_CREATE_CONNECTION.indication

```
(
service type,
service flow parameters,
sequence number
)
```

Parameters: See MAC_CREATE_CONNECTION.request.

The encryption and CRC flags are not delivered with the indication primitive since the lower layers would have already acted on it to decrypt the data or to check a CRC, before the MAC SDU is passed up to the SSS.

3.3.2.3 *When Generated*

This primitive is generated by the MAC sublayer of the non-initiating side of the protocol when it receives a Dynamic Service Addition Request message from the initiating side of the connection.

3.3.2.4 *Effect of Receipt*

When the SSS receives this primitive, it checks the validity of the request from the point of view of its own resources. It accepts or rejects the request via the MAC_CREATE_CONNECTION.response primitive. If the connection request has originated on the *ST* side, the *S* sends the CID to the *ST* side in this RESPONSE primitive. Otherwise, if the origin was *S* itself, the RESPONSE contains the CID in the DSA header bearing the indication.

3.3.3 **MAC_CREATE_CONNECTION.response**

3.3.3.1 *Function*

This primitive is issued by a non-initiating MAC entity in response to a MAC_CREATE_CONNECTION.indication requesting the creation of a new connection.

3.3.3.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_CREATE_CONNECTION.response

(
 Connection ID,
 response code,
 response message,
 sequence number,
 ARQ parameters
)

The Connection ID is returned to the requester for use with the traffic specified in the request. If the request is rejected, then this value shall be ignored. The response code indicates success or the reason for rejecting the request. The response message provides additional information to the requester, in type-length-value (TLV) format. The sequence number is returned to the requesting entity to correlate this response with the original request.

The ARQ parameters are: whether or not ARQ is used for the connection, maximum retransmission limit and acknowledgment window size.

3.3.3.3 *When Generated*

This primitive is generated by the non-initiating SSS entity when it has received a MAC_CREATE_CONNECTION.indication.

3.3.3.4 *Effect of Receipt*

The receipt of this primitive causes the MAC sublayer to send the Dynamic Service Addition Response message to the requesting MAC entity. Once the Dynamic Service Addition Acknowledgement is received, the MAC is prepared to pass data for this connection on to the air link.

3.3.4 **MAC_CREATE_CONNECTION.confirmation**

3.3.4.1 *Function*

This primitive confirms to a convergence entity that a requested connection has been provided. It informs the *ST* or *S* of the status of its request and provides a CID for the success case.

3.3.4.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_CREATE_CONNECTION.confirmation

```
(
Connection ID,
response code,
response message,
sequence number
)
```

Parameters: see MAC_CREATE_CONNECTION.response.

3.3.4.3 *When Generated*

This primitive is generated by the initiating side MAC entity when it has received a Dynamic Service Addition Response message.

3.3.4.4 *Effect of Receipt*

The receipt of this primitive informs the convergence entity that the requested connection is available for transmission requests.

3.3.5 **MAC_Terminate_CONNECTION.request**

3.3.5.1 *Function*

This primitive is issued by a *SSS* entity in a *S* or *ST* unit to request the termination of a connection.

3.3.5.2 *Semantics of the service primitive*

MAC_TERMINATE_CONNECTION.request

```
(
```

Connection ID

Sequence number

) The Connection ID parameter specifies which connection is to be terminated.

3.3.5.3 *When Generated*

This primitive is generated by a SSS of a *S* or *ST* unit to request the termination of an existing connection.

3.3.5.4 *Effect of Receipt*

If the primitive is generated on the *ST* side, the receipt of this primitive causes the MAC to pass the request to the MAC entity in the *S* via the Dynamic Service Deletion Request message. The *S* checks the validity of the request, and if it is valid it terminates the connection. If the primitive is generated at *S*, it has already been validated and the MAC at *S* informs the *ST* by issuing a Dynamic Service Deletion Request message.

3.3.6 **MAC_Terminate_CONNECTION.indication**

3.3.6.1 *Function*

This primitive is issued by the MAC entity on the non-initiating side to request the termination of a connection in response to the receipt of a Dynamic Service Deletion—Request message.

3.3.6.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_TERMINATE_CONNECTION.indication

(

Connection ID

)

The Connection ID parameter specifies which connection is to be terminated.

3.3.6.3 *When Generated*

This primitive is generated by the MAC sublayer when it receives a Dynamic Service Deletion—Request message to terminate a connection, or when it finds it necessary for any reason to terminate a connection.

.

3.3.6.4 *Effect of Receipt*

If the protocol was initiated at the *ST*, when it receives this primitive, the *S* checks the validity of the request. In any case, the receiving SSS returns the MAC_TERMINATE_CONNECTION.response primitive and deletes the CID from the appropriate polling and scheduling lists.

3.3.7 MAC_Terminate_CONNECTION.response

3.3.7.1 Function

This primitive is issued by a SSS entity in response to a request for the termination of a connection.

3.3.7.2 Semantics of the service primitive

The parameters of the primitive are as follows:

MAC_TERMINATE_CONNECTION.response

```
(
Connection ID,
response code,
response message
sequence number
)
```

The Connection ID is returned to the requesting entity. The response code indicates success or the reason for rejecting the request. The response message provides additional information to the requester, in TLV format.

3.3.7.3 When Generated

This primitive is generated by the SSS entity when it has received a MAC_TERMINATE_CONNECTION.indication from its MAC sublayer.

3.3.7.4 Effect of Receipt

The receipt of this primitive causes the MAC sublayer to pass the message to the initiating side via the Dynamic Service Deletion—Response message. The initiating MAC in turn passes the CONFIRM primitive to the requesting convergence entity. The convergence entity shall no longer use this CID for data transmission

3.3.8 MAC_Terminate_CONNECTION.confirmation

3.3.8.1 Function

This primitive confirms to a convergence entity that a requested connection has been terminated.

3.3.8.2 Semantics of the service primitive

The parameters of the primitive are as follows:

MAC_TERMINATE_CONNECTION.confirmation

```
(
Connection ID,
response code,
```

response message

sequence number

) Parameters: see MAC_TERMINATE_CONNECTION.response.

3.3.8.3 *When Generated*

This primitive is generated by the MAC entity when it has received a Dynamic Service Deletion—Response message.

3.3.8.4 *Effect of Receipt*

The receipt of this primitive informs the convergence entity that a connection has been terminated. The convergence entity shall no longer use this CID for data transmission.

3.3.9 ***Changing a Connection***

The following primitives are used:

MAC_CHANGE_CONNECTION.request

MAC_CHANGE_CONNECTION.indication

MAC_CHANGE_CONNECTION.response

MAC_CHANGE_CONNECTION.confirmation

The semantics and effect of receipt of these primitives are the same as for the corresponding CREATE primitives, except that a new CID is not generated.

3.3.10 ***MAC_Data.request***

3.3.10.1 *Function*

This primitive defines the transfer of data to the MAC entity from a SSS SAP.

3.3.10.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_DATA.request

(
 Connection ID,
 length,
 data,
 discard-eligible flag,
 encryption flag

)

The Connection ID parameter specifies the connection over which the data is to be sent; the service class is implicit in the Connection ID. The length parameter specifies the length of the MAC SDU in bytes. The

data parameter specifies the MAC SDU as received by the local MAC entity. The discard-eligible flag specifies whether the MAC SDU is to be preferentially discarded in the event of link congestion and consequent buffer overflow. The encryption flag specifies that the data sent over this connection is to be encrypted, if ON. No encryption is used, if OFF.

3.3.10.3 *When Generated*

This primitive is generated by a SSS whenever a MAC SDU is to be transferred to a peer entity or entities.

3.3.10.4 *Effect of Receipt*

The receipt of this primitive causes the MAC entity to process the MAC SDU through the MAC sublayer and pass the appropriately formatted PDUs to the PHY transmission SSS for transfer to peer MAC sublayer entities, using the CID specified.

3.3.11 **MAC_Data.indication**

3.3.11.1 *Function*

This primitive defines the transfer of data from the MAC to the SSS. The specific SSS to receive the indicate message is implicit in the CID.

3.3.11.2 *Semantics of the service primitive*

The parameters of the primitive are as follows:

MAC_DATA.indication

```
(
Connection ID,
length,
data,
reception status,
encryption flag
)
```

The Connection ID parameter specifies the connection over which the data was sent. The length parameter specifies the length of the data unit in bytes. The data parameter specifies the MAC SDU as received by the local MAC entity. The reception status parameter indicates transmission success or failure for those PDUs received via the MAC_DATA.indication.

3.3.11.3 *When Generated*

This primitive is generated whenever an MAC SDU is to be transferred to a peer convergence entity or entities.

3.3.11.4 *Effect of Receipt*

The effect of receipt of this primitive by a convergence entity is dependent on the validity and content of the MAC SDU. The choice of SSS is determined using the CID over which the MAC SDU was sent.

4 MAC DETAILED DESCRIPTION

The addressing, protocol actions at the ST and BS, and PDU formats (Protocol Data Units) are specified in this subsection. During the interaction between a ST and a BS, the MAC PDU(s) exchanged between them fall under three categories: (i) Network Initialization, (ii) Connection Management and (iii) Data Transport. All stations shall be able to properly construct PDU(s) for transmission and decode PDU(s) upon reception.

4.1 Addressing and connection identification

Each ST shall have a 48-bit universal MAC address. This address uniquely defines the ST from within the set of all possible vendors and equipment types. It is used during the registration process to establish the appropriate connections for an ST. It is also used as part of the authentication process by which the BS and ST each verify the identity of each other.

Connections are identified by a 16-bit Connection Identifier (CID). The use of a 16-bit CID permits a total of 64K connections within each downlink and uplink channel. The CID serves as a pointer to context and destination information. It is assigned even for nominally connectionless traffic like IP. The type of service may be implicitly specified in the CID itself. In order to avoid the overhead in creating and deleting the context for a CID, many higher-layer sessions may use the same CID over a period of time, sequentially one after another.

At ST initialization, two management connections in each direction (uplink and downlink) shall be established between the ST and the BS. These CID(s) shall be assigned in the Ranging Response messages and reflect the fact that there are inherently two different types of management traffic between an ST and the BS. One of them is the *basic CID*, used by the System S MAC and ST MAC to exchange short, time-urgent MAC management messages, such as ranging. The other is the *primary CID*, used by the BS MAC and ST MAC to exchange longer, more delay tolerant MAC management messages, such as creation of data connections. When the higher layer at BS or ST requests for a data connection as per one of the supported service flow types, a *data CID* is assigned by S to that connection. Since the MAC is connection-oriented, there are as many *data CID(s)* as there are active data connections, at any given point of time.

The reason for having different types of CID(s) is mainly to facilitate the QoS scheduler. A scheduler could give different levels of importance to the messages in the queue(s) depending on the connection CID(s).

The format of the CID is shown in

Figure 16. The first two bits implicitly identify the type of the CID: (00) implies it is a basic CID; (01) implies primary CID; both (10) and (11) imply data CID. In case of data CID, the next two bits implicitly identify the type of the associated service flow: (00) for UGS, (01) for rtPS, (10) for nrtPS and (11) for BE. (The exact semantics of these types are defined in section 4.2).

Type of CID (2 bits)	Type of associated service flow (2 bits)	Identifier (12 bits)
-------------------------	---	-------------------------

Figure 16: CID format

4.2 Bandwidth Request Grant Service

The following specifies how the uplink is scheduled for bandwidth requests from ST(s) and how bandwidth grants are provided to ST(s).

4.2.1 Types of services

WiFiRe provides following types of bandwidth request services:

- Unsolicited Grant Service:* The Unsolicited Grant Service (UGS) is designed to support real-time flows that generate fixed size data packets on a periodic basis, such as T1/E1 and Voice over IP without silence suppression. When a data CID is associated with UGS service flow type, the ST does not have to send periodic bandwidth request to the BS for that connection (data CID). The UGS service offers fixed size grants on a real-time periodic basis, which eliminate the overhead and latency of ST requests and assure that grants are available to meet the flow's real-time needs. The BS shall provide fixed size data grant slots at periodic intervals to the service flow. The UGS shall be specified using the following parameters: the Unsolicited Grant Size, the Nominal Grant Interval, the Tolerated Grant Jitter, and the Request/Transmission Policy.
- Real-time Polling Service:* The Real-Time Polling Service (rtPS) is designed to support real-time flows that generate variable size data packets on a periodic basis, such as MPEG video. The service offers real-time, periodic, unicast request opportunities, which meet the flow's real-time needs and allow the ST to specify the size of the desired grant. This service requires more request overhead than UGS, but supports variable grant sizes for optimum data transport efficiency. The BS shall provide periodic unicast request opportunities, by assigning appropriate *polling slots* in the uplink. The key service information elements are the Nominal Polling Interval, the Tolerated Poll

Jitter, and the Request/Transmission Policy.

- *Non real time Polling Service*: The Non-Real-Time Polling Service (nrtPS) is designed to support non real-time flows that require variable size data grant slots on a regular basis, such as high bandwidth FTP. The service offers unicast polls on a regular basis, which assures that the flow receives request opportunities even during network congestion. The BS typically polls nrtPS CIDs on an interval (periodic or non-periodic). The BS shall provide timely unicast request opportunities by assigning appropriate *polling slots* in the uplink. The key service elements are Nominal Polling Interval, Minimum Reserved Traffic Rate, Maximum Sustained Traffic Rate, Request/Transmission Policy, and Traffic Priority.
- *Best Effort Service*: The intent of the Best Effort (BE) service is to provide efficient service to best effort traffic. Request/Transmission Policy setting should be such that the ST is allowed to use contention request opportunities. The key service elements are the Minimum Reserved Traffic Rate, the Maximum Sustained Traffic Rate, and the Traffic Priority.

4.2.2 Types of Grants

Regarding the grant of the bandwidth requested, there are three modes of operation:

- *Grant per Connection mode (GPC)*: In GPC, the BS grants bandwidth explicitly to each connection.
- *Grant per Subscriber Terminal mode (GPST)*: In GPST, the bandwidth is granted collectively to all the connections belonging to an ST. This allows for smaller UL-MAP(s) and provides freedom to ST to make real-time scheduling decisions and perhaps utilize the bandwidth differently than it was originally granted by the BS.
- *Grant per Service Flow type (GPSF)*: GPST is an intermediate between GPC and GPST. In GPST, the bandwidth is granted collectively to all the connections *of a particular flow type* belonging to an ST. This avoids the need for transmitting a detailed UL-MAP as in GPC. It also avoids the need for a complex scheduler at ST as in GPST.

The BS and ST exchange capabilities and agree upon the type of grant mechanism during registration. Also, in case of GPC or GPSF, the BS uses the ST-id field in the UL-MAP to inform the ST about the CID or flow type for which it has allocated slots in the uplink.

4.2.3 Polling process

Polling is the process by which the BS allocates bandwidth to the ST(s), specifically for the purpose of making bandwidth requests. These allocations may be to an individual connection at an ST, a group of connections at an individual ST or to a group of ST(s). These are indicated as Polling Slots in the UL-MAP.

When a connection is polled individually, no explicit message is transmitted for polling it. Instead, the ST is

allocated (in the UL-MAP), sufficient bandwidth in order to transmit a bandwidth request for that connection. If the ST does not need bandwidth for that connection, it returns stuff bytes (0xFF).

When a ST is polled individually, no explicit message is transmitted for polling it. Instead, the ST is allocated (in the UL-MAP), sufficient bandwidth in order to transmit a bandwidth request for some of its connections. ST decides the choice of connections based on the service flow type associated with them. If the ST does not need bandwidth for any of its data connections, it returns stuff bytes (0xFF). ST(s) operating in GPST mode that have an active UGS connection of sufficient bandwidth shall not be polled individually unless they set the Poll Me (PM) bit in the header of a packet on the UGS connection. This saves bandwidth over polling all ST(s) individually. Similarly, a More Data (MD) bit is set in the header of an active uplink transmission whenever the ST wants to be polled for rtPS and nrtPS connections.

When the allocation is to a group of ST(s), it actually defines the bandwidth request polling slot(s) among that group. The BS may schedule one or more of the polling slot(s) in uplink to be shared by many ST(s) to transmit bandwidth requests. An ST may randomly choose one of these slots to transmit its request. In case the ST does not receive the bandwidth grant correspond to this request within a timeout, it assumes that there was a bandwidth request collision. In this case, a standard backoff algorithm is used. This backoff algorithm is similar to that defined for timed-out ranging and registration requests (section 4.7.1).

4.3 MAC PDU format

MAC PDU(s) are of the form illustrated in Figure 17. Each PDU shall begin with a fixed-length Generic MAC Header. The header may be followed by the Payload of the MAC PDU. If present, the Payload shall consist of zero or more sub-headers and zero or more MAC SDU(s) (Service Data Units). The payload information may vary in length, so that a MAC PDU may represent a variable number of bytes. A MAC PDU may contain a CRC (Cyclic Redundancy Check). The maximum size of a single MAC PDU is bounded by the maximum size payload accepted by the WiFi PHY. Larger MPDU(s) may be fragmented and transmitted.




Figure 17: MAC PDU format

4.4 MAC header format

Two MAC header formats are defined – Generic MAC Header and Beacon Header. The Generic MAC Header is used for Management and Data PDU(s). The Beacon Header used to transmit a beacon message. The single-bit Header Type (HT) field distinguishes the Generic and Beacon Header formats. The HT field shall be set to zero for the Generic Header and to one for a Beacon Header.

Additionally, there may be sub-Header(s) defined for packing multiple MAC SDU(s) into a single MAC PDU.

4.4.1 Generic MAC header

The fields of the Generic MAC header are as shown in Figure 18.

HT	Len	Type	CID	Reserved
----	-----	------	-----	----------

Figure 18: Generic MAC Header

Generic MAC header:

HT is 1 bit header type and is set to 0 for Generic MAC header.

Len is 7 bits and represents length of the MAC PDU including header length.

CID is 2 bytes and represents Connection to which the MPDU belongs to.

Reserved is 1 byte and is reserved for future use. It may carry information regarding presence or absence of CRC in the MPDU, authentication, encryption etc.

Type is 1 byte and has the following Values:

0x00 – no sub headers present

0x01 – sub header present

0x03- Management PDU of type Initial Ranging Request

0x04- Management PDU of type Initial Ranging Request Response

0x05- Management PDU of type Registration Request

0x06- Management PDU of type Registration Response

0x07- Management PDU of type Dynamic Service Addition Request

0x08- Management PDU of type Dynamic Service Addition Response

0x09- Management PDU of type Dynamic Service Change Request

0x10- Management PDU of type Dynamic Service Change Response

0x11- Management PDU of type Dynamic Service Deletion Request

0x12- Management PDU of type Dynamic Service Deletion Response

0x14 – MAC PDU with data Payload

4.4.2 Beacon header

The fields of the Beacon header are as shown in Figure 19.

HT	Len	Reserved
-----------	------------	-----------------

Figure 19: Beacon Header

Beacon Header

HT is 1 bit and the value is set to 1.

Len is 7 bits and represents length of the Beacon including header length.

Reserved is 1 byte and reserved for future use.

4.5 MAC Management PDU(s)

MAC Management messages are carried in the Payload of the MAC PDU. The type of MAC Management Message is specified in Type field of generic MAC header. MAC management messages on the Basic, Broadcast, and Initial Ranging connections shall neither be fragmented nor packed. MAC management messages on the Primary Management connection may be packed. MAC management messages shall not be carried on the data transport connections.

As mentioned earlier, the MAC procedures can be categorized mainly under: (i) Network Initialization, (ii) Connection Management and (iii) Data Transport. Each of these phases involves management messages to be exchanged between ST and S. The control plane includes (i) and (ii) while the Data plane includes (iii). These messages exchanged in these phases are described below. The detailed description of MAC procedures in ST and in S is provided in the subsequent sections.

4.5.1 Beacon Message

The format of the Beacon Message is as shown in Figure 20.

Header	Opr ID	Sys ID	BS ID	Rng Slot	DL MAP	UL MAP
---------------	---------------	---------------	--------------	-----------------	---------------	---------------

Figure 20: Beacon Message

Beacon

Header is the 2 bytes defined earlier.

Opr ID is a 1 byte value identifying the Operator of the network.

Sys ID is a 1 byte a value identifying the System (S).

BS ID is a 7 bits value identifying the BS in the System that is transmitting this Beacon.

DL-MAP is 50 bytes. It is a 50 element vector of <ST ID = (1 byte)>. ST ID = 0x11 value implies that the message in the corresponding DL slot is a broadcast message for all ST(s).

UL-MAP is 50 bytes. It is a 25 element vector of <ST ID = (1 byte), Slot id= (1 byte)>.

Rng Slot is a 1 bit indicating if there are any ranging blocks allocated in the UL-MAP. The value is set to 1 if any ranging block is present in the UL sub-frame, 0 otherwise. Ranging block information is transmitted in the first few entries of the DL-MAP. These entries are identified by a specific value in the ST ID field of the DL-MAP vector. Ranging block(s) (if present) are always the starting slot(s) of the UL. An ST-ID of all 1s is used to denote a ranging block, while an ST-ID of all 0's is used to denote a contention slot.

4.5.2 Network Initialization Messages

In addition to the Beacon, the management messages used in this phase are:

- 1) Initial Ranging Request
- 2) Initial Ranging Response
- 3) Initial Authentication Request
- 4) Initial Authentication Response
- 5) Registration Request
- 6) Registration Response

The PDU Formats for these messages are given along with the network initialization procedure in section 4.7.

4.5.3 Connection Management Messages

The management messages used in this phase are:

- 1) Dynamic Service Addition Request
- 2) Dynamic Service Addition Response
- 3) Dynamic Service Change Request
- 4) Dynamic Service Change Response
- 5) Dynamic Service Deletion Request
- 6) Dynamic Service Deletion Response

The PDU Formats for these messages are given along with the connection management procedure in section 4.8.

4.6 MAC Data PDU(s)

The format of the MAC Data PDU is as shown in Figure 21.

Gen. Header (5 bytes)	Sub Header (Optional)	Payload (MAC SDU(s)) (Maximum 2312 bytes)	CRC (optional)
---------------------------------	---------------------------------	---	--------------------------

Figure 21: MAC Data PDU

MAC Data PDU

Header is the 5 bytes Generic MAC header defined earlier.

Payload is the MSDU(s). On the uplink, multiple MSDU(s) of one or more data connections in a ST, may be packed into one MAC Data PDU. On the downlink, multiple MSDU(s) of one or more data connections to the *same or different ST(s)*, may be packed into one MAC Data PDU. The maximum size of a MSDU is 2312 bytes.

CRC is the (optional) cyclic redundancy check. The reserved bits in the Header are used to indicate the presence or absence of CRC.

4.7 Network Initialization sub-procedures

Network Initialization consists of Ranging, Authentication, and Registration sub-procedures. These are described below.

4.7.1 Ranging

Upon completion of power-up sequence and self-initialization, the ST enters the process of 'Ranging' in order to synchronize clock and other physical layer parameters with system S. This is important since the links can be over several Kilometers (RF propagation latency can be in the order of 50-100 micro-seconds). Ranging is also performed periodically so that the ST is kept in-sync with S.

The System S periodically transmits a *beacon*, having the structure described in section 4.5.1. An ST first listens for a beacon and then sends a *ranging request*. The System S then sends a *ranging response*. In the ranging response, the System S assigns the ST two connection-IDs (CIDs) called the *primary CID* and the *basic CID*. The primary CID is used for further exchange of management messages, while the basic CID is used for further periodic ranging exchanges. The detailed steps involved in ranging are as follows:

1. An ST in the Ranging phase detects the beacons from the System and Operator ID it is configured

for. From the UL-MAP in the beacon, it determines the location of ranging block(s) in the UL sub-frame. Ranging block information is transmitted in the first few entries of the DL-MAP. Ranging block(s) if present, are always the starting slot(s) of the UL. These are identified by a specific value in the ST ID field of the UL-MAP vector. There should be at least one Ranging block per TMax seconds. Slotted ALOHA is to be used as the multiple access mechanism. A standard backoff mechanism is to be used in case of collision.

2. The ST constructs an Initial Ranging Request (IRR) PDU having the structure given below:

Initial Ranging Request:

MAC generic Header: with type field = 0x04

Operator ID : 1 byte

System ID: 1 byte

ST ID: 4 byte (MAC address)

BS ID 1: 1 byte

SignalStrength1 = 2 bytes

BS ID 2: 1 byte

SignalStrength2 = 2 bytes

BS ID 3: 1 byte

SignalStrength3 = 2 bytes

CID = 2 bytes

Bkoff = 1 byte

Field Description:

Operator ID is a configured value at ST indicating which operator the ST should associate with.

System ID is a configured value at ST indicating which System the ST should associate with.

ST ID is the MAC address of ST.

BS id 1/2/3 are the BS ID(s) which are audible from ST.

Signal Strength 1/2/3 are the signal strengths received by the ST from each BS.

CID is basic connection ID set to a fixed value for PDU(s) used for re-ranging. This is left as blank for first ranging request for an ST.

Bkoff contains the backoff time of the ST, in case the ST does not transmit immediately at the start of the ranging block.

3. After construction of the Initial Ranging Request PDU (IRR), the ST waits for the ranging contention slot. It transmits the IRR MAC PDU in the appropriate slot during the UL sub-frame. One or more

ST(s) may transmit a 'Ranging Request' message in a ranging block.

4. Upon receiving the IRR PDU, the system S extracts the ST ID and signal strength measurements from the IRR PDU. It determines the best BS that the ST should be associated with and forms the Initial Ranging Request Response (IRRe) PDU. The structure of the IRRe PDU is given below:

Initial Ranging Response

MAC Header: with type field = 0x05

BS ID: 1 byte

Basic CID: 2 bytes; Primary CID: 2 bytes

Time Advance (Tadv) = 4 bytes

Field Description:

BS ID is the id of BS to which the ST should associate (register).

Basic CID is a connection id used for subsequent periodic ranging.

Primary CID is a connection id used for exchanging management messages.

Tadv is the time by which the ST is required to advance its slot timing. This is calculated as shown in Figure 11.

5. This IRRe PDU is transmitted in DL sub-frame in the ranging response slot (size = 1 slot). The BS first indicates the ranging response slot in DL-MAP entry $\langle \text{ST ID}, \text{Slot ID} \rangle$, where ST ID is the ID of the ST for which the response is being sent. The BS then waits for the appropriate DL-slot and transmits the IRRe PDU.
6. On the client side, the ST continues to scan DL sub-frame and read beacons. In order to find the IRRe PDU, it processes the DL-MAP in all the beacons it receives to determine if there is a targeted IRRe PDU for it. If it finds its ST ID in any of the DL-MAP, it identifies the corresponding ranging response slot and waits to receive the IRRe PDU in that DL-slot.
7. In case it does not find such a DL-MAP within a specific timeout period ($T_{\text{max-Rg}}$), it waits for a random amount of time (C_w) and transmits a IRR PDU once again. Since multiple ST(s) may perform ranging simultaneously, it is possible that the IRR PDU(s) collide. Hence the value for C_w is chosen from a window of $C_w\text{-Min}$ and $C_w\text{-Max}$ using a standard backoff algorithm. A flag is set in the IRR PDU to indicate that it is a retransmitted or duplicate request.

The appropriate values of C_w , $T_{\text{max-Rg}}$, $C_w\text{-MIN}$, $C_w\text{-Max}$ and detailed specification of backoff algorithm are deferred to a later release

8. Ranging is complete when the ST is able to complete processing the IRRe PDU to determine the BS it should associate with, the timing advance value and record the primary and basic Connection ID. The primary CID is used for further exchange of management messages, while the basic CID is used for further periodic ranging exchanges.

Note: The ranging block(s) occur at the same time in the synchronized frames for all BS(s) of system S. All ST(s) will transmit their respective IRR packets using a contention resolution protocol. The IRR packet should contain the IDs of the BS beacons the ST is able to receive in decreasing order of signal strength. One or more BS will successfully receive the IRR packet transmitted by an ST. The BS then transmits the IRRe packet to the ST in the ranging response slot of the downlink sub-frame. The ST will continue to transmit the IRR packet till such time as the BS whose beacon is the strongest received by the ST sends an IRRe packet, or till a specified timer expires.

Each time a BS receives the IRR packet from an ST, it measures the time delay from the start of the ranging block as determined by the System's slot clock to the arrival of the IRR packet. This requires the PHY to provide information about the time of arrival of a received packet. This could be a specific signal provided by the PHY, or it could be a measurement inferred from some other signal (e.g. the start of transfer of received data by the PHY) provided by the PHY. This measurement will have associated with it a margin of error $\pm m$ bits (@11 Mbps)

After the System receives the IRR, the System determines which BS to associate the ST with. This will normally be the strongest BS which successfully received the IRR, though this is not a must. The System will then register the ST to the BS it selects for the SS. Using the IRRe, the System will inform the ST about the timing advance (in number of bit periods at 11 Mbps) it must employ for uplink transmissions. Every ST will advance its uplink sub-frame by the number of bits specified in its timing advance.

The guard time between uplink transmissions from different ST(s) must be at least $2m$ ($\pm m$ bits is margin for measurement error). This is best provided by ensuring that the last slot of an ST's transmission has at least $2m$ bits of silence at the end. Every ST can then start its transmission at the slot boundary, (after advancing the uplink sub-frame by the timing advance specified for it).

4.7.2 Authentication and Security

After Ranging, the ST authenticates itself to the S. The authentication process is required prior to registration. The process involves a authentication request from the ST, followed by multiple message exchange between the BS and ST depending on the authentication scheme chosen. The 802.1x

authentication and security mechanisms shall be used. The primary CID will be used for any such exchange. Encryption and other privacy mechanisms may be required for the transfer of various PDU(s).

Mechanisms are provided for:

1. Mutual authentication between BS and a ST.
2. Encryption of the MSDU(s) and/or MPDU(s) by the LCS before they are transmitted over the air and decryption by the peer LCS. All the security services are performed in LCS and no SAP needs to be exported to SSS.

There are three possible security schemes:

- a) MAC based authentication and no encryption.
- b) Globally Shared key based WPA authentication and encryption.
- c) Pairwise Shared key base TKIP authentication and encryption.

While in option b) and c) both BS and ST authenticate each-other by way of a shared secret, we recommend that in option a) only the ST needs to authenticate itself to BS and the BS does not have to authenticate itself. This is because in our deployment scenario, a ST cannot communicate with multiple BS due to interference. Hence a ST will either get service or it will not. The reward for spoofing a BS is not high. Hence a ST may transact with any BS that is willing to provide it service.

The options b) and c) also provide for appropriate encryption schemes. The authentication process results in the initialization of the appropriate seed values, if any, for the chosen security scheme. In the initial stages of deployment, privacy concerns are not important in general and hence no encryption needs to be used. Any privacy concerns can be taken care of at application layer.

The steps involved in authentication are as follows:

1. ST constructs an Authentication Request PDU. The PDU has the following structure:

MAC Header: with type field set to Authentication Request (chosen in 4.3.1) and

CID = Primary CID for that ST.

MAC Address = 6 bytes - MAC address of the sender

Authentication Algorithm Identification = 1 byte.

Authentication Transaction Sequence No = 1 byte.

Len = 1 byte.

2. The Authentication PDU can be sent in the contention slot(s) allocated in UL sub-frame. These slots are indicated in the Beacon's UL-MAP, by the ST ID value 0x0. Since Authentication typically follows

immediately after Initial Ranging, the scheduler may optionally allocate UL-slot(s) specifically to the ST, in the next frame(s). This is again indicated in the UL-MAP.

3. The MAC sub-layer of S receives the Authentication Request PDU and depending on the Authentication Algorithm, performs the authentication. This process may involve multiple message exchange with the ST. The structure of these messages is same as that described in step 1 above. The Authentication Transaction Sequence No is incremented at every step by both sides.

4. These Authentication Response PDUs are transmitted on the DL, in a manner similar to the Ranging Response (IRRe).

5. The ST processes the DL-MAP, identifies the corresponding authentication message slot and waits to receive the Authentication Response PDU in that DL-slot.

6. In case the ST does not receive the Authentication Response PDU within a specific timeout period, it performs a backoff in a manner similar to that defined for Ranging and retransmits the Authentication PDU with the duplicate flag set to 1. When the S receives an Authentication PDU with the duplicate flag set to 1, it simply processes the PDU again.

7. After successful authentication S associates the related authentication information with the primary CID. When new connections are created using this CID, they inherit the security parameters of the primary CID.

8. Authentication is complete when the ST receives a message from S with Authentication Result set to 'successful'.

9. In case of authentication failure, primary CID can be released and a warning logged.

Further detailed specifications of appropriate security sub-procedures are deferred to a later release.

4.7.3 Registration

Registrant of ST to S happens after ST is authenticated with S. Through this procedure, the ST informs the System S that it is entering the set of ST serviced by S. The link between S and ST is connection-oriented: one or more connections can be established for data exchange. The registration process is required prior to any data connection formation. The process involves a *registration request* from the ST, followed by a *registration response* from S. During this process, ST and S exchanges operational parameters and

capabilities. This process enables the ST to communicate packet protocol specification such as IP version and acquire IP address from S for set up of provisioned connections. The detailed steps involved in registration are as follows:

1. ST constructs Registration Request (RegR) PDU to S. The PDU has the following structure:

Registration Request

MAC Header: with type field = 0x06 and CID = Primary CID for that ST.

IP version = 1 byte.

ParamSet = 38 bytes (44 bytes (one slot) – 5 bytes Header – 1 byte IP version).

Field Description:

The value of CID is set to the primary CID as received in the Ranging Response (IRRe).

IP Version is the IP version supported by ST.

Paramset is a Type-Length-Value parameter which can be used for representing operational parameters of ST. The data is byte-stream serialized and represented as type-value pairs.

2. The RegR PDU can be sent in the contention slot(s) allocated in UL sub-frame. These slots are indicated in the Beacon's UL-MAP, by the ST ID value 0x1. Since Registration typically follows immediately after Initial Ranging, the scheduler may optionally allocate UL-slot(s) specifically to the ST, in the next frame(s). This is again indicated in the UL-MAP.
3. The MAC sub-layer of S receives the RegR PDU and a) checks for appropriate version, b) Generates IP address and c) installs resource for provisioned connection. Thereafter it constructs the Registration Response (RegRe) PDU with following structure:

Registration Response

MAC Header: with type field = 0x07

IP Version: 1 byte; IP Address: 4/ 6 bytes

ParamSet: (44 – above) bytes

Field Description:

IP Version is required for ST to interpret the IP address correctly.

Paramset is a Type-Length-Value parameter which can be used for representing operational parameters of S. The byte stream is serialized data, represented as type value pair. The result (success or failure) of connection provisioning is given in the byte stream. In case of success, the duration of registration validity may be provided here.

4. This RegRe PDU is transmitted on the DL, in a manner similar to the Ranging Response (IRRe).

5. The ST processes the DL-MAP identifies the corresponding registration response slot and waits to receive the RegRe PDU in that DL-slot.
6. In case it does not receive the RegRe PDU within a specific timeout period, it performs a backoff in a manner similar to that defined for Ranging and retransmits the RegR PDU with the duplicate flag set to 1. When the BS receives a RegR PDU with the duplicate flag set to 1, it checks if it had received any valid RegR from the same ST. If yes, it simply retransmits the corresponding RegRe. Otherwise the same ST may get multiple IP addresses.
7. Registration is complete when the ST is able to process the RegRe PDU, determine its IP address and the secondary CID assigned to it.

After completion of the registration process, a ST has an IP address, provisioned connections, active operation parameters and access to the network for future data communication. Now the ST may enter the connection request phase depending on connection demand from higher layers.

4.8 Connection Management sub-procedures

After registration, the ST can request for any number of further connections. A *connection request* from ST to S elicits a *connection response* from S to the ST. The number of connections may be restricted by S in an implementation specific fashion. The MAC is connection-oriented and data flow between BS and ST occurs as per the service flow type associated with that particular data flow. For example, a real-time VoIP data flow may be associated with one type of service flow while a best-effort TCP data flow may be associated with another type of service flow. The various types of service flows supported are described in section 4.2.

An active service flow is identified uniquely by a connection identifier (CID). There are two ways to create and change service flow with intended QoS parameters: a) create the connection with the desired QoS, using a Dynamic Service Addition message or b) create a generic connection (by specifying only the type of Bandwidth request service) and then use the CID to send Dynamic Service Change message to add specified QoS parameters to the connection. Thus, Connection Management consists of Service Addition, Change and Deletion sub-procedures.

4.8.1 Service Addition

This may also be termed as the Connection Creation phase. In this phase, the entity (BS or ST) wishing to create a data connection exchanges a management message which installs a CID at BS and informs the

destination about the nature of bandwidth request service to be used with the connection. The destination responds with a either acceptance or rejection of the request. The detailed steps involved in service addition are as follows:

1. A ST wishing to create a data connection sends a Dynamic Service Addition Request (DSA-Req) PDU to the BS. The DSA-Req PDU has the following structure:

Dynamic Service Addition Request

MAC Header: with type field = 0x08 and CID as the primary CID for that ST.

CID: 2 bytes

QosParamSet = (44 – above) bytes

Field Description:

CID is Primary Connection ID for that ST.

QosParamSet is a Type-Length-Value parameter which can be used for representing QoS parameters for the requested connection.

2. The BS processes the DSA-Req PDU, assigns a data CID and responds with a Dynamic Service Addition Response (DSA-Resp) PDU. The DSA-Resp PDU has the following structure:

Dynamic Service Addition Response

MAC Header: with type field = 0x09 and CID as primary CID for that ST.

CID: 2 bytes

Accpetd QosParamSet = (44 – above) bytes

Field Description:

CID is Data CID for admitted connection requested by ST.

Accpetd QosParamSet is a Type-Length-Value parameter which represents the allotted QOS parameters for the requested connection.

3. If the BS is initiating a connection creation then it generates a data CID and sends a DSA-Req PDU containing this CID in the CID field and the QoS parameter description. The ST responds to this message with a DSA-Resp PDU containing same CID and acceptance of QoS parameters.

After this process, a data CID has been created for data transmission. Also, the BS knows the service flow type of the connection. Hence it can appropriately schedule slots in uplink for ST to send data as well as resource (slot) request messages for that connection.

Note that in case of flows such as TCP, the higher layer may request for a separate data connection ID in

order to send the ACK(s). Otherwise the ACK(s) may be sent in the contention slots, leading to reduced TCP throughput.

4.8.2 Service Change

This may also be termed as the QoS Management phase. It is applicable to a new connection having a CID but not having any specified/allocated bandwidth resource. It is also applicable to an existing connection having some allocated resources but wanting a change in the allocation. The detailed steps involved in service change are as follows:

1. The ST sends a Dynamic Service Change Request (DSC-Req) PDU to the BS. The DSC-Req PDU has the following structure:

Dynamic Service Change Request

MAC Header: with type field = 0x10

CID : 2 bytes

QosParamSet = 36 (44- 8) bytes

Field Description:

CID is Data Connection ID of the active connection.

QosParamSet is a Type-Length-Value parameter which represents the change required in QoS parameters for the connection.

2. The BS may admit or reject the request, depending upon the admission control scheduler. It then responds with a Dynamic Service Change Response (DSC-Resp) PDU. The DSC-Resp PDU has the following structure:

Dynamic Service Change Response

MAC Header: with type field = 0x11

CID : 2 bytes

AccpetdQosParamSet = 36 (44- 8) bytes

Field Description:

CID is Data Connection ID of the active connection.

QosParamSet is a Type-Length-Value parameter which represents change accepted in QoS parameters for the connection.

4.8.3 Service Deletion

This may also be termed as the Connection Termination phase. In this phase, the entity (BS or ST) wishing

to terminate a data connection exchanges a management message to inform the peer entity. The steps involved in service deletion are as follows:

1. A ST wishing to terminate a data connection sends a Dynamic Service Deletion Request (DSD-Req) PDU to the BS. The DSD-Req PDU has the following structure:

Dynamic Service Deletion Request

MAC Header: with type field = 0x

CID : 2 bytes

ParamSet = 36 (44 – 8) bytes

Field Description:

CID is data Connection ID for the connection that is being terminated.

ParamSet may contain authentication information to guard against bogus deletion requests.

2. The BS processes the DSD-Req PDU, releases the resources assigned to that data connection ID and responds with a Dynamic Service Deletion Response (DSD-Resp) PDU. The DSD-Resp PDU has the following structure:

Dynamic Service Deletion Response

MAC Header: with type field = 0x

CID : 2 bytes

Status = 36 (44 – 8) bytes

Field Description:

CID is Data Connection ID for connection being terminated.

Status is a Type-Length-Value parameter which represents the success or error flag as a result of the deletion.

3. The BS may unilaterally decide to terminate a connection. In this case it simply sends a DSD-Req PDU to the ST. The ST responds to this message with a DSD-Resp PDU containing the status.

4.9 Data Transport sub-procedures

These include procedures for concatenation, fragmentation, ARQ and reassembly of data. They have been described in brief in sections 3.1.2 and 3.1.3 respectively.

Further detailed description of these sub-procedures is deferred to the next release.

4.10 Protocol Summary: State-Transition Diagrams

A high-level summary of the actions performed at the BS and ST is shown in Figure 22. A more detailed view of the ranging and registration process is shown in Figure 23.

Figure 24 represents the state transition diagram for the subscriber terminal (ST). Since ST may have limited battery power, it may go to power-saving mode. Hence, the initial state of ST is either ST_PowerOn, when the ST wakes up or the initial state is Idle, when ST is already 'on' and is waiting for some action to happen.

While in Idle state, ST can receive a notification that network layer has an IP packet to transmit. ST goes to ST_NetworkLayerHasPacketToSend state. It appends MAC header to the packet, does concatenation or fragmentation as required and adds the packet to the uplink traffic queue. ST transmits data and request while it is in ST_TransmitDataAndRequest state.

For the downlink sub-frame duration, ST PHY continuously listens to downlink channel to discover if there are any downlink packets intended for it. A message is sent by physical layer to notify the MAC for receiving a packet sent by the BS and the ST goes to ST_ReceivePacketFromPhy state. Downlink packets addressed to the ST are received and processed based on their type. Downlink and uplink control messages sent by BS on downlink channel are used for determining various control parameters for uplink and downlink channels. Downlink control message is decoded to determine start time of the frame, frame duration, etc. ST determines its uplink transmission time and duration of transmission in the current frame by decoding uplink control message. Downlink data packets are handed over to higher layer after removing MAC header.

Figure 25 represents the state transition diagram for the Base Station (BS). The initial state for BS is the Idle state. Periodically, BS transmits beacon by going to the BS_TransmitBeacons state.

When the PHY layer receives a packet from any ST, it notifies the MAC layer. BS goes to the BS_ReceivePacketFromPhy state. The packet is processed based on whether it is a data packet or a request packet. Uplink data packets handed to BS MAC layer are sent to higher layer after removing MAC header. Uplink bandwidth request packets are classified and placed in uplink grant queues.

On receiving a message from network layer for sending an IP packet, BS goes to the BS_NetworkLayerHasPacketToSend state and BS adds it to one of the downlink traffic queues. BS invokes

the multi-sector scheduling algorithm by going to the BS_PerformMultiSectorScheduling state and constructs DL MAP and UL MAP that are sent in the next beacon. The data is transmitted in the appropriate DL slot. BS generates periodic request grants by going to the BS_GeneratePeriodicGrant state.

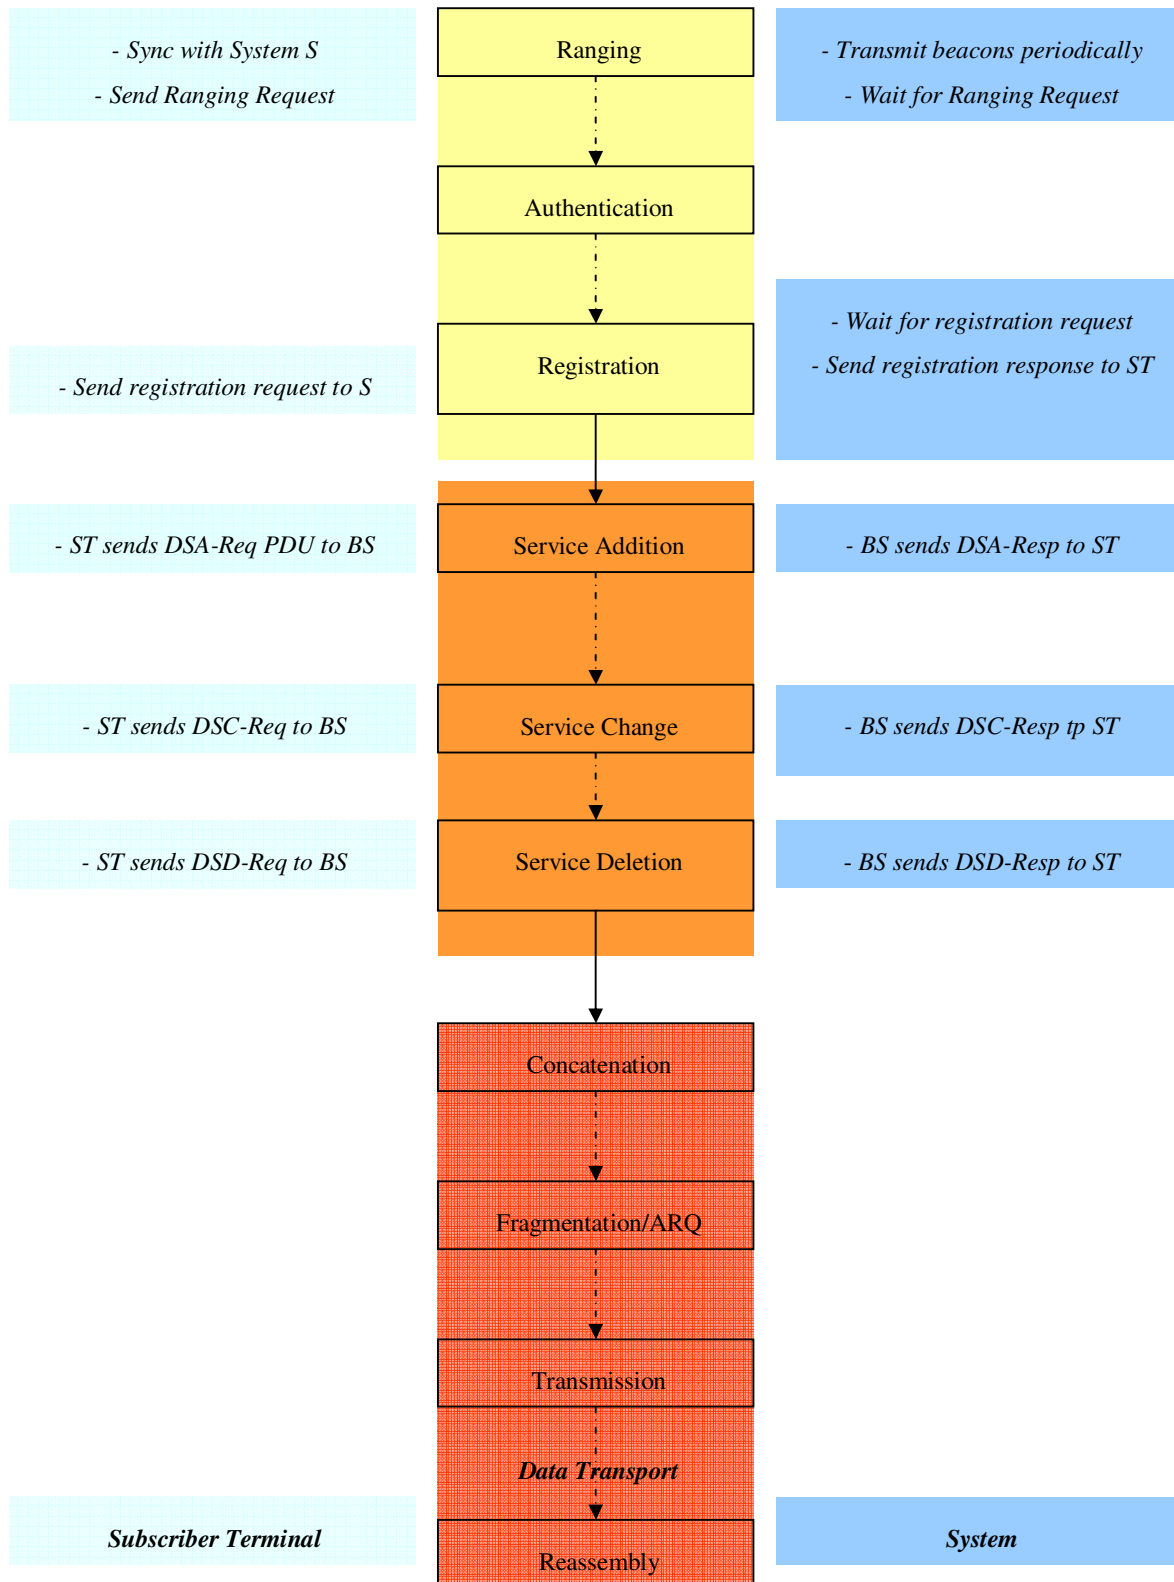


Figure 22: Summary of Actions

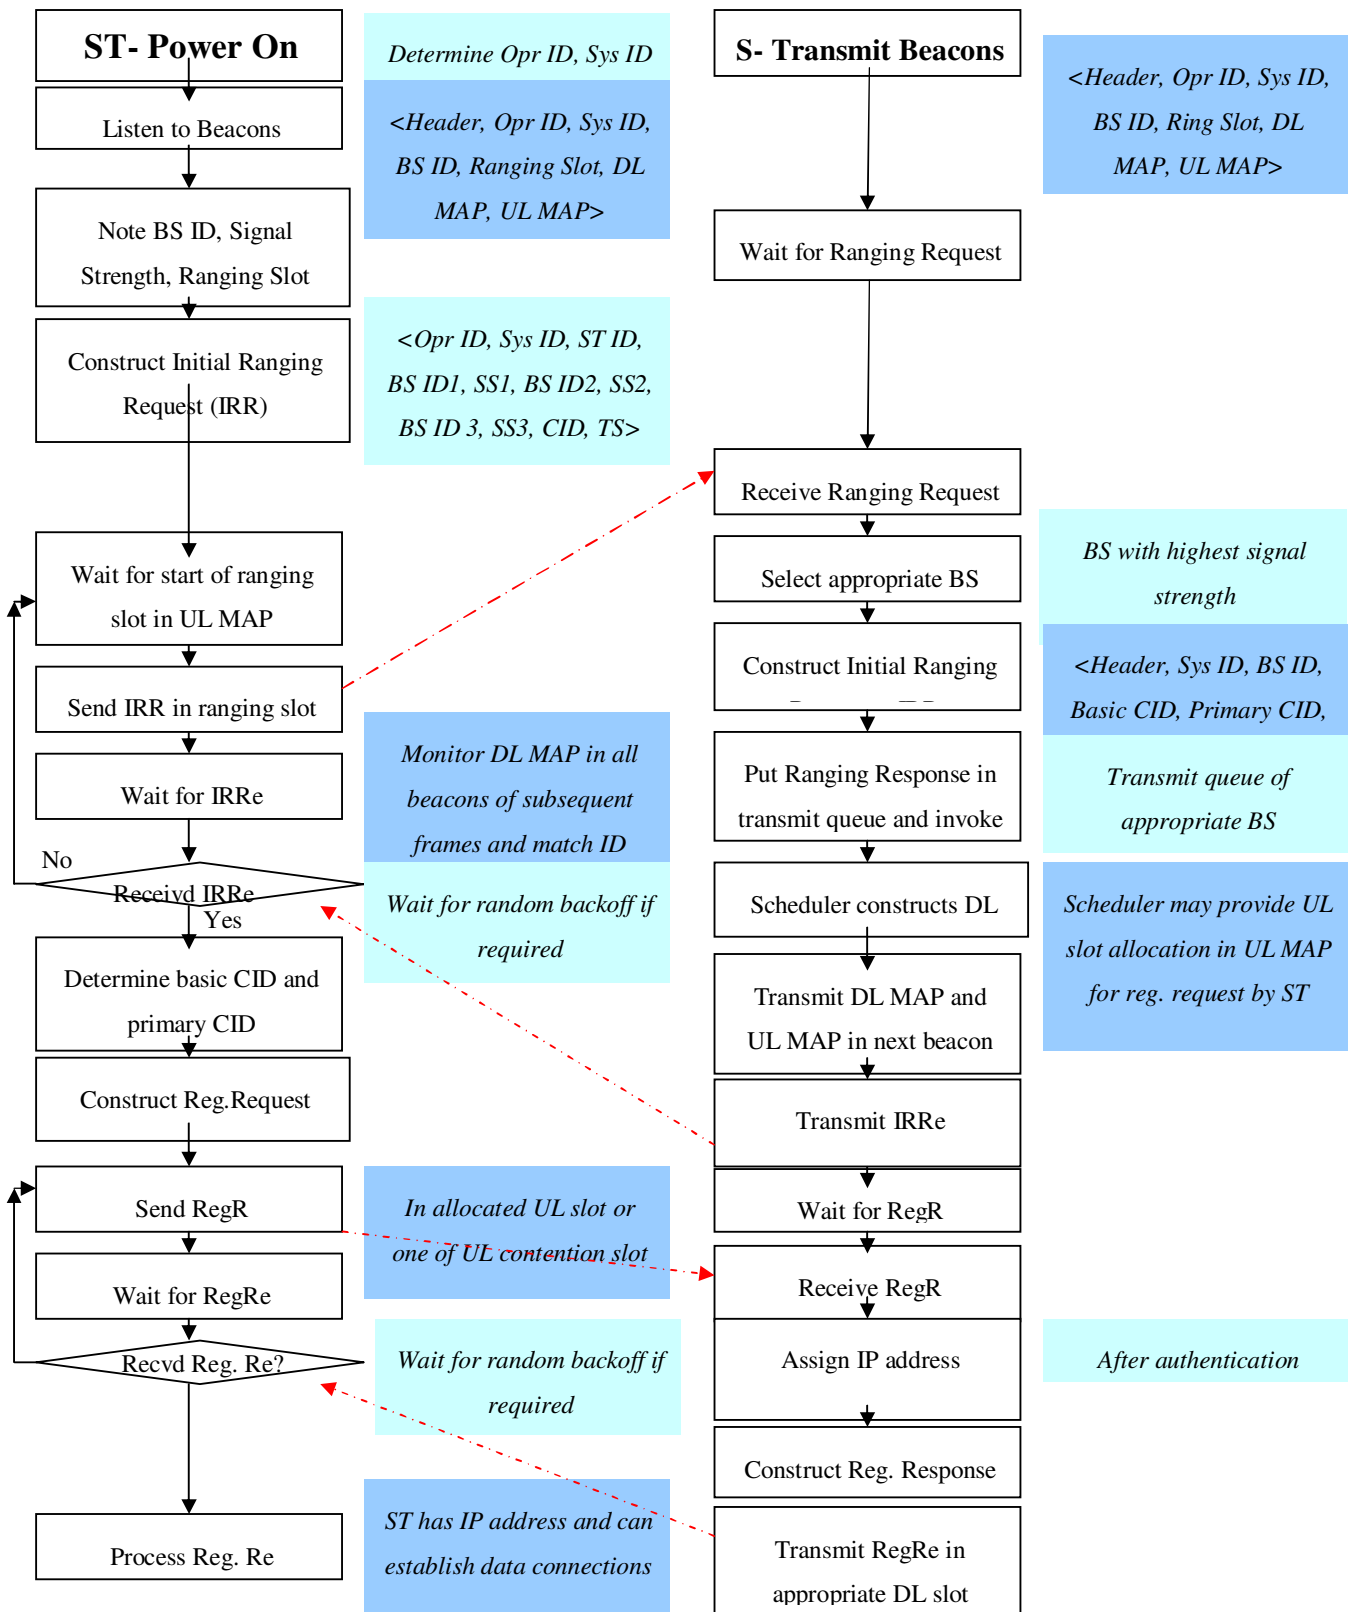


Figure 23: Ranging and Registration

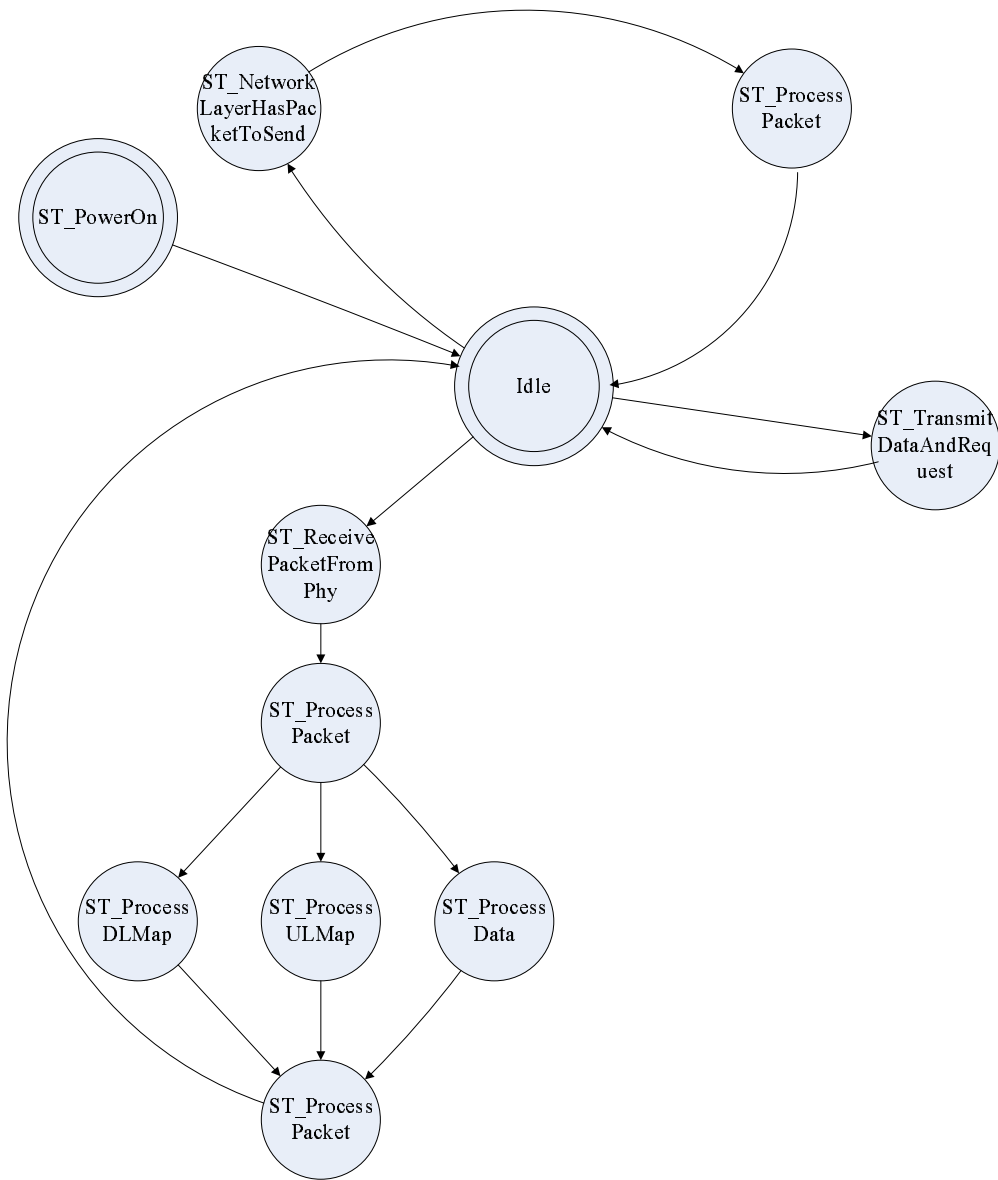


Figure 24: ST State-Transition Diagram

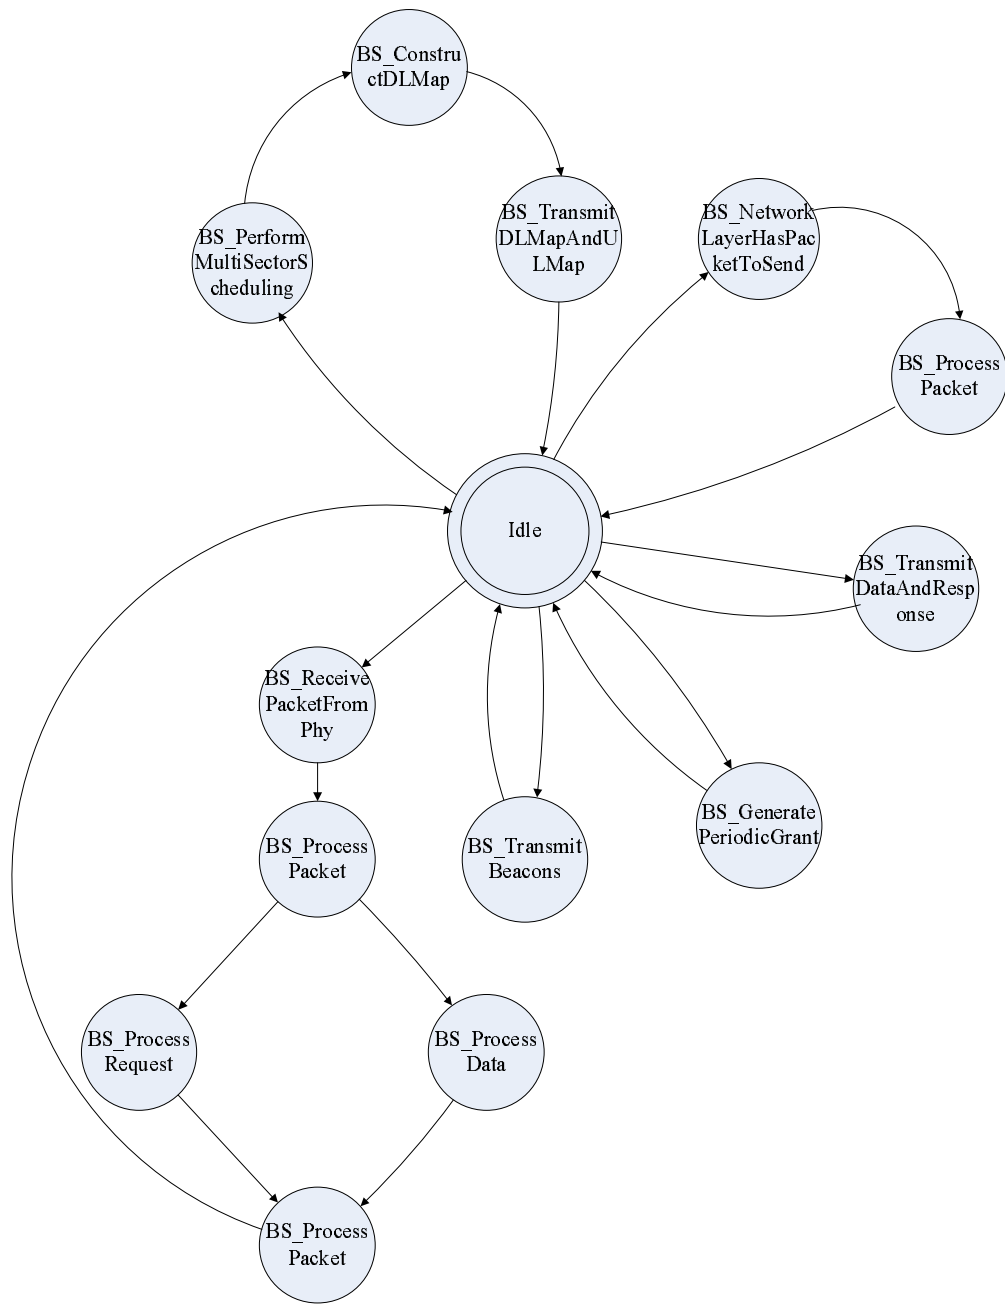


Figure 25: BS State-Transition Diagram

5 AUTHENTICATION AND PRIVACY

These mechanisms have been described in brief in section 4.7.2. Further detailed specification of the authentication and privacy mechanism is deferred to the next release.

6 MAC MANAGEMENT

This section describes the API(s) to be provided for configuring and management of a WiFiRe compliant device. This includes API(s) for setting the operator ID, system ID, WiFi channel, slot duration, frame duration, authentication parameters, encryption keys etc.

The detailed specification of the MAC management API(s) is deferred to the next release.

7 PHY SERVICE SPECIFICATION AND MANAGEMENT

The PHY is identical to the IEEE 802.11b Direct Sequence Spread Spectrum (DSSS). The PHY services used by the WiFiRe MAC are:

1. **PHY-DATA.request:** Issued by MAC to transfer data from the MAC sublayer to the local PHY entity.
2. **PHY-DATA.confirm:** Issued by PHY to confirm transfer data from the MAC sublayer to the local PHY entity.
3. **PHY-DATA.indication:** Issued by PHY to transfer data from the PHY sublayer to the local MAC entity.
4. **PHY-TXSTART.request:** Issued by MAC sublayer to local PHY entity to start the transmission of an MPDU.
5. **PHY-TXSTART.confirm:** Issued by PHY to confirm the start of transmission of an MPDU.
6. **PHY-TXEND.request:** Issued by MAC sublayer to local PHY entity to complete the transmission of the current MPDU.
7. **PHY-TXEND.confirm:** Issued by PHY to confirm completion of transmission of the current MPDU.
8. **PHY-RXSTART.indication:** Issued by PHY sublayer to local MAC entity to indicate receipt of a valid start frame delimiter.
9. **PHY-RXEND.indication:** Issued by PHY sublayer to local MAC entity to indicate that the MPDU currently being received is complete.

8 GLOSSARY OF TERMS

8.1 Abbreviations and Acronyms

ACK	acknowledgment
ARQ	automatic repeat request
BE	best effort
BR	bandwidth request
BS	base station
BSID	base station identification
BWA	broadband wireless access
C/I	carrier-to-interference ratio
C/N	carrier-to-noise ratio
CRC	cyclic redundancy code
CS	carrier sense
CSMA	carrier sense multiple access
DA	destination address
DCF	distributed co-ordination function
DL	downlink
DL-MAP	downlink slot allocation map
DL-TB	downlink transport block
DLL	data link layer
DSSS	direct sequence spread spectrum
ETSI	European Telecommunications Standards Institute
GPC	grant per connection
GPSF	grant per service flow type
GPST	grant per subscriber terminal
ID	identifier
IETF	Internet engineering task force
IFS	Inter frame space
IP	Internet protocol
LAN	local area network
LLC	logical link control
LoS	line of sight
MAC	medium access control layer
MAN	metropolitan area network

nrtPS	non-real-time polling service
PBR	piggyback request
PCF	point co-ordination function
PDU	protocol data unit
PHY	physical layer
PoP	point of presence
PS	physical slot
PSH	packing subheader
QoS	quality of service
RF	radio frequency
RSSI	received signal strength indication
rtPS	real-time polling service
Rx	reception
SAP	service access point
SDU	service data unit
SF	service flow
ST	subscriber terminal
TCP	transmission control protocol
TDD	time division duplex
TDM	time division multiplex
TDMA	time division multiple access
Tx	transmission
UDP	user datagram protocol
UE	user equipment
UGS	unsolicited grant service
UL	uplink
UL-MAP	uplink slot allocation map
UL-TB	uplink transport block
VoIP	voice over IP
WAN	wide area network
WDM	wireless distribution media
WDS	wireless distribution system
WiFi	wireless fidelity
WiMax	wireless microwave access

8.2 Definitions

1. **access control:** The mechanisms to prevent unauthorized usage of resources.
2. **authentication:** The service used to establish the identity of a subscriber terminal (ST) to the base station (BS) and vice versa.
3. **base station (BS):** The equipment used for providing wireless connectivity, management, and control of the subscriber terminals. It typically has a sectorized antenna.
4. **beacon:** A control packet transmitted by the BS at the start of every time frame.
5. **broadband:** Having bandwidth greater than 1 MHz and supporting data rates more than 256 Kbit/s.
6. **broadcast address:** A unique multicast address that specifies all ST(s).
7. **cell:** A set of co-located BS that provide wireless service to a given geographical area.
8. **channel:** An instance of medium use for the purpose of passing protocol data units (PDUs).
9. **concatenation:** The act of combining multiple medium access control (MAC) PDU(s) into a single time division multiplex (TDM) burst.
10. **connection:** A unidirectional mapping between BS and ST MAC layer peers for the purpose of transporting a service flow's traffic. All traffic is carried on a connection, even for service flows that implement connectionless protocols, such as internet protocol (IP). Connections are identified by a connection identifier (CID).
11. **connection identifier (CID):** Identifies a connection uniquely. It maps to a *service flow identifier* (SFID), which defines the quality of service (QoS) parameters of the service flow associated with that connection.
12. **deauthentication:** The service that voids an existing authentication relationship.
13. **downlink:** The direction from the base station (BS) to the subscriber terminal (ST).
14. **downlink map (DL-MAP):** Defines the mapping between ST identifier and slot start times for traffic sent on the downlink.
15. **dynamic service:** The set of messages and protocols that allow the base station and subscriber terminal to add, modify, or delete the characteristics of a service flow.
16. **frame:** A periodic, fixed duration, structured data transmission sequence. A frame contains both an uplink subframe and a downlink subframe.
17. **grant per connection (GPC):** A bandwidth allocation method in which grants are allocated to a specific connection within a ST. Note that bandwidth requests are always made for a connection.
18. **grant per service flow (GPSF):** A bandwidth allocation method in which grants are aggregated for all connections of the same service flow type, within a subscriber terminal.
19. **grant per subscriber terminal (GPST):** A bandwidth allocation method in which grants are

aggregated for all connections within a subscriber terminal and are allocated to the subscriber terminal as that aggregate.

- 20. MAC protocol data unit (MPDU):** The unit of data exchanged between two peer MAC entities using the services of the physical layer (PHY).
- 21. MAC service data unit (MSDU):** Information that is delivered as a unit between MAC service access points (SAPs).
- 22. multicast:** A medium access control (MAC) address that has the group bit set.
- 23. packing:** The act of combining multiple service data units (SDUs) from a higher layer into a single medium access control protocol data unit (PDU).
- 24. privacy:** The service used to prevent the content of messages from being read by other than the intended recipients.
- 25. ranging:** The service used by a subscriber terminal (ST) to notify the base station (BS) of its presence in the network.
- 26. registration:** The service used to establish mapping between ST and BS and enable ST to invoke the system services.
- 27. rural area:** An area about 15-20 Km radius, having a low density population.
- 28. service access point (SAP):** The point in a protocol stack where the services of a lower layer are available to its next higher layer.
- 29. slot:** A unit of time for allocating bandwidth.
- 30. station (ST):** Any device that contains a WiFiRe conformant medium access control (MAC) and physical layer (PHY) interface to the wireless medium (WM). It typically has a directional antenna.
- 31. time division duplex (TDD):** A duplex scheme where uplink and downlink transmissions occur at different times but may share the same frequency.
- 32. time division multiplex (TDM):** A scheme where the total number of available time slots are shared between multiple transmitters and receivers.
- 33. unicast:** A PDU that is addressed to a single recipient, not a broadcast or multicast.
- 34. uplink:** The direction from a subscriber terminal to the base station.
- 35. uplink map (UL-MAP):** Defines the mapping between ST identifier and slot start times for traffic on the uplink, for a scheduling interval.
- 36. wireless medium (WM):** The medium used to implement the transfer of protocol data units (PDUs) between peer physical layer (PHY) entities.

ANNEX A : Design Drivers

1 The Indian Rural Scenario

About 70% of India's population, or 750 million, live in its 600,000 villages. More than 85% of these villages are in the plains or on the Deccan plateau. The average village has 200-250 households, and occupies an area of 5 sq. km. Most of this is farmland, and it is typical to find all the houses in one or two clusters. Villages are thus spaced 2-3 km apart, and spread out in all directions from the market towns. The market centers are typically spaced 30-40 km apart. Each such centre serves a catchment of around 250-300 villages in a radius of about 20 km. As the population and the economy grow, several large villages are continually morphing into towns and market centres.

The telecommunication backbone network, mostly optical-fiber based, which passes through these towns and market centers, is new and of high quality. The state-owned telecom company has networked exchanges in all these towns and several large villages with optical fiber that is rarely more than 10-15 years old. The mobile revolution of the last four years has seen base stations sprouting in all these towns, with three or more operators, including the state-owned company. These base stations are also networked using mostly optical fiber laid in the last 5 years. There is a lot of dark fiber, and seemingly unlimited scope for bandwidth expansion.

The solid telecom backbone that knits the country together ends abruptly when it reaches the towns and larger villages. Beyond that, cellular coverage extends mobile telephone connectivity up to a radius of 5 km, and then telecommunications simply peters out. Cellular telephony will expand further as it becomes affordable to the rural populace. It is a highly sought after service, and the only reason for the service not spreading as rapidly in rural areas as in urban areas is the lack of purchasing power in the the rural areas. Fixed wireless telephones have been provided in tens of thousands of villages as a service obligation; however, the wireless technologies currently being deployed can barely support dial-up speeds as far as Internet access is concerned.

The rural per capita income is distinctly lower than the national average, and rural income distribution is also more skewed. About 70% of the rural households earn less than Rs 3000 per month, and only 4% have incomes in excess of Rs 25000 per month. Only the latter can be expected to even aspire to have a personal computer and Internet connection. For the rest, the key to Internet access is a public kiosk providing a basket of services. Provision of basic telecommunications as well as broadband Internet services is imperative, since ICT is known to be an enabler for wealth creation.

2 Affordability

When considering any technology for rural India, it is clear that the question of affordability must be addressed first. Given the income levels, one must work backwards to determine the cost of any economically sustainable solution. It is reasonable to expect an expenditure on telecommunication

services of only around Rs 60 per month on the average (2% of household income) from about 70% of the 200-250 households in a typical village. Thus, the revenue of a public kiosk can only be of the order of Rs 4500 per month (assuming two kiosks per village on the average). Apart from this, a few wealthy households in each village can afford private connections. Taking into account the cost of the personal computer, power back-up, peripherals, etc, it is estimated that a cost of at most Rs 15000 per broadband connection is sustainable for the kiosk. This includes the User Equipment, as well the per-subscriber cost of the Network Equipment connecting the user to the optical fiber PoP.

A typical wireless system for servicing such a rural area will have a BTS at the fiber PoP. A BTS can be expected to serve about 250-300 connections initially, going upto a 1000 connections as the service becomes stable and popular and the wealthy households decide to invest in a computer. Growth to full potential will take several years. Given the cost target mentioned above, it is found that a wireless technology becomes economically viable in the rural areas only when it has reached maturity and volumes worldwide are high enough to bring the cost down. New technologies at the early induction stage are too costly, particularly since the slow growth in the subscriber base keeps the per-subscriber cost of the BTS and associated equipment high.

3 Coverage, Towers and System Cost

We have already mentioned that we need to cover a radius of 15-25 km from the PoP using wireless technology. The system gain is a measure of the link budget available for overcoming propagation and penetration losses (through foliage and buildings) while still guaranteeing system performance. Mobile cellular telephone systems have a system gain typically of around 150-160 dB, and achieve indoor penetration within a radius of about 3-5 km. They do this with Base Station towers of 40 m height, which cost about Rs 5 lakhs each. If a roof-top antenna is mounted at the subscriber end at a height of 6m from the ground, coverage can be extend upto 15-20 km with this system gain. When the system gain is lower at around 135 dB, as with many low-power systems such as those based on the WiFi standard or the DECT standard [19], coverage is limited to around 10 km and antenna-height at the subscriber-end has to be at least 10m. This increases the cost of the installation by about Rs 1000.

In any case, we see that fixed terminals with roof-top antennas are a must if one is to obtain the required coverage from the fiber PoP. A broadband wireless system will need a system gain of around 140 - 150 dB at bit-rates in excess of 256 kbps, if it is to be easily deployable. This system gain may be difficult to provide for the higher bit-rates supported by the technology, and one may have to employ taller poles in order to minimize foliage loss.

There is an important relationship between coverage and the heights of the towers and poles, and indirectly their cost. The Base Station tower must usually be at least 40 m high for line-of-sight deployment, as trees have a height of 10-12m and one can expect a terrain variation of around 20-25m even in the plains over a 15-20 km radius. Taller Base Station towers will help, but the cost goes up exponentially with height. A shorter tower will mean that the subscriber-end will need a 20 m mast. At Rs 15000 or more, this is substantially costlier than a pole, even if the mast

is a guyed one and not self-standing. The cost of 250-300 such masts is very high compared to the additional cost of a 40 m tower vis--vis a 30 m one. With the 40m towers, simple poles can be deployed at the subscriber-end, and these need be only than 12m high.

In summary, one can conclude that for a cost-effective solution the system gain should be of the order of 145 dB, (at least for the reasonable bit-rates, if not the highest ones supported), a 40 m tower should be deployed at the fiber PoP, and roof-top antennas with 6-12m poles at the subscriber-end. The system gain can be lower at around 130 dB, provided repeaters are used to cover areas beyond 10 km radial distance, and assuming antenna poles that are 10-12m high are deployed in the villages. The cost per subscriber of the tower and pole (assuming a modest 300 subscribers per tower) is Rs 2500. This leaves about Rs 12500 per subscriber for the wireless system itself.

4 Definition of Broadband

The Telecom Regulatory Authority of India has defined broadband services as those provided with a minimum downstream (towards subscriber) data rate of 256 kbps. This data-rate must be available unshared to the user when he/she needs it. At this bit-rate, browsing is fast, video-conferencing can be supported, and applications such as telemedicine and distance education using multi-media are feasible. There is no doubt that a village kiosk could easily utilise a much higher bit-rate, and as technology evolves, this will become available too. However, it is important to note that even at 256 kbps, since kiosks can be expected to have a sustained rate not much lower, 300 kiosks will generate of the order of 75 Mbps traffic to evacuate over the air per Base Station. This is non-trivial today even with a spectrum allocation of 20 MHz.

The broadband wireless access system employed to provide Internet service to kiosks must also provide telephony using VoIP technology. Telephony earns far higher revenue per bit than any other service, and is an important service. The level of teletraffic is limited by the income levels of the populace. Assuming that most of the calls will be local, charged at around Rs 0.25 per minute, even if only one call is being made continuously from each kiosk during the busy hours (8 hours per day), this amounts to an expenditure of Rs 120 per day at each kiosk. This is a significant fraction of the earnings of the kiosk, and a significant fraction of the total communications expenditure of the village.

Thus, depending on the teledensity in the district, one can expect around 0.5-1 Erlang traffic per kiosk. This works out to a total of around 100-200 Erlangs traffic per BTS. Assuming one voice call needs about 2x16 kbps with VoIP technology, this traffic level requires 2x1.5 Mbps to 2x3.0 Mbps of capacity. If broadband services are not to be significantly affected, the system capacity must be several times this number. It is to be noted here that if the voice service either requires a higher bit-rate (say, 64 kbps) per call, or wastes system capacity due to MAC inefficiency when handling short but periodic VoIP packets, we will have significant degradation of other broadband services. Thus, an efficient VoIP capability is needed, with QoS guaranteed, that eats away from system capacity only as much as is unavoidably needed to support the voice traffic. Such a capability must be built into the wireless system by design. It is also important not to discourage use of the system for telephony since it is the major revenue earner as well as most popular service.

5 Broadband Wireless Technologies circa 2006

One of the pre-requisites for any technology for it to cost under Rs 12500 per connection is that it must be a mass-market solution. This will ensure that the cost of the electronics is driven down by volumes and competition to the lowest possible levels. As an example, both GSM and CDMA mobile telephone technologies can today easily meet the above cost target, except that they do not provide broadband access.

The third-generation evolution of cellular telephone technologies may, in due course, meet the cost target while offering higher bit-rate data services. However, they are in the early induction stage at present, and it is also not clear whether they will right away provide the required system capacity. However, the third-generation standards are constantly evolving, and the required system capacity is likely to be reached at some time. The only question is regarding when the required performance level will be reached and when the cost will drop to the required levels.

If we turn our attention next to some proprietary broadband technologies such as iBurst [7], and Flash-OFDM [8], or a standard technology such as WiMAX-d (IEEE 802.16d) [10], we find that volumes are low and costs high. Of these, WiMAX-d has a lower system gain. All of them will give a spectral efficiency of around 4 bps/Hz/cell (after taking spectrum re-use into account), and thus can potentially evacuate 80 Mbps with a 20 MHz allocation. High cost is the inhibitory factor.

It is likely that one or more OFDMA-based broadband technologies will become widely accepted standards in the near future. WiMAX-e (IEEE 802.16e) [10] is one such that is emerging rapidly. The standards emerging as the Long-Term Evolution (LTE) of the 3G standards are other candidates. These will certainly have a higher spectral efficiency, and more importantly, when they become popular and successful, they will become mass-market technologies, and the cost will be low. Going by the time-to-maturity of mass-market wireless technologies till date, none of these technologies are likely to provide an economically viable solution for India's rural requirements for several years yet.

6 Alternative Broadband Wireless Technologies in the Near Term

While wide-area broadband wireless technologies will be unavailable at the desired price-performance point for some time, local-area broadband technologies have become very inexpensive. A well-known example is WiFi (IEEE 802.11) technology. These technologies can provide 256 kbps or more to tens of subscribers simultaneously, but can normally do so only over a short distance, less than 50m in a built-up environment. Several groups have worked with the low-cost electronics of these technologies in new system designs that provide workable solutions for rural broadband connectivity.

One of the earliest and most widely deployed examples of such re-engineering is the corDECT Wireless Access System [9] developed in India. A next-generation broadband corDECT system has also been launched recently, capable of evacuating upto 70 Mbps per cell in 5 MHz bandwidth (supporting 144 full-duplex 256 kbps connections simultaneously). These systems are built around

the electronics of the European DECT standard, which was designed for local area telephony and data services. Proprietary extensions to the DECT standard have been added in a manner that the low-cost mass-market ICs can continue to be used. These increase the bit-rate by three times, while being backward compatible to the DECT standard.

The system gain in Broadband corDECT for 256 kbps service is 125-130 dB, depending on the antenna gain at the subscriber-end. This is sufficient for 10 km coverage under line-of-sight conditions (40 m tower for BS and 10-12 m pole at subscriber side). A repeater is used for extending the coverage to 25 km. The system meets the price-performance requirement, but with the additional encumbrance of taller poles and one level of repeaters.

The WiFiRe standard proposed by CEWiT is an alternative near-term solution, with many similarities. It, too, is a re-engineered system based on low-cost low-power mass-market technology. Cost structures are similar, and deployment issues too are alike. There is one key aspect in which WiFiRe differs from Broadband corDECT. The spectrum used for WiFiRe is unlicensed without fees, whereas the spectrum used by Broadband corDECT is licensed with a fee. The spectral efficiency of the WiFiRe system is poorer, and the cell capacity per MHz of bandwidth is lower. However, this is offset by the fact that the spectrum used by it is in the unlicensed WiFi band of 2.4-2.485 MHz. This unlicensed use is subject to certain conditions, and some modifications to these conditions will be needed to support WiFiRe in rural areas (see section on Conditional Licensing below). WiFiRe technology is best suited for local niche operators who can manage well the conditionalities associated with unlicensed use. It does not afford the blanket protection from interference that a system operating with licensed spectrum enjoys.

7 Motivation for WiFiRe

In recent years, there have been some sustained efforts to build a rural broadband technology using the low-cost, mass-market WiFi chipset. WiFi bit rates go all the way up to 54 Mbps. Various experiments with off-the-shelf equipment have demonstrated the feasibility of using WiFi for long-distance rural point-to-point links [12]. One can calculate that the link margin for this standard is quite adequate for line-of-sight outdoor communication in flat terrain for about 15 kms of range. The system gain is about 132 dB for 11 Mbps service, and as in corDECT, one requires a 40 m tower at the fiber PoP and 10-12 m poles at the subscriber-end.

The attraction of WiFi technology is the de-licensing of spectrum for it in many countries, including India. In rural areas, where the spectrum is hardly used, WiFi is an attractive option. The issues related to spectrum de-licensing for WiFiRe are discussed separately in the next section. Before that, we turn our attention to the suitability of the WiFi standard as it exists for use over a wide rural area. We have already seen that the limitations of the Physical Layer of WiFi can be demonstrably overcome. We turn our focus now to the MAC in WiFi.

The basic principle in the design of MAC in Wi-Fi is fairness and equal allocation to all sources of demand for transmission. This leads to the DCF mode which operates as a CSMA/CA with random backoff upon sensing competing source of tx. On the other hand there is also a PCF mode, which assumes mediation by access points. This gives rise to the possibility of enterprise-owned

and managed networks with potential for enhanced features like security and quality of service guarantees.

The CSMA/CA DCF MAC has been analysed and turns out to be inefficient for a distribution service that needs to maximize capacity for subscribers and maintain quality of service [13]. The delays across a link are not bounded and packet drops shoot up rapidly in such a system while approaching throughputs of the order of 60% of rated link bandwidth. The PCF MAC will perform better than the DCF MAC. However, both the MACs in the WiFi standard become very inefficient when the spectrum is re-used in multiple sectors of a BTS site.

Fundamentally, in a TDD system, wherein uplink and downlink transmissions take place in the same band in a time-multiplexed manner, the down-link (and similarly uplink) transmissions of all the sectors at a BTS site must be synchronized. Otherwise the receivers in one sector will be saturated by the emissions in another. This can be avoided only by physical isolation of the antennas, which is very expensive if all the antennas must also be at a minimum height of 40 m. Further, this synchronization must be achieved with minimal wastage of system capacity due to the turnaround from uplink to downlink and vice versa, as well as due to varying traffic characteristics (packet sizes, packet arrival rates) in different sectors at different times.

It is thus clear that a new MAC is needed which is designed to maximize the efficiency in a wide-area rural deployment supporting both voice and data services with modest use of spectrum (see next section for the need for limiting the use of spectrum). Fortunately, most Wi-Fi chipsets are designed so that the Physical and MAC layers are separate. Thus one can change the MAC in ways that enable high- efficiency outdoor systems that can be used for rural internet service provisioning or voice applications, while retaining the same PHY. Thus without significantly changing radio costs, one can arrive at entirely different network level properties by changing the MAC, sectorization and antenna design choices and tower/site planning. Taking a cue for this approach, we design a new wireless system, WiFiRe, which shares the same PHY as WiFi, but with a new MAC. The principle of Access Points, or special nodes which control the channel and allocate bandwidth to individual nodes, and tight synchronization based on the time-slotting principle used in cellular voice systems such as GSM or upcoming data systems like WiMax, can be combined to guarantee efficiency and quality of service.

8 Conditional Licensing of Spectrum

The spectrum allotted for WiFi, in the 2.4-2.485 MHz band, can be employed by anyone for indoor or outdoor emissions, without a prior license provided certain emission limits are met [www.dotindia.com/wpc]. The 5 GHz band, also universally allotted to WiFi, can be used in India only for indoor emissions. In the 2.4 Ghz band, the maximum emitted power can be 1W in a 10 MHz (or higher) bandwidth, and the maximum EIRP can be 4W. The outdoor antenna can be no higher than 5 m above the rooftop. For antenna height higher than the permissible level, special permission has to be obtained. Further, if the emissions interfere with any licensed user of spectrum in the vicinity, the unlicensed user may have to discontinue operations.

It is clear that some modifications of the rules are needed for WiFiRe. A higher EIRP will need

to be permitted in rural areas, and further, antenna deployment at 40 m must be permitted at the PoP, and possibly for repeaters (in due course). Antenna deployments at 10-15 m will have to be permitted at the villages.

The relaxations may be restricted to WiFiRe-compliant technology. It may be given only for one specified carrier per operator, and a maximum of two operators may be permitted in an area. The BTS and repeaters of the second operator (in chronological order of deployment) may be restricted to be at least one kilometer from those of the first operator in an area. This will prevent mutual adjacent-channel interference, as well as permit maximum use of the two conditionally licensed carriers by others in the vicinity of the BTSs. If an unlicensed WiFi user in the vicinity of the BTS or village kiosk/private subscriber interferes with the WiFiRe system, the unlicensed user will have to switch over to a non-interfering carrier in the same band or in the 5 GHz band. This last condition is not very restrictive, as only around 15 MHz of the available 85 MHz in the 2.4 GHz band is blocked in the vicinity of any one BTS or village kiosk/subscriber. Further, if the unlicensed user is an indoor user, the area where there is noticeable interference to/from the WifiRe system is likely to be fairly small.

ANNEX B: Capacity Analysis and Optimisation

1 Spatial Reuse Model

Maximising the cardinality of independent sets used in a schedule need not necessarily increase the throughput, since as the cardinality of the set increases, the prevailing SINR drops, thereby resulting in an increase in the probability of error, decreasing the throughput. Hence it is necessary to limit the cardinality of the independent set used so as to satisfy the SINR requirements. i.e., there is a limit to the number of simultaneous transmissions possible.

In this section the problem of finding the maximum number of simultaneous transmissions possible in different sectors in the uplink and the downlink is being considered. There is no power control in the downlink. The BTS transmits to all the STs at the same power. There is static power control in the uplink. Each ST transmits to the BTS at a fixed power, such that the average power received from different STs at the BTS is the same. The STs near the BTS transmit at a lower power and the ones farther away transmit at a higher power.

A typical antenna pattern used in the deployment is as shown in Figure 26. Based on the antenna pattern, one can divide the region into an *association region*, a *taboo region* and a *limited interference region* with respect to each BTS.

The radial zone over which the directional gain of the antenna is above -3dB is called the association region. In the analysis, the directional gain is assumed to be constant over this region. Any ST which falls in this region of a BTS antenna j is associated to the BTS j .

The region on either side of the association region where the directional gain is between -3dB and -15dB is called the taboo region. Any ST in this region of BTS j causes significant interference to the transmissions occurring in Sector j . When a transmission is occurring in Sector j , no transmission is allowed in this region.

In the limited interference region the directional gain of the BTS antenna is below -15dB. A single transmission in this region of BTS j may not cause sufficient interference to the transmission in Sector j . But a number of such transmissions may add up causing the SINR of a transmission in Sector j to fall below the threshold required for error free transmission. This is taken care of by limiting the total number of simultaneous transmissions in the system as explained in Sections 1.1 and 1.2.

As an example, for the antenna pattern shown in Figure 26, the association region is a 60° sector centered at the 0° mark, the taboo region is 30° on either side of this association region, and the limited interference region covers the remaining 240° .

1.1 Uplink

In the uplink, there is static power control. All STs transmit at a power such that the power received at the BTS is P times noise power. Let the maximum power that can be transmitted by an ST be

RADIATION PATTERN - HORIZONTAL BEAMWIDTH




Figure 26: Radiation pattern for a typical antenna that could be used in the deployment.

P_t times noise power. Let R_0 be the distance such that when P_t is transmitted by an ST at distance R_0 , the average power received at the BTS is P_0 times noise power, where P_0 is the minimum SNR required to decode a frame with a desired probability of error. Also, let R be such that when P_t is transmitted from an ST at distance R , the power received at the BTS is P times noise power, i.e.,

$$\frac{P}{P_0} = \left(\frac{R}{R_0} \right)^{-\eta}$$

In the presence of interferers, the power required at the receiver will be greater than P_0 times noise. Let P be the power required, so that the receiver decodes the frame with a desired probability of error, in the presence of interferers. The directional gain of the BTS antenna is -15dB in the other non taboo directions. Hence, the interference power from a transmission in any other sector would be $10^{-\frac{3}{2}}P$. If there are $n_0 - 1$ simultaneous transmissions, the path loss factor being η , the signal

to interference ratio at the BTS receiver is

$$\begin{aligned}\Psi_{rcv} &= \frac{P}{1 + \sum_{i=1}^{n_0-1} 10^{-\frac{3}{2}} P} \\ &= \frac{P_0 \left(\frac{R}{R_0}\right)^{-\eta}}{1 + (n_0 - 1) 10^{-\frac{3}{2}} P_0 \left(\frac{R}{R_0}\right)^{-\eta}}\end{aligned}$$

For decoding a frame with less than a given probability of error, we need an SINR of P_0 at the receiver. So, R should be such that

$$\begin{aligned}\Psi_{rcv} &\geq P_0 \\ \frac{P_0 \left(\frac{R}{R_0}\right)^{-\eta}}{1 + (n_0 - 1) 10^{-\frac{3}{2}} P_0 \left(\frac{R}{R_0}\right)^{-\eta}} &\geq P_0 \\ n_0 &\leq 1 + \frac{\left(\frac{R}{R_0}\right)^{-\eta} - 1}{10^{-\frac{3}{2}} P_0 \left(\frac{R}{R_0}\right)^{-\eta}}\end{aligned}$$

To provide a margin for fading, consider a reduced range R' such that

$$10 \log \left(\frac{R'}{R} \right)^{-\eta} \geq 2.3\sigma$$

where σ is the fade variance. In this case, 99% of the STs in a circle of radius R' around the BTS can have their transmit power set so that the average power P is received at the BTS in the uplink.

Evidently, n_0 can be increased by reducing R' . But then, spatial reuse increases at the expense of coverage. This tradeoff can be captured by the spatial capacity measure $C = nR'^2$, which has units slots.km² (or packets.km²)

The variation of the maximum number of transmissions, n_0 and system capacity, C , with coverage is as shown in Figure 27. One can see that for each η , there is an optimal n_0 and R' such that C is maximum. The coverage for which capacity is maximum can be obtained by equating the derivative of C , with respect to (R'/R_0) to be zero. Take $r' = \frac{R'}{R_0}$ and set

$$\frac{dC}{dr'} = 0$$

Then we get the optimum value of r' and n_0 as

$$\begin{aligned}r' &= \left(10^{-\frac{2.3\sigma}{10}} \frac{1 + 10^{-\frac{3}{2}} P_0}{1 + \frac{\eta}{2}} \right)^{\frac{1}{\eta}} \\ n_0 &= \frac{(1 + 10^{-\frac{3}{2}} P_0)\eta}{10^{-\frac{3}{2}} P_0 (\eta + 2)}\end{aligned}$$

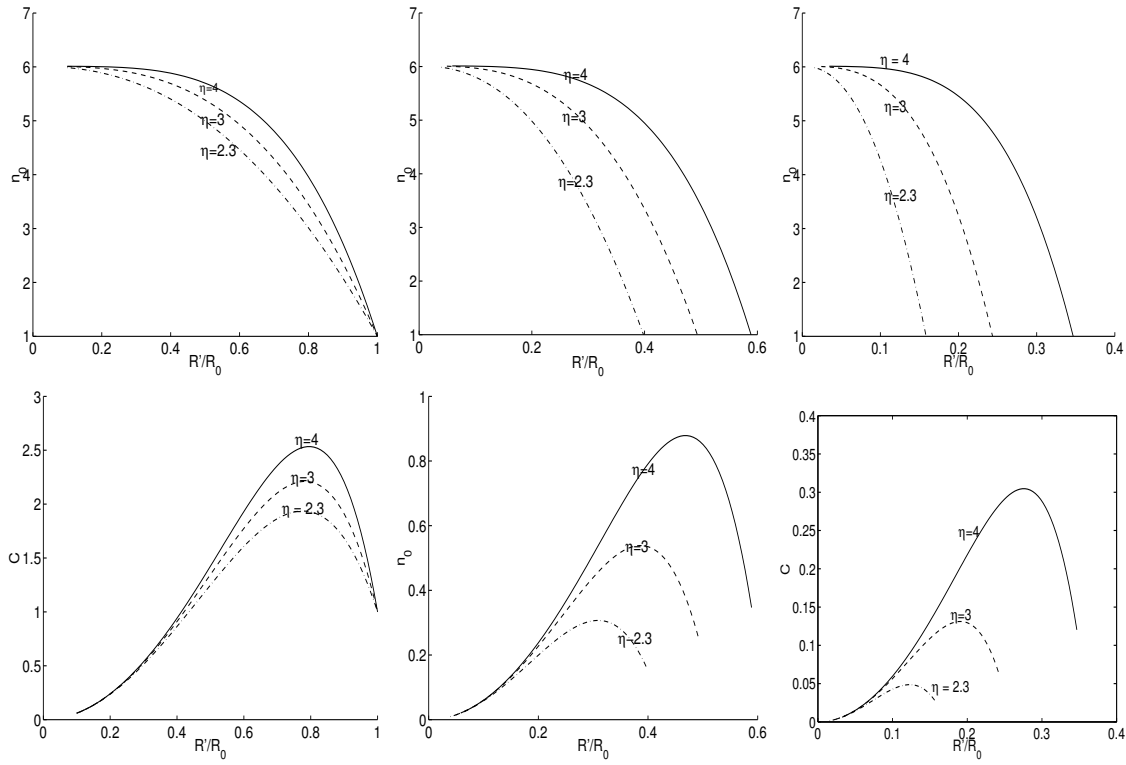


Figure 27: Variation of the number of simultaneous transmissions possible (n_0) and system capacity(C) with coverage relative to a reference distance R_0 for $\eta = 2.3, 3, 4$ and $\sigma = 0, 4, 8$. Plots for $\sigma = 0, 4, 8$ are shown left to right.

σ	0	4	8
η			
2.3	0.77	0.31	0.12
3	0.78	0.39	0.20
4	0.80	0.47	0.28

η	2.3	3	4
n_0	3	3	4

Table 1: The optimum values of C and n_0 for different values of η and σ .

Table 1 gives the optimum coverage C and maximum number of simultaneous transmissions possible for different values of η and σ . P_0 is taken to be $8dB$. Directional gain of the antenna is taken to be 1 in the associated sector and $-15dB$ in non-taboo directions. For path loss factor $\eta = 4$, the number of simultaneous transmissions is seen to be 4. For a given value of η , the maximum number of simultaneous transmissions is found to be independent of the fade variance σ .

1.2 Downlink

In the downlink, the transmit power is kept constant. The BTS antennas transmit at a power P_t times noise. Let R_0 be the distance at which the average power received is P_0 times noise. R be

the distance such that the average power received is P . Then,

$$\begin{aligned} P_0 &= P_t \left(\frac{R_0}{d_0} \right)^{-\eta} \\ P &= P_t \left(\frac{R}{d_0} \right)^{-\eta} \\ \frac{P}{P_0} &= \left(\frac{R}{R_0} \right)^{-\eta} \end{aligned}$$

Allowing $n_0 - 1$ interferers,

$$\begin{aligned} \Psi_{rcv} &= \frac{P}{1 + (n_0 - 1)10^{-\frac{3}{2}}P} \\ &= \frac{P_0 \left(\frac{R}{R_0} \right)^{-\eta}}{1 + (n_0 - 1)10^{-\frac{3}{2}}P_0 \left(\frac{R}{R_0} \right)^{-\eta}} \end{aligned}$$

which is the same as in uplink. So, the optimum number of transmissions and optimum coverage in uplink and downlink are the same. The plots and tables for uplink apply for downlink also.

1.3 Number of Sectors

Once the maximum number of simultaneous transmissions possible, n_0 is obtained, one gets some idea about the number of sectors required in the system. In an n_0 sector system, when a transmission occur in the taboo region between Sector j and Sector $j + 1$, no more transmissions can occur in Sectors j and $j + 1$. So, the number of simultaneous transmissions can be at most $n_0 - 1$, one in Sector j and $j + 1$ and at most one each in each of the other sectors. Thus the maximum system capacity cannot be attained with $n_0 - 1$ sectors. With $n_0 + 1$ sectors, one can choose maximal independent sets such that the sets are of cardinality n_0 . So, at least $n_0 + 1$ sectors are needed in the system. From the spatial reuse model it can be seen that there can be up to 4 simultaneous transmissions in the system, for path loss $\eta = 4$. So, the system should have at least 5 sectors.

2 Characterising the Average Rate region

There are m STs. Suppose a scheduling policy assigns $k_j(t)$ slots, out of t slots, to ST j , such that $\lim_{t \rightarrow \infty} \frac{k_j(t)}{t}$ exists and is denoted by r_j . Let $\mathbf{r} = (r_1, r_2, \dots, r_m)$ be the rate vector so obtained. Denote by $\mathcal{R}(n)$ the set of achievable rates when the maximum number of simultaneous transmission permitted is n . Notice that for $n_1 > n_2$, $\mathcal{R}_{n_1} \supset \mathcal{R}_{n_2}$. This is evident because any sequence of scheduled slots with $n = n_2$ is also schedulable with $n = n_1$. In the previous section, we have determined the maximum value of n , i.e., n_0 . Denote $\mathcal{R}_0 = \mathcal{R}(n_0)$. A scheduling policy will achieve an $\mathbf{r} \in \mathcal{R}_0$. In this section, we provide some understanding of \mathcal{R}_0 via bounds.

2.1 An Upper Bound on Capacity

Suppose each ST has to be assigned the same rate r . In this subsection an upperbound on r is determined. In general, the rate vector $(r, r, \dots, r) \notin \mathcal{R}_0$. The upper bound is obtained via simple linear inequalities. Consider the case $n \geq 3$. Suppose one wishes to assign an equal number of slots k to each ST in the uplink. There are N_U uplink slots in a frame. Consider Sector j , which contains m_j STs. Thus $k \cdot m_j$ slots need to be allocated to uplink transmission in Sector j . When STs in the interference region j^- or j^+ transmit, then no ST in Sector j can transmit. Suppose $k_{j\pm}$ slots are occupied by such interference transmission. Now it is clear that

$$k \cdot m_j + k_{j\pm} = N_U$$

because whenever there is no transmission from the interference region for sector j there can be a transmission from sector j . Let m_{j-} and m_{j+} denote the number of STs in the interference regions adjacent to Sector j . Since the nodes in j^- and j^+ can transmit together, we observe that

$$k_{j\pm} \geq \max(k \cdot m_{j-}, k \cdot m_{j+})$$

with equality if transmission in j^- and j^+ overlap wherever possible. Hence one can conclude that for any feasible scheduler that assigns k slots to each ST in the uplink

$$k \cdot m_j + \max(k \cdot m_{j-}, k \cdot m_{j+}) \leq N_U$$

For large frame time N , divide the above inequality by N and denote the rate of allocation of slots by r . Thus if out of t slots, each ST is allocated k slots, then $r = \lim_{t \rightarrow \infty} \frac{k}{t} \leq 1$

$$r \cdot m_j + r \cdot \max(m_{j-}, m_{j+}) \leq \phi_u$$

where ϕ_u is the fraction of frame time allocated to the uplink or

$$r \leq \frac{\phi_u}{m_j + \max(m_{j-}, m_{j+})}$$

This is true for each j . So,

$$r \leq \frac{\phi_u}{\max_{1 \leq j \leq n} (m_j + \max(m_{j-}, m_{j+}))}$$

For the case $n = 2$ for $j \in \{1, 2\}$ denote the interfering nodes in the other sector by m_j . One easily sees that

$$r \leq \frac{\phi_u}{\max(m_1 + m'_1, m_2 + m'_2)}$$

2.2 An Inner Bound for the Rate Region

In this section a rate set \mathcal{R}_L is obtained such that $\mathcal{R}_L \subset \mathcal{R}$, i.e., \mathcal{R}_L is an inner bound to the achievable rate set.

The following development needs some graph definitions.

Reuse constraint graph: Vertices represent links. In any slot all links are viewed as uplinks or all are downlinks. Two vertices in the graph are connected, if a transmission in one link can cause interference to a transmission in the other link. The reuse constraint graph is represented as $(\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the set of vertices and \mathcal{E} is the set of edges.

Clique: A fully connected subgraph of the reuse constraint graph. A transmission occurring from an ST in a clique can interfere with all other STs in the clique. At most one transmission can occur in a clique at a time.

Maximal clique: A maximal clique is a clique which is not a proper subgraph of another clique.

Clique incidence matrix: Let κ be the number of maximal cliques in $(\mathcal{V}, \mathcal{E})$. Consider the $\kappa \times m$ matrix \mathcal{Q} with

$$\mathcal{Q}_{i,j} = \begin{cases} 1 & \text{if link } j \text{ is in clique } i \\ 0 & \text{o.w.} \end{cases}$$

By the definition of \mathbf{r} and \mathcal{Q} , a necessary condition for \mathbf{r} to be feasible is (denoting by $\mathbf{1}$, the column vector of all 1s.)

$$\mathcal{Q} \cdot \mathbf{r} \leq \mathbf{1}$$

since at most one link from a clique can be activated. In general, $\mathcal{Q} \cdot \mathbf{r} \leq \mathbf{1}$ is not sufficient to guarantee the feasibility of \mathbf{r} . It is sufficient if the graph is linear. A linear graph is one in which links in each clique are contiguous. A linear clique will have a clique incidence matrix of the form

$$\mathcal{Q} = \begin{bmatrix} 1 & 1 & 1 & 1 & \dots & \dots & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & \dots \\ \vdots & & & & & & & & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 1 \end{bmatrix}$$

The reuse constraint graph in the multisector scheduling problem being considered has a ring structure. $\mathcal{Q} \cdot \mathbf{r} \leq \mathbf{1}$ gives an upper bound on the rate vector. The reuse constraint graph is linear except for the wrapping around at the end. If the nodes in one sector are deleted, the graph becomes linear. Let \mathbf{m}_i be the set of STs in Sector i . There is a feasible \mathbf{r}_i such that $\mathcal{Q} \cdot \mathbf{r}_i \leq \mathbf{1}$ and all STs in \mathbf{m}_i are given rate 0. Linear combination of feasible vectors is also feasible. Thus, defining

$$\mathcal{R}_L := \left\{ \mathbf{x} : \mathbf{x} = \sum_{i=1}^m \alpha_i \mathbf{r}_i; \quad \mathcal{Q} \cdot [\mathbf{r}_1 \mathbf{r}_2 \dots \mathbf{r}_m] \leq \mathbf{1}; \quad \sum_{i=1}^m \alpha_i = 1; \quad \mathbf{r}_1(\mathbf{m}_1) = \dots = \mathbf{r}_m(\mathbf{m}_m) = \mathbf{0} \right\}$$

We see that $\mathcal{R}_L \in \mathcal{R}$

2.3 Optimum Angular positioning of the Antennas

As can be seen from the previous section, feasible rates set, \mathcal{R}_0 , of the system depends on the spatial distribution of the STs around the BTS. Thus the \mathcal{R}_0 varies as the sector orientation is changed. A system where the antennas are oriented in such a way that most STs fall in the association region of BTSs rather than in the taboo region will have more capacity than one in which more STs are in the taboo regions.

One sector boundary is viewed as a reference. Let $\mathcal{R}_0(\theta)$ denote the feasible rate set, when this boundary is at an angle θ with respect to a reference direction. Then, for each $0 \leq \theta \leq \frac{360^\circ}{n}$, we have $\mathcal{R}_0(\theta)$, where n is the number of sectors. Since $\mathcal{R}_0(\theta)$ is not known, the inner bound $\mathcal{R}_L(\theta)$ is used in the following analysis. If each vector \mathbf{r} is assigned a utility function $U(\mathbf{r})$, then one could seek to solve the problem

$$\max_{0 \leq \theta \leq \frac{360^\circ}{n}} \max_{\mathbf{r} \in \mathcal{R}_L(\theta)} U(\mathbf{r})$$

and then position the antenna at this value of θ .

The optimization can be done so as to maximise the average rate allocated to each ST, with the constraint that each ST gets the same average rate. The bound evaluated with average rate to each ST, for antenna positions differing by 5° is given below. $ub(i)$ gives the upperbound on capacity of the system with antenna placed at $((i - 1) * 5)^\circ$ from the reference line. Similarly $lb(i)$ is the lower bound for each position.

$$\text{Upper bound, } \mathbf{ub} = \begin{bmatrix} 0.0714 & 0.0769 & 0.0714 & 0.0714 & 0.0667 & 0.0769 & 0.0714 \\ 0.0667 & 0.0667 & 0.0625 & 0.0625 & 0.0667 & 0.0667 & 0.0769 \end{bmatrix}$$

$$\text{Lower bound, } \mathbf{lb} = \begin{bmatrix} 0.0714 & 0.0769 & 0.0714 & 0.0714 & 0.0667 & 0.0769 & 0.0714 \\ 0.0667 & 0.0667 & 0.0625 & 0.0625 & 0.0667 & 0.0667 & 0.0769 \end{bmatrix}$$

The bounds are seen to be very tight, and the maximum rate is obtained when antennas are at 5° , 25° or -5° from the reference line. The maximum rate so obtained is 0.0769, giving a sum capacity of 3.076. Only 14 different positions of the antenna are considered for a 5 sector system, since the pattern would repeat itself after that.

Trying to optimize the rates such that the rate to each ST is maximized will adversely affect the sum capacity of the system. So, take $U(\mathbf{r}) = \sum_{j=1}^m (r_j)$.

For example, the sum capacity evaluated for antenna position varied in steps of 5° is as follows. It can be seen that the system capacity does not vary much with the position of the antenna. But, there seems to be some positions which are worse than the others.

$$\text{Upper bound, } \mathbf{ub} = [4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4]$$

$$\text{Lower bound, } \mathbf{lb} = [4 \ 4 \ 3.8046 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 4 \ 3.8883 \ 3.8699 \ 4 \ 3.9916]$$

Maximising $\sum_{i=1}^m \log(r(i))$ under the given constraints for upper bound and lower bound gives the utility functions for different positions of the antenna as

$$U_{lb} = \begin{bmatrix} -96.2916 & -95.7459 & -97.0998 & -95.112 & -95.9083 & -98.0191 & -96.1752 \\ -99.1991 & -101.1465 & -102.5186 & -102.0872 & -99.52350 & -98.37440 & -96.58240 \end{bmatrix}$$




Figure 28: Variation of sum rate and fairness index with antenna orientation for different utility functions.

$$U_{ub} = \begin{bmatrix} -96.2916 & -95.5485 & -97.0998 & -95.1122 & -95.9083 & -98.0191 & -96.1752 \\ -99.1991 & -101.1465 & -102.5186 & -102.0872 & -99.5235 & -97.3909 & -96.4809 \end{bmatrix}$$

The sum capacities for each of the rates above are bounded by

$$\text{Sum of rates for upperbound} = \begin{bmatrix} 3.9102 & 3.8463 & 3.7333 & 3.8611 & 3.8535 & 3.6159 & 3.7896 \\ 3.6663 & 3.3637 & 3.30270 & 3.3513 & 3.6291 & 3.7565 & 3.7844 \end{bmatrix}$$

$$\text{Sum of rates for lower bound} = \begin{bmatrix} 3.9102 & 3.8350 & 3.7333 & 3.8612 & 3.8535 & 3.6157 & 3.7896 \\ 3.6663 & 3.3636 & 3.3027 & 3.3512 & 3.6291 & 3.704 & 3.7734 \end{bmatrix}$$

The utility function is maximum when the antenna is positioned at 15° from the reference line. The sum of rates at this position is 3.86. This gives a trade-off between maximising the system capacity and providing fairness.

The sum of the rate given to STs and the fairness index vs antenna orientation is plotted in Figure 28. for different utility functions (the lower bounds are plotted here). Fairness index varies from 0 to 1. For a rate vector \mathbf{r} , the fairness index is given by

$$\gamma = \frac{(\sum_{i=1}^m x_i)^2}{m \sum_{i=1}^m x_i^2}$$

If the rates to different STs are equal, then fairness index would be 1, and it decreases as the rates are made unfair. The plots for maximum $\sum_{i=1}^m r_i$, maximum $\sum_{i=1}^m \log r_i$ are shown. It can be seen that maximising the sum rate gives high overall capacity, but poor fairness. On the other hand, maximising the average rate to each ST gives good fairness, but low sum capacity. Maximising $\sum_{i=1}^m \log r_i$ gives a good tradeoff between maximising the system capacity and providing fairness. It is interesting to note that in maximum $\sum_{i=1}^m \log r_i$ case, the sum capacity is higher when fairness is lower and viceversa. For example, at $\theta = 10$, we can see that the sum rate is close to 4. The fairness index is also close to 1. So, we may choose this orientation as optimum.

ANNEX C: Scheduler Design

1 Slot Scheduling

The scheduling problem is the following.

First partition the frame of size N slots into a contiguous part with N_D downlink slots and an uplink part with N_U uplink slots, such that $N_D + N_U = N - N_B$, where N_B is the number of slots required for the periodic beacon. Typically, $N_D \gg N_U$ as TCP data traffic is highly asymmetric since users download a lot more than they upload, and during downloads, long TCP packets (upto 1500 bytes) are received in the downlink and one 40 byte TCP ACK is sent in the uplink for alternate recieved packets.

Now, when $m_{v,i}$, $1 \leq i \leq m$, VOIP calls are admitted for ST i , one needs to determine the number of slots C_i to be reserved in the uplink and downlink subframes for ST i , such that the QoS targets are met for all the voice calls. For doing this, evidently the set of vectors $\mathbf{C} = (C_1, \dots, C_m)$ that are feasible (i.e., can be scheduled) needs to be known. For each deployment, there will be an optimal set of such vectors \mathcal{C}_{opt} , and for any practical scheduler, there will be an achievable set of admissible vectors $\mathcal{C} \in \mathcal{C}_{opt}$.

Once the required vector of voice payload slots has been scheduled, one has to schedule as many additional payload slots, so as to maximize the traffic carrying capacity for TCP while ensuring some fairness between the flows.

1.1 Representation of a Schedule

A feasible schedule for the deployment in Figure 29 is given in the table below, where each row stands for a sector and each column stands for a slots. There are 15 slots in the example. The entry in the table denotes the index of the ST that is transmitting in each sector in each slot. ST 1 in sector 2 and ST 2 in sector 1 transmit for the first 4 slots. At the end of 4th slot, ST 1 stops transmitting and ST 4 starts transmitting. The matrix representation of the schedule is also shown. The first 3 columns stands for the 3 sectors. Each entry shows the index of the ST that is transmitting. The last column indicates the number of slots for which that row is operative. Here, the sum of elements in column 4 is 15, indicating that there are 15 slots. This notation is followed throughout the examples.

			slots →														
2	2	2	2	2	2	2	2	2	0	0	0	0	0	0	0	0	$\begin{bmatrix} 2 & 1 & 0 & 4 \\ 2 & 4 & 0 & 4 \\ 0 & 4 & 6 & 7 \end{bmatrix}$
1	1	1	1	4	4	4	4	4	4	4	4	4	4	4	4		
0	0	0	0	0	0	0	0	0	6	6	6	6	6	6	6		




Figure 29: A system example showing the distribution of 6 stations in 3 sectors

1.2 A Greedy Heuristic Scheduler for the Uplink

The STs are scheduled such that the one with the longest queue is scheduled first. Find an activation vector which includes the ST with the largest voice queue. Next include a non interfering ST with the longest queue and so on until the number of STs in the activation set is equal to the number of simultaneous transmissions possible or till the activation set is maximal. A maximal activation set is one to which one cannot add any more links such that there is no interference between the links in the set. Use this maximal activation vector until one of the STs in the set completes transmission. Once one of the STs complete transmission, we remove that ST from the set and schedule another ST, that does not interfere with the STs in the set. Repeat the procedure until all the STs completes transmitting their voice packets. When all the STs complete their voice transmission, the remaining slots are used for TCP transmission. If at any stage during voice transmission, a situation occurs where there are no more noninterfering STs in a sector which can transmit voice, but there is one that can transmit data, schedule data for that interval.

In the beginning of each frame, for each slot k , heuristically build an activation vector $\mathbf{u}_k \in \mathcal{U}$ starting from an ST in $\{i : q_{k,i} = \max_j q_{k,j}\}$, i.e., the non interfering station with the maximum voice queue. Here, $q_{k,j}$ denotes the queue length of the j th ST at the beginning of the k th slot. k varies from 1 to N over a frame. Build a maximal activation vector beginning with that link, and augmenting the vector every time with a non interfering link.

$\mathcal{I}(\mathbf{u})$ denote the interference set of activation set \mathbf{u} , the set of links that can interfere with the STs in \mathbf{u} .

Algorithm 1.1

1. Modify the voice queue lengths to include the overhead slots required. i.e., If an ST has a voice queue of 2 packets, add 3 slots of PHY overhead to make the queue length 5.
2. Initially, slot index $k = 0$. Let ST i be such that

$$q_{ki} = \max_{l=1 \dots m} \{q_{kl}\}$$

i.e., The ST with longest voice queue at the beginning of slot k is i . Form activation vector \mathbf{u} with link i activated. i.e., $\mathbf{u} = \{i\}$

3. Let ST j be such that

$$q_{kj} = \max_l \{q_{kl} : l \notin \mathcal{I}(\mathbf{u})\}$$

j is such that it is the non interfering ST with maximum queue length. Augment \mathbf{s} with link j . Now, find $\mathcal{I}(\mathbf{u})$ corresponding to the new \mathbf{u} .

4. Repeat step 3 until activation vector that we get is a maximal activation vector.

5. Let

$$n = \{q_{kl} : \min_{l=1, \dots, m} (q_{kl}, l \in \mathbf{u})\}$$

i.e., n is the minimum number of slots required for the first ST in \mathbf{u} to complete its transmission. Use \mathbf{u} in the schedule from k th to $(k + n)$ th slot.

$$q_{k+n,i} = \begin{cases} q_{k,i} - n & \text{for } i \in \mathbf{u}' \\ q_{k,i} & \text{for } i \notin \mathbf{u}' \end{cases}$$

and $k = k+n$ i.e., slot index advances by n , and the queue length for the STs at the beginning of $k + n$ th slot is n less

6. At the end of $k + n$ th slot,

$$\mathbf{u} = \mathbf{u} - \{l : q_{kl} = \min(q_{kl}, l \in \mathbf{u})\}$$

i.e., remove from the activation vector, those STs that have completed their voice slot requirement.

7. Go back to Step 3 and form maximal activation vector including \mathbf{u} . Continue the above procedure until $\mathbf{q} = \mathbf{0}$ or $n = N_U$. In this step, we form a new activation vector with the remaining STs in the activation vector (which need more slots to complete their requirement).
8. Once the voice packets are transmitted, we serve the TCP packets in the same way, except that if in forming a maximal activation set, it is found that the only schedulable ST has only TCP packets to send, then TCP packets are scheduled.

If $\mathbf{q} > \mathbf{0}$ when $n = N_U$, the allocation is infeasible.

1.3 A Greedy Heuristic Scheduler for the Downlink

The difference of the downlink scheduling problem from the uplink scheduling problem is that in downlink, a transport block can contain packets to multiple STs. By combining the voice packets to different STs to a single TB, we save considerable PHY overhead. For transmitting a single voice packet needs 4 slots, where 3 slots are for the PHY header. Transmitting 2 voice packets need only 5 slots. So, it is always advantageous to have transmissions in longer blocks. This can

be done by grouping together the STs to those which are heard only by i th BTS, those heard by i th and $(i - 1)$ th BTS, but associated to the i th BTS and those heard by i th and $(i - 1)$ th BTS, but associated to the $(i - 1)$ th BTS, for all values of i .

For example in Figure 30 STs 3 and 4 are associated with BTS 1 and hears only form BTS 1. So, any ST in the interference set of 3 will also be in the interference set of 4. Any transmission to ST 3 can equivalently be replaced by a transmission to 4. So, they form a group for the down linkn schedule. Similarly, STs 8 and 9 are associated to BTS 2 and interfere with BTS 3. They are associated to the same BTS and cause interference to the same STs. So, ST 8 and 9 also form a group.




Figure 30: Typical deployment of a system with 3 sectors and 15 STs

The STs are grouped together based on the above criterion. The queue length of each group would be the sum of queue lengths of the STs forming the group. The greedy heuristic scheduler for the uplink scheduling problem can then be used over these groups. This is made clearer in Example 2.3

1.4 Round Robin Scheduling

A scheduler with least complexity can be designed as follows. The uplink and downlink parts of the frame may further be divided into two contiguous parts. Alternate sectors are served in these two parts. For example, Sectors 1, 3, 5 are served in the first part, and Sectors 2, 4, 6 can be served in the second part of the frame. Interference between adjacent sectors can be eliminated in this way. With the number of sectors close to $2n_0$, performance of this scheduler would be equivalent to that of the scheduler discussed in Section 1.2, since we can have n_0 transmissions going on in each slot, with this scheduler. But, with $n_0 = 4$, this would require 8 sectors in the system. With number of sectors less than $2n_0$, the number of simultaneous transmissions would be less than n_0 with the round robin algorithm, where as we can have upto n_0 transmissions with the greedy algorithm.

1.5 Fair Scheduling

Having all voice transmission in the beginning and data afterwards is wasteful in terms of overhead. Following the principle of having longer TBs, it is advantageous to have the voice and data transfer to an ST in a block. To provide fairness to users, keep track of the average rates allocated to STs over time. The STs with low average rate over frames are given a chance to transmit first. Maximal

independent sets are formed starting from the ST with the lowest average rate. Once the slots for voice transmission are scheduled, the attempt should be to have TCP transmission in blocks of size T_{max} , so that PHY overhead per slot is minimised.

Let \mathbf{R}_k be the vector of average rates allocated to STs till the k th frame and \mathbf{r}_k be the vector of rates allocated to the STs in the k th slot, i.e, the fraction of slots allocated to STs in the k th frame. The average rate achieved by the STs is obtained by computing the following geometrically weighted average. A large value of α places less weight on the previous frames.

$$\mathbf{R}_{k+1} = \alpha \mathbf{R}_k + (1 - \alpha) \mathbf{r}_k$$

1. Given a rate vector \mathbf{R} , obtain a maximal independent set as follows

(a) $\mathbf{u}_1 = \{i_1\}$

$$i_1 = \arg \min_{1 \leq j \leq n} \mathbf{R}_j$$

$\mathcal{I}(\mathbf{u}_1)$ is the set of links interfering with the links in \mathbf{u}_1 . In this step, we select the ST with the smallest average rate R_k for transmission.

(b) Choose $i_2 \in \arg \min_{1 \leq j \leq n, i_2 \notin \mathcal{I}(\mathbf{u}_1)} \mathbf{R}_j$

$\mathbf{u}_1 = \{i_1, i_2\}$. In this step, we select one of the non interfering STs with minimum average rate for transmission.

(c) Repeat the above until a maximal independent set is obtained. Now, we have a set with STs which have received low average rates in the previous slots. So, once all STs transmit their voice packets, we schedule these STs for data packets.

2. Let l_1 denote the number of nodes in \mathbf{u}_1 at the end of step 1. Repeat the above for the remaining $n - l_1$ nodes. Now we have a maximal independent set from the remaining $N_u - l_1$ nodes. If any one of the l_1 nodes can be activated along with the maximal independent set formed from the $N_u - l_1$ nodes, add that till one get a maximal independent set. This yields $\mathbf{u}_1, \mathbf{u}_2 \dots \mathbf{u}_k$ such that each node is included atleast once. Each node is included atleast once since a given number of slots is to be reserved for each ST in every frame.

3. Now, we need to schedule \mathbf{u}_1 for t_1 , \mathbf{u}_2 for t_2 , etc. To maximize throughput, we take $t_j = T_{max}$ or number of voice slots required. The vectors in the initial part of the schedule had low average rate over frames. So, they get priority to send data packets. So, starting from $j=1$, i.e., from the first activation vector, if the sum of number of slots allocated to STs in the frame is less than N_u , $t_j = T_{max}$. Else, $t_j =$ number of voice slots required. Therefore transmission takes place in blocks of length equal to T_{max} as long as it is possible.

4. Update the rate vector as

$$\mathbf{R}_{k+1} = \alpha \mathbf{R}_k + (1 - \alpha) \mathbf{r}_k$$

2 Design Methodology: Examples

The deployment considered in the examples are described by the \mathbf{A} and \mathbf{I} matrices given in Table 2. The \mathbf{A} matrix shows the BTS to which each ST is associated. Columns in both matrices correspond

to sectors and rows correspond to STS. $A_{ij} = 1$ indicates that the i th ST is associated to the j th BTS. Each ST is associated to one and only one BTS. $I_{ij} = 1$ indicates that the i th ST can hear from the j th BTS. An ST might hear from more than one BTS.

Example 2.1

Equal Voice Load

In this example there are 40 stations in 5 sectors. Any station within 72° is associated with the BTS and any station within 32° of the center of a mainlobe can cause interference at that BTS. We take the maximum burst length, T_{max} to be 15 and $N_U = 112$. Voice packet (including MAC header) is 44 bytes long. So, each voice packet requires one slot time for transmission. In this example, we take $n_0 = 4$.

Consider all stations to have the same voice slot requirement of 4 packets. The uplink transmission from each station could be done in a single block , which requires a total of 3 slots of overhead. So, the schedule should give at least 7 slots for each station. One schedule that attains this, (as obtained Algorithm 1.1) is given by S_2 .

The first 5 columns of the matrix S_2 gives the links that can transmit in each sector. A 0 indicates that there cannot be a transmission by any of the stations in this sector without causing interference to at least one of the ongoing transmissions. The entries in the 6th column indicates the number of consecutive slots for which these vector is used in the schedule. The 7th column is the number of slots (over all the sectors) used for transmitting voice packets. The last column shows the number of slots, over all the sectors used for transmitting TCP packets.

For example, consider the first row of S_2 . This indicates that the ST 1 from Sector 1, 2 from Sector 2, 11 from Sector 3, 6 from Sector 4 and none from Sector 5 are scheduled for the first 7 slots. At the end of 7 slots, the voice queues of these stations are exhausted, so we schedule other stations.

From S_2 , one sees that ST 40 gets a chance for transmission only in the last row, i.e., by the 98th slot. It transmits till the 105th slot to complete transmitting its voice queue. This indicates that the scheduling of voice queues requires 105 slots.

In the 1st row, ST 11 has been scheduled for 7 slots. So the voice queue of ST 11 is exhausted by the end of that row. But, since there are no other STs in Sector 3, that can use the slot for voice transmission, we schedule ST 11 for TCP transmission further in the schedule. In this algorithm, priority is to STs that had been transmitting in the previous slot, and next priority to stations with lower index. But, better fairness can be ensured by round robin between the STs in a sector.

Since the voice queue of ST 11 is exhausted by the end of 1st row in S_2 the ST 11 scheduled after 1st row, i.e., from 29th slot is for TCP transmission. Similarly, ST 7 scheduled after the 11th row, ST 39 scheduled after 9th row etc. are for TCP transmission.

In this example, 4 STs transmit in each slot (i.e., 3×26 slots of transmission.). Thereby, of the 112×5 slots available (since in each slot, at most one ST from a sector can transmit.), 112×4 slots are used for transmission, giving a throughput of 78% and a goodput efficiency of 46%. In this schedule, transmission of TCP packets occur in 105 slots.

$\mathbf{A} =$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}$	$\mathbf{I} =$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \end{bmatrix}$	$\mathbf{Q}_1 =$	$\begin{bmatrix} 3 \\ 2 \\ 3 \\ 3 \\ 4 \\ 5 \\ 5 \\ 2 \\ 4 \\ 4 \\ 4 \\ 5 \\ 3 \\ 2 \\ 5 \\ 2 \\ 4 \\ 5 \\ 5 \\ 4 \\ 4 \\ 3 \\ 2 \\ 4 \\ 4 \\ 3 \\ 5 \\ 5 \\ 3 \\ 4 \\ 5 \\ 2 \\ 2 \\ 5 \\ 3 \\ 3 \\ 2 \\ 5 \\ 4 \end{bmatrix}$	$\mathcal{S}_1 =$	$\begin{bmatrix} 0 & 10 & 15 & 6 & 5 & 7 & 28 & 0 \\ 0 & 10 & 24 & 6 & 5 & 1 & 4 & 0 \\ 16 & 0 & 24 & 13 & 27 & 6 & 24 & 0 \\ 16 & 39 & 0 & 13 & 27 & 2 & 8 & 0 \\ 32 & 39 & 11 & 20 & 0 & 6 & 24 & 0 \\ 32 & 4 & 12 & 20 & 0 & 2 & 8 & 0 \\ 8 & 4 & 12 & 0 & 26 & 3 & 12 & 0 \\ 8 & 4 & 11 & 0 & 26 & 2 & 6 & 2 \\ 8 & 19 & 11 & 0 & 26 & 2 & 6 & 2 \\ 28 & 19 & 11 & 0 & 26 & 1 & 3 & 1 \\ 28 & 19 & 0 & 34 & 30 & 4 & 16 & 0 \\ 28 & 18 & 0 & 34 & 30 & 1 & 4 & 0 \\ 7 & 18 & 0 & 34 & 30 & 3 & 12 & 0 \\ 7 & 18 & 0 & 35 & 25 & 2 & 8 & 0 \\ 31 & 18 & 0 & 35 & 25 & 1 & 4 & 0 \\ 31 & 0 & 40 & 35 & 25 & 3 & 12 & 0 \\ 31 & 0 & 40 & 35 & 38 & 1 & 4 & 0 \\ 7 & 0 & 40 & 35 & 38 & 1 & 3 & 1 \\ 7 & 0 & 40 & 36 & 38 & 2 & 6 & 2 \\ 17 & 0 & 11 & 36 & 38 & 1 & 3 & 1 \\ 17 & 0 & 11 & 36 & 38 & 3 & 6 & 6 \\ 17 & 0 & 11 & 37 & 38 & 4 & 8 & 8 \\ 9 & 21 & 11 & 37 & 0 & 2 & 6 & 2 \\ 9 & 21 & 11 & 37 & 0 & 4 & 8 & 8 \\ 9 & 2 & 11 & 37 & 0 & 1 & 2 & 2 \\ 0 & 2 & 11 & 37 & 23 & 4 & 8 & 8 \\ 0 & 2 & 11 & 37 & 23 & 3 & 21 & 7 \\ 7 & 2 & 11 & 0 & 29 & 7 & 21 & 7 \\ 7 & 2 & 11 & 33 & 0 & 7 & 21 & 7 \\ 1 & 2 & 11 & 33 & 0 & 6 & 18 & 6 \\ 1 & 2 & 11 & 3 & 0 & 6 & 18 & 6 \\ 1 & 2 & 14 & 0 & 0 & 5 & 15 & 5 \\ 1 & 2 & 0 & 22 & 0 & 5 & 15 & 5 \end{bmatrix}$
$\mathbf{Q}_2 =$	$\begin{bmatrix} 4 \\ 4 \\ 4 \\ \vdots \\ 4 \end{bmatrix}$	$\mathcal{S}_2 =$	$\begin{bmatrix} 1 & 2 & 11 & 6 & 0 & 7 & 28 & 0 \\ 7 & 4 & 12 & 3 & 0 & 7 & 28 & 0 \\ 8 & 0 & 15 & 13 & 5 & 7 & 28 & 0 \\ 9 & 19 & 24 & 34 & 0 & 7 & 28 & 0 \\ 0 & 10 & 11 & 35 & 25 & 7 & 28 & 0 \\ 16 & 21 & 14 & 0 & 27 & 7 & 28 & 0 \\ 17 & 0 & 11 & 36 & 30 & 7 & 21 & 7 \\ 28 & 18 & 0 & 37 & 38 & 7 & 28 & 0 \\ 31 & 39 & 11 & 20 & 0 & 7 & 21 & 7 \\ 32 & 39 & 0 & 22 & 5 & 7 & 21 & 7 \\ 0 & 39 & 0 & 22 & 23 & 7 & 7 & 14 \\ 7 & 39 & 11 & 0 & 26 & 7 & 7 & 21 \\ 7 & 39 & 11 & 0 & 29 & 7 & 21 & 7 \\ 7 & 39 & 11 & 33 & 0 & 7 & 21 & 7 \\ 7 & 0 & 40 & 33 & 0 & 7 & 14 & 7 \end{bmatrix}$				

Table 2: Data and results for the Example 2.1 and 2.2. \mathcal{S}_i is the schedule for voice slot requirements given by $\mathbf{Q}_i, i \in \{1, 2\}$. Note that the last two columns of \mathcal{S}_i are the total number of slots in all sectors for voice and TCP packets respectively scheduled in each row.

Example 2.2

Unequal Voice Load

Also given is the schedule obtained by the Algorithm 1.1, when the voice slot requirements are different for different stations. Given is one such vector (\mathbf{Q}_1) and the corresponding schedule (\mathbf{S}_1).

In the example, in the first row, i.e., links 10,15,6,5 are scheduled for the first 7 slots. At the end of the 7th slot, ST 15 exhausts its voice queue. So, remove ST 15 from the vector, and add ST 24 which can take its place. This continues till one of the other queues become empty. Then, remove the one whose queue is empty, and schedule some other ST which does not interfere with the ongoing transmissions, and so on till all STs are served.

All the STs get a chance to transmit by the last row of the given schedule. But, the last row finishes transmission by the 108th slot (as seen by summing up the last column of \mathbf{S}_1). The slots from 108 to 112 may be used for transmitting TCP packets for the STs in the last row. Here, the throughput 78% and a goodput efficiency of 46%. In this schedule, transmission of TCP packets occurs in 147 slots.

Example 2.3

Downlink

In downlink, transmission to multiple STs can be done in a block. In the deployment in this example, STs 2, 4, 19, 21, 39 all hear only from BTS 2. This can be seen from the \mathbf{A} and \mathbf{I} matrices. All these STs have 1 in the second column, indicating that they are associated to BTS 2. The \mathbf{I} matrix has 1s only in the second column, indicating that they hear only from BTS 2. These are grouped in to b . The transmissions to these STs can take place in a TB, since an ST which interferes with any ST in h will interfere with every other ST in h , and vice versa. So every ST in h is equivalent with respect to interference constraints. An h occurring in the schedule in Table 3 indicates a transmission to this group of STs. The STs which are in the same group are associated to the same BTS and are in the exclusion region of same sectors. The groups so formed for the example are given in Table 3. Similarly, the STs 3,20 and 33 hear from BTSs 4 and 5; all of them are associated to BTS 4. They are grouped to c . All STs in c are equivalent with respect to the interference constraints. The transmissions to these STs can occur in a TB. Now, the voice slot requirement of a group is the sum of the voice slot requirements of individual STs constituting the group. The same algorithm as for the uplink is then used over the groups a, b, c, d, \dots with their respective queue lengths to obtain the schedule in downlink. Aggregating STs in the association and exclusion regions of each BTS like this has the advantage of increasing the overall system throughput, since transmissions occur in longer TBs. The goodput efficiency in downlink is about 64%.

$$\mathbf{Q}_{2d} = \begin{bmatrix} 2 \\ 2 \\ \vdots \\ 2 \end{bmatrix} \quad \mathbf{S}_{2d} = \begin{bmatrix} f & b & 0 & e & d & 31 \\ f & b & h & e & 0 & 5 \\ f & b & h & c & 0 & 18 \\ a & b & h & c & 0 & 13 \\ f & b & h & 0 & n & 13 \\ 0 & g & h & 0 & n & 8 \\ 0 & g & i & 0 & n & 8 \\ j & 0 & i & 0 & n & 8 \\ f & k & 0 & 0 & n & 8 \\ f & k & 0 & l & d & 8 \\ 0 & k & 0 & l & m & 8 \\ 0 & 0 & o & e & m & 8 \end{bmatrix}$$

$a = \{1, 9\}$	$b = \{2, 4, 19, 21, 39\}$	$c = \{3, 20, 33\}$
$d = \{5, 25, 27, 30, 38\}$	$e = \{6, 13, 34, 35, 36, 37\}$	$f = \{7, 8, 16, 28, 31, 32\}$
$g = \{10\}$	$h = \{11, 12, 15, 24\}$	$i = \{14\}$
$j = \{17\}$	$k = \{18\}$	$l = \{22\}$
$m = \{23\}$	$n = \{26, 29\}$	$o = \{40\}$

Table 3: The part of the downlink schedule for completing the voice slot requirement when 2 slots are reserved for voice transmission from each ST.

ANNEX D: Simulation Analysis

The fair scheduling algorithm and the round robin algorithm discussed in ANNEX C were implemented in MATLAB. The voice and data capacity of the system using these algorithms is provided in this section.

1 Simulation Scenario

We consider the distribution of n STs in m sectors. Results for different values of m and n are given. All STs are assumed to be carrying the same number of voice calls. One VOIP call requires one slot every alternate frame. A voice packet that arrives in the system is scheduled within the next two frames. If the system capacity and scheduling constraints do not allow the voice packet to be transmitted within two frame times of its arrival, the packet is dropped. In this simulation, we assume synchronous arrival of voice packets. i.e., if two voice calls are going on from an ST, packets for both calls arrive synchronously, in the same frame. The data throughput considered is the saturation throughput, i.e., the STs have packets to be transmitted throughout. All STs have infinite queues and the schedule is driven by the average rates obtained over time. The STs which received low service rates in the previous frames are scheduled first. (See ANNEX C, Section 1.5)

2 Numeric Results

Following the fair scheduling algorithm described earlier, the data throughputs attainable for n STs in m sectors are calculated for different number of voice calls, widths of taboo regions, n and m . For each value of n , m , width of taboo region (θ°), and voice call requirement, the attainable data throughput is calculated in terms of fraction of downlink bandwidth, for 30 different deployments and the averages are tabulated. The minimum average throughput (in slots per slot time) attained among the STs, the maximum average throughput among STs, the sum of the average throughputs to STs and also the probability of a voice packet being dropped are tabulated.

The first column of Tables 4 to 8 gives the number of sectors, number of STs, width of taboo region and the maximum number of simultaneous transmissions possible, n_0 . *minrate* is the minimum average rate over STs, averaged over random deployments, *maxrate* is the maximum average rate over STs, *sumrate* is the sum of average rates to STs, and the average fraction of voice packets dropped.

For example, one can see from Table 4 that for 40 STs in 6 sectors, with a taboo region of width 10° and $n_0 = 3$, the minimum average downlink throughput provided to an ST is 0.0507 of the downlink bandwidth. If one slot is reserved every two frames for voice calls, the downlink bandwidth being 66% of the total bandwidth of 11 Mbps, the ST gets 370 Kbps. Similarly, if two slots are reserved per ST every two frames, each ST receives an average of at least 369 Kbps. The average fraction of voice packets dropped is zero when the number of STs is as small as 40.

		Number of voice calls per station			
number of sectors, number of sSTs, n_0 width of taboo region, θ		1	2	3	4
6, 40, $n_0=3$, $\theta = 10^\circ$	min d/l rate	0.0507	0.0504	0.0486	0.0453
	max d/l rate	0.0557	0.0545	0.0538	0.0524
	sum d/l rates	2.1379	2.0909	2.0342	1.9766
	min u/l rate	0.0337	0.0288	0.0277	0.0203
	max u/l rate	0.0733	0.0636	0.0578	0.0511
	sum u/l rate	2.0641	1.8120	1.6330	1.4013
	packet drop u/l	0	0	0	0
6, 40, $n_0=3$, $\theta = 20^\circ$	min d/l rate	0.0507	0.0487	0.0472	0.0452
	max d/l rate	0.0547	0.0552	0.0535	0.0528
	sum d/l rates	2.1086	2.0574	2.0171	1.9272
	min u/l rate	0.0347	0.0299	0.0238	0.0190
	max u/l rate	0.0712	0.0664	0.0596	0.0547
	sum u/l rate	2.0544	1.8157	1.6002	1.3898
	packet drop u/l	0	0	0	0
6, 40, $n_0=3$, $\theta = 30^\circ$	min d/l rate	0.0526	0.0479	0.0472	0.0437
	max d/l rate	0.0633	0.0613	0.0570	0.0598
	sum d/l rates	2.0865	2.0023	1.9881	1.9740
	min u/l rate	0.0298	0.0276	0.0208	0.0172
	max u/l rate	0.0843	0.0759	0.0749	0.0624
	sum u/l rate	2.0484	1.8204	1.5909	1.3926
	packet drop u/l	0	0	0	0
6, 60, $n_0=3$, $\theta = 10^\circ$	min d/l rate	0.0326	0.0312	0.0297	0.0270
	max d/l rate	0.0380	0.0359	0.0339	0.0322
	sum d/l rates	2.0792	2.0034	1.8873	1.7249
	min u/l rate	0.0137	0.0094	0.0050	0.0026
	max u/l rate	0.0448	0.0357	0.0274	0.0197
	sum u/l rate	1.4863	1.2020	0.8714	0.6368
	packet drop u/l	0	0	0.0022	0.0156
6, 60, $n_0=3$, $\theta = 20^\circ$	min d/l rate	0.0325	0.0317	0.0293	0.0267
	max d/l rate	0.0363	0.0368	0.0340	0.0321
	sum d/l rates	2.0472	1.9924	1.8861	1.7338
	min u/l rate	0.0130	0.0091	0.0055	0.0018
	max u/l rate	0.0444	0.0398	0.0291	0.0198
	sum u/l rate	1.5189	1.2006	0.9036	0.6327
	packet drop u/l	0	0	0.0029	0.0167
6, 60, $n_0=3$, $\theta = 30^\circ$	min d/l rate	0.0338	0.0326	0.0293	0.0260
	max d/l rate	0.0375	0.0358	0.0337	0.0318
	sum d/l rates	2.1148	2.0344	1.8997	1.7368
	min u/l rate	0.0139	0.0087	0.0049	0
	max u/l rate	0.0452	0.0360	0.0326	0.0257
	sum u/l rate	1.4985	1.1587	0.9299	0.6301
	packet drop u/l	0	0	0.0033	0.0179
6,80, $n_0=3$, $\theta = 10^\circ$	min d/l rate	0.0244	0.0220	0.0200	
	max d/l rate	0.0265	0.0271	0.0248	
	sum d/l rates	2.0453	1.9118	1.7390	
	min u/l rate	0.0056	0.0025	0	
	max u/l rate	0.0245	0.0160	0.0091	
	sum u/l rate	1.0388	0.6563	0.3505	
	packet drop u/l	0	0.0029	0.0229	
6,80, $n_0=3$, $\theta = 20^\circ$	min d/l rate	0.0242	0.0224	0.0203	
	max d/l rate	0.0267	0.0257	0.0263	
	sum d/l rates	2.0150	1.9039	1.7552	
	min u/l rate	0.0048	0	0	
	max u/l rate	0.0252	0.0185	0.0097	
	sum u/l rate	1.0604	0.6305	0.3352	
	packet drop u/l	0	0.0033	0.0312	
6,80, $n_0=3$, $\theta = 30^\circ$	min d/l rate	0.0249	0.0228	0.0204	
	max d/l rate	0.0268	0.0258	0.0240	
	sum d/l rates	2.0653	1.9339	1.7480	
	min u/l rate	0.0045	0.0022	0	
	max u/l rate	0.0249	0.0174	0.0114	
	sum u/l rate	1.0319	0.6797	0.3399	
	packet drop u/l	0	0.0042	0.0346	

Table 4: The average throughput per ST averaged over 30 random deployments. The data throughputs are tabulated for different number of sectors, number of STs, width of taboo region and $n_0 = 3, 4$. Throughputs are in terms of fraction of uplink or downlink bandwidth respectively

number of sectors, number of sSTs, n_0 width of taboo region, θ		Number of voice calls per station	
		1	2
6, 100, $n_0=3$, $\theta = 10^\circ$	min d/l rate	0.0187	0.0168
	max d/l rate	0.0219	0.0202
	sum d/l rates	2.0074	1.8184
	min u/l rate	0.0011	0
	max u/l rate	0.0146	0.0068
	sum u/l rate	0.6103	0.2755
6, 100, $n_0=3$, $\theta = 20^\circ$	min d/l rate	0.0192	0.0174
	max d/l rate	0.0217	0.0202
	sum d/l rates	2.0010	1.8327
	min u/l rate	0.0010	0
	max u/l rate	0.0129	0.0063
	sum u/l rate	0.5927	0.2576
6, 100, $n_0=3$, $\theta = 30^\circ$	min d/l rate	0.0194	0.0173
	max d/l rate	0.0211	0.0205
	sum d/l rates	2.0005	1.8133
	min u/l rate	0	0
	max u/l rate	0.0142	0.0098
	sum u/l rate	0.5860	0.3040
6, 120, $n_0=3$, $\theta = 10^\circ$	min d/l rate	0.0155	
	max d/l rate	0.0184	
	sum d/l rates	1.9903	
	min u/l rate	0	
	max u/l rate	0.0057	
	sum u/l rate	0.2305	
6, 120, $n_0=3$, $\theta = 20^\circ$	min d/l rate	0.0157	
	max d/l rate	0.0174	
	sum d/l rates	1.9538	
	min u/l rate	0	
	max u/l rate	0.0052	
	sum u/l rate	0.2177	
6, 120, $n_0=3$, $\theta = 30^\circ$	min d/l rate	0.0139	
	max d/l rate	0.0153	
	sum d/l rates	1.7393	
	min u/l rate	0	
	max u/l rate	0.0052	
	sum u/l rate	0.2117	
	packet drop u/l	0.0394	

Table 5: The average throughput per ST averaged over 30 random deployments. The data throughputs are tabulated for different number of sectors, number of STs, width of taboo region and $n_0 = 3, 4$ (Contd)

number of sectors, number of sSTs, n_0 width of taboo region, θ		Number of voice calls per station			
		1	2	3	4
6, 40, $n_0=4$, $\theta = 10^\circ$	min d/l rate	0.0658	0.0658	0.0661	0.0632
	max d/l rate	0.0808	0.0782	0.0772	0.0777
	sum d/l rates	2.9443	2.8665	2.8532	2.7495
	min u/l rate	0.0452	0.0371	0.0333	0.0256
	max u/l rate	0.1214	0.1125	0.1027	0.1029
	sum u/l rate	2.9113	2.6910	2.4730	2.3426
	packet drop u/l	0	0	0	0
6, 40, $n_0=4$, $\theta = 20^\circ$	min d/l rate	0.0616	0.0607	0.0606	0.0587
	max d/l rate	0.0866	0.0756	0.0758	0.0708
	sum d/l rates	2.9434	2.7094	2.6901	2.5764
	min u/l rate	0.0405	0.0338	0.0267	0.0219
	max u/l rate	0.1462	0.1526	0.1491	0.1513
	sum u/l rate	2.7302	2.4593	2.3504	2.1716
	packet drop u/l	0	0	0	0
6, 40, $n_0=4$, $\theta = 30^\circ$	min d/l rate	0.0550	0.0536	0.0530	0.0480
	max d/l rate	0.0681	0.0666	0.0653	0.0690
	sum d/l rates	2.4473	2.3727	2.3198	2.1754
	min u/l rate	0.0326	0.0277	0.0227	0.0173
	max u/l rate	0.1110	0.0923	0.0972	0.0866
	sum u/l rate	2.1184	1.8977	1.7520	1.5281
	packet drop u/l	0	0	0	0
6, 60, $n_0=4$, $\theta = 10^\circ$	min d/l rate	0.0469	0.0436	0.0404	0.0382
	max d/l rate	0.0556	0.0552	0.0497	0.0473
	sum d/l rates	3.0059	2.8453	2.6891	2.5362
	min u/l rate	0.0233	0.0144	0.0088	0.0022
	max u/l rate	0.0624	0.0602	0.0513	0.0484
	sum u/l rate	2.2345	1.8461	1.5598	1.3567
	packet drop u/l	0	0	0.0004	0.0062
6, 60, $n_0=4$, $\theta = 20^\circ$	min d/l rate	0.0465	0.0429	0.0401	0.0379
	max d/l rate	0.0567	0.0547	0.0506	0.0500
	sum d/l rates	3.0010	2.8456	2.6901	2.5312
	min u/l rate	0.0192	0.0126	0.0069	0.0039
	max u/l rate	0.0944	0.0870	0.0795	0.0695
	sum u/l rate	2.1408	1.7864	1.4975	1.2942
	packet drop u/l	0	0	0	0.0083
6, 60, $n_0=4$, $\theta = 30^\circ$	min d/l rate	0.0352	0.0334	0.0305	0.0283
	max d/l rate	0.0469	0.0439	0.0419	0.0343
	sum d/l rates	2.3692	2.2139	2.0396	1.8635
	min u/l rate	0.0139	0.0078	0.0033	0.0018
	max u/l rate	0.0521	0.0490	0.0417	0.0350
	sum u/l rate	1.5640	1.2581	0.9693	0.7319
	packet drop u/l	0	0	0	0.0083
6, 80, $n_0=4$, $\theta = 10^\circ$	min d/l rate	0.0333	0.0304	0.0282	
	max d/l rate	0.0438	0.0413	0.0384	
	sum d/l rates	2.9465	2.7337	2.5300	
	min u/l rate	0.0105	0.0069	0	
	max u/l rate	0.0292	0.0255	0.0236	
	sum u/l rate	1.4752	1.1493	0.9122	
	packet drop u/l	0	0.0029	0.0283	
6, 80, $n_0=4$, $\theta = 20^\circ$	min d/l rate	0.0304	0.0288	0.0264	
	max d/l rate	0.0421	0.0380	0.0408	
	sum d/l rates	2.8728	2.6641	2.4441	
	min u/l rate	0.0069	0.0022	0	
	max u/l rate	0.0399	0.0389	0.0477	
	sum u/l rate	1.4386	1.1273	0.8810	
	packet drop u/l	0	0.0025	0.0304	
6, 80, $n_0=4$, $\theta = 30^\circ$	min d/l rate	0.0256	0.0246	0.0209	
	max d/l rate	0.0316	0.0309	0.0282	
	sum d/l rates	2.3167	2.0943	1.85944	
	min u/l rate	0.0040	0.0016	0	
	max u/l rate	0.0263	0.0205	0.0155	
	sum u/l rate	1.0279	0.6997	0.4316	
	packet drop u/l	0	0.0029	0.0254	

Table 6: The average throughput per ST averaged over 30 random deployments. The data throughputs are tabulated for different number of sectors, number of STs, width of taboo region and $n_0 = 3, 4$ (Contd)

number of sectors, number of sSTs, n_0 width of taboo region, θ		Number of voice calls per station	
		1	2
6, 100, $n_0=4$, $\theta = 10^\circ$	min d/l rate	0.0260	0.0239
	max d/l rate	0.0324	0.0301
	sum d/l rates	2.8934	2.6314
	min u/l rate	0.0035	0
	max u/l rate	0.0176	0.0105
	sum u/l rate	0.9422	0.6522
6, 100, $n_0=4$, $\theta = 20^\circ$	min d/l rate	0.0257	0.0226
	max d/l rate	0.0332	0.0315
	sum d/l rates	2.8303	2.5684
	min u/l rate	0.0017	0
	max u/l rate	0.0318	0.0270
	sum u/l rate	0.9711	0.6803
6, 100, $n_0=4$, $\theta = 30^\circ$	min d/l rate	0.0210	0.0181
	max d/l rate	0.0250	0.0240
	sum d/l rates	2.2563	1.9626
	min u/l rate	0	0
	max u/l rate	0.0159	0.0107
	sum u/l rate	0.6055	0.3280
6, 120, $n_0=4$, $\theta = 10^\circ$	min d/l rate	0.0184	
	max d/l rate	0.0269	
	sum d/l rates	2.5225	
	min u/l rate	0	
	max u/l rate	0.0099	
	sum u/l rate	0.5529	
	packet drop u/l	0.0750	

Table 7: The average throughput per ST averaged over 30 random deployments. The data throughputs are tabulated for different number of sectors, number of STs, width of taboo region and $n_0 = 3, 4$ (Contd)

As the number of STs goes up to 80, n_0 remaining as 3, one can reserve at most 2 slot per ST in alternate frames, with acceptable probability of dropping voice packets. With 80 STs, the probability of dropping a voice packet is more than 2 percent when we reserve slots for 2 calls per ST. With 80 STs, and slots reserved for one call per ST, each ST gets at least 123kbps of downlink bandwidth. But the uplink bandwidth is as small as 17 kbps. But, since the internet traffic is predominantly downlink, the capacity required in the uplink is small.

The rates given in Tables 4 to 8 are in terms of fraction of uplink of downlink bandwidth available for data transfer. Total uplink bandwidth is one third of 11 Mbps and downlink bandwidth is two thirds of 11 Mbps. So, a downlink rate of 0.0507 should be read as an uplink downlink bandwidth of a fraction of 0.0507 of $11 \times \frac{2}{3}$ Mbps. Similarly an uplink bandwidth of 0.0307 should be read as a fraction 0.0307 of $11 \times \frac{1}{3}$ Mbps. i.e., 11 kbps.

As seen from the tables, with 6 sectors, and $n_0 = 3$, the capacity remains almost the same irrespective of the width of taboo region. This happens because even if the taboo region of a sector covers half of the adjacent sector, we can form maximal sets with 3 links. But, with $n_0 = 4$, the capacity decreases with an increase in the width of taboo region. With a taboo region of width 30° , in a 6 sector system, we can have at most 3 links scheduled at a time.

no. of STs		Number of voice calls per station			
		1	2	3	4
40	min d/l rate	0.0531	0.0508	0.0479	0.0457
	max d/l rate	0.0683	0.0657	0.0629	0.0604
	sum d/l rates	2.2938	2.1899	2.0859	1.9813
60	min d/l rate	0.0344	0.0319	0.0295	0.0267
	max d/l rate	0.0421	0.0390	0.0364	0.0341
	sum d/l rates	2.2438	2.0875	1.9313	1.7750
80	min d/l rate	0.0255	0.0229	0.0202	
	max d/l rate	0.0296	0.0271	0.0245	
	sum d/l rates	2.1917	1.9833	1.7750	
100	min d/l rate	0.0199	0.0172	0.0147	
	max d/l rate	0.0227	0.0202	0.0175	
	sum d/l rates	2.1396	1.8792	1.6188	
120	min d/l rate	0.0110			
	max d/l rate	0.0132			
	sum d/l rates	1.4625			

Table 8: The average data throughput per ST with round robin scheduler for number of sectors=6, number of STs 40, 60, 80, 100 and 120

Depending upon the antenna pattern and the terrain, n_0 may go up to 4. With $n_0 = 4$, we can schedule slots for 2 calls per ST with a probability of voice packet dropping as small as 0.29%.

The minimum, maximum and sum of average data throughput over STs over different deployments are tabulated in Table 8 for the round robin scheduler, for different values of voice slot requirement. The average behavior of the round robin scheduler is almost the same as the greedy heuristics scheduler. But, in a given deployment, the greedy heuristic scheduler is found to be more fair. For example consider a deployment, where the adjacent sectors have 3 and 10 STs respectively. With a scheduler that round robins among the sectors, all STs in the sector with 3 STs will get 1/3 of the downlink bandwidth, whereas all those in a sector with 10 STs get only 1/10 of the downlink bandwidth. But, if we were allowed to select the STs in the activation set based on queue lengths, or on average rates, with the condition that no two interfering transmissions are allowed simultaneously, it is possible to vary the fraction of time each sector is allowed to transmit, based on the number of STs in each sector (which will be reflected in queue lengths or average rates), and thereby provide better fairness.

3 OPNET model and Simulation

Development of an OPNET model for the WiFiRe protocol, including the fair scheduling algorithm, is ongoing. Such a development will not only be useful to further understand the behaviour of the WiFiRe System under various conditions, but also may serve as a building block for the actual implementation of the protocol.

A description of the preliminary work completed so far is given in this section. This makes the following assumptions: (i) The ST(s) and BS are initialized at the same time. No ranging is performed. (ii) BS after becoming ready, starts transmitting MAPs. Each ST receives MAPs and gets associated with the BS antenna from which it receives maximum power. (iii) *Dynamic Service Addition* messages are exchanged and the simulation starts; *Dynamic Service Change* messages are not supported at present.

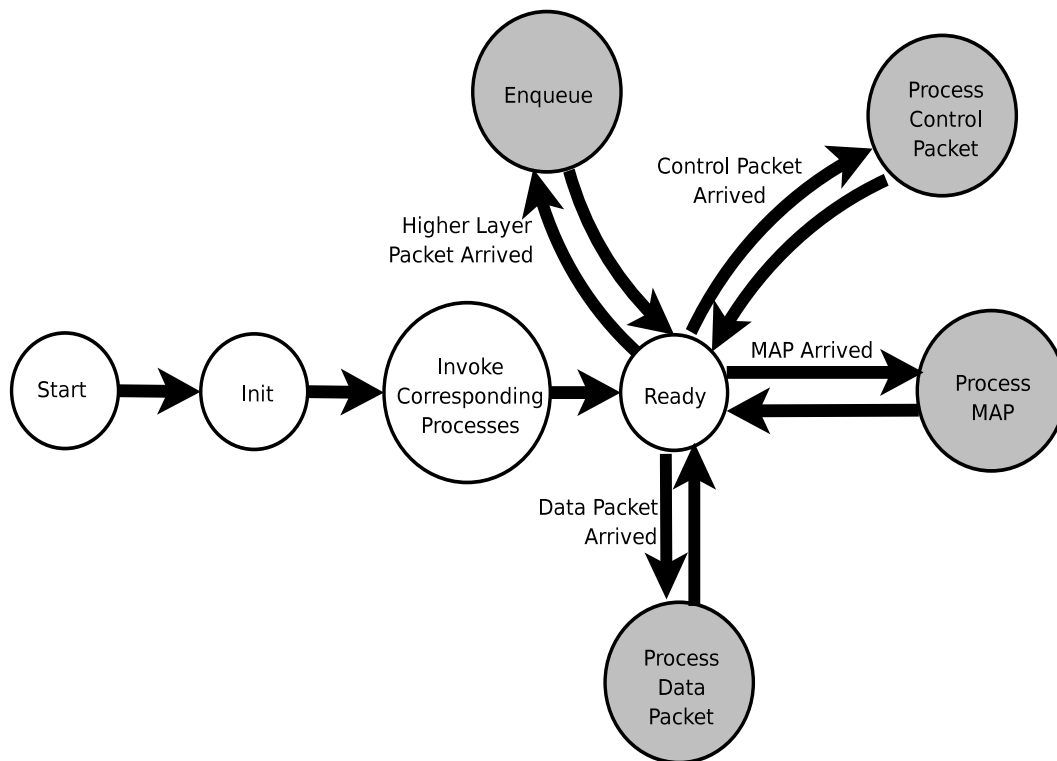


Figure 31: OPNET Model: WiFiRe MAC State Diagram

3.1 MAC State Diagrams

Figure 31 shows the generic state diagram of the MAC. It shows states common to both BS and ST. Different actions are taken in each state based on whether the node is the BS or an ST.

In the *Init State* all variables are initialised. MAC and IP addresses are obtained. The *Invoke Corresponding Process State* invokes a `bs_control()` module, in case the node is the BS. This

state initializes the directional antennas (six in current implementation) and then enters the *Ready State* to wait for events. In case the node is an ST, a `st_control()` module is invoked. This orients the ST antenna towards the BS and then enters the *Ready State* to wait for events.

When the MAC (in the *Ready State*) receives a packet from the higher layer, it makes transition to *Enqueue State*. Here, a `classify_packet()` module maps each incoming higher layer packet to one of the outgoing connections, identified by a CID. A CID has a queue associated with it. A `get_queue()` module is invoked to determine the appropriate queue and the packet is then inserted into it. A packet which does not match any condition specified in classifier is inserted in a default best-effort (BE) connection queue.

When the MAC receives a data packet PDU from the PHY, it makes a transition to the *Process Data Packet State*. Here the PDU is de-multiplexed into appropriate higher layer packets, based on the CID(s). These packets are passed to the higher layer for further processing.

When the MAC receives a control packet PDU from the PHY, it makes a transition to the *Process Control Packet State*. In case of the BS, this corresponds to *Dynamic Service Addition Request* and *Bandwidth Request* messages. In case of the ST, this corresponds to *Dynamic Service Addition Response* messages. When the BS receives a *Bandwidth Request* message, the CID of the flow and the bandwidth requested are noted and the MAC returns to the *Ready State*. When the BS receives a *Dynamic Service Addition* message, it determines the service flow type, assigns a CID and a queue. For UGS flows, an `admission_control()` module is invoked. This takes into account the UGS flow requirements and the number of free slots per frame, to determine whether the flow can be admitted or not. An appropriate *Dynamic Service Response* message is created. This is given to a `schedule_pk()` module and the MAC returns to the *Ready State*. When the ST receives a *Dynamic Service Response* message, it creates a queue for the CID and returns to the *Ready State*. Note: One limitation of the simulation model implemented so far is that an ST needs to send all its *Dynamic Service Addition Request* messages at the start of the simulation itself.

At a ST node, when a Beacon is received from the lower layer (PHY), it enters the *Process MAP State*. A child `st_control` process is invoked and the packet is passed to it. The process checks each element in the UL-MAP. The each element in the UL-MAP contains: (i) CID, (ii) Start Slot (Slot number at which transmission should start) and (iii) Number of Slots (Number of slots allotted for transmission). If the CID in UL-MAP element belongs to that ST node, the number of slots allocated for transmission is checked and an appropriate number of packets are extracted from the queue associated with that CID. A MAC PDU (protocol data unit) is constructed, taking into account the PHY overhead and the PDU is given to a `schedule_pk()` module. The control then returns from the child to the parent `st_control` process which returns to the *Ready State*. This `schedule_pk()` module independently transmits the PDU in the appropriate slot.

The BS enters the *Process MAP State* periodically, at the end of each frame (time interval). Here it invokes the scheduler, which takes into account the admitted UGS flows and pending bandwidth requests, to construct the DL-MAP and UL-MAP for the next frame. It then constructs data PDU(s) for downlink transmission in the next frame, based on the DL-MAP. This is similar to the mechanism followed by the ST.

4 Simulation Setup

A preliminary validation of the WiFiRe OPNET model was carried out using the scenario shown in Figure 32. The scenario consists of one BS, surrounded by 16 ST(s), which are placed randomly around the BS. Each ST has 4 flows registered at the BS: one UGS downlink flow, one UGS uplink flow, one BE downlink flow and one BE uplink flow. The following parameters can be specified while setting up the simulation: (i) Classifier Definition, (ii) Service Class Definition, (iii) Packet inter-arrival time (exponential distribution) and (iii) Packet size (Uniform Distribution).

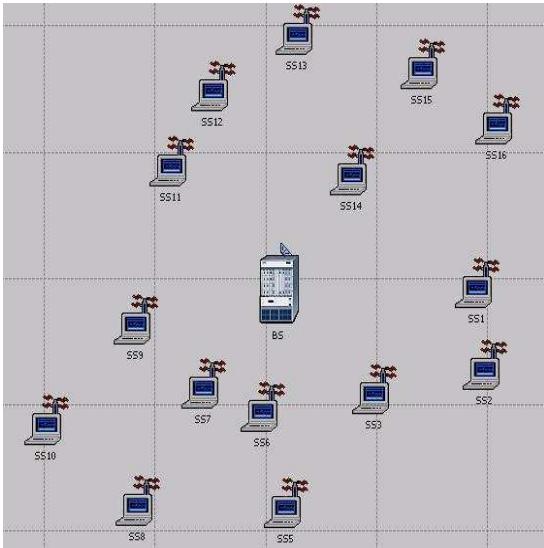


Figure 32: Simulation scenario; unit 5 km.




Figure 33: Throughput (on Y-axis in bps) v/s Load (on X-axis in 10 Mbps)

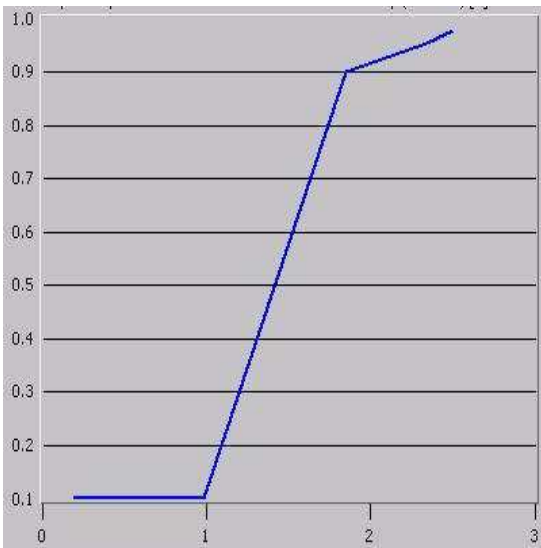


Figure 34: Delay of UGS flows (on Y-axis in seconds) v/s Throughput (on X-axis in 10 Mbps)




Figure 35: Delay of BE flows (on Y-axis in seconds) v/s Throughput (on X-axis in 10 Mbps)

Figure 33 shows how throughput of system varies with increasing load, while Figure 34 and Figure 35 show the delay v/s throughput for UGS and BE flows, respectively. The behaviour observed in graphs is close to that expected by simple theoretical analysis, thereby validating the WiFiRe OPNET model developed so far. Further development of the model as well as detailed experimentation is underway.

References

- [1] Ashok Jhunjhunwala, Anand Kannan and Bhaskar Ramamurthi. Interim Broadband Wireless System for India. CEWiT Technical Report, CEWiT-2005-02, March 2005, available from <http://www.cewit.org.in/>
- [2] Broadband Policy 2004. Govt. Of India, Ministry of Communications and Information Technology, Department of Telecommunications, available from <http://www.dotindia.com/ntp/broadbandpolicy2004.htm>
- [3] J. Klutto Milleth and Bhaskar Ramamurthi. Study of Imminent and Emerging Broadband Wireless Technologies. CEWiT Technical Report, CEWiT-2005-01, March 2005, available from <http://www.cewit.org.in/>
- [4] Digital cellular telecommunications system (Phase 2+) (GSM); GSM Release 1999 Specifications. available from <http://www.etsi.org/>
- [5] Digital cellular telecommunications system (Phase 2+) (GSM); General Packet Radio Service (GPRS) Release 1999 Specifications. available from <http://www.etsi.org/>
- [6] cdma2000 Specifications. available from <http://www.3gpp2.org/>
- [7] iBurst Specifications. available from <http://www.arraycomm.com/docs/iBurst.pdf>
- [8] Flash-OFDM Specifications. available from http://www.flarion.com/products/flash_ofdm.asp
- [9] corDECT, "Wireless access system". Technical report of Midas Communication Technologies Private Ltd. December 2000.
- [10] IEEE Std 802.16-2004. IEEE Standard for Local and metropolitan area networks "Part 16: Air Interface for Fixed Broadband Wireless Access Systems". (Revision of IEEE Std 802.16-2001), 2004.
- [11] IEEE Std 802.11-1999. IEEE Standard for Local area networks "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications". 1999.
- [12] Pravin Bhagwat, Bhaskar Raman and Dheeraj Sanghi. Turning 802.11 Inside-Out. *ACM SIGCOMM Computer Communication Review*, 34(1):33–38, Jan 2004.
- [13] Leena Chandran-Wadia, Shruti Mahajan and Sridhar Iyer. Performance of 802.11 Wireless LANs. IEEE Conference on Computer Communications, Mumbai, 2002.
- [14] <http://www.smartbridges.com/products/ap2np.asp>
- [15] http://www.brovis.com/files/product_as1000.htm
- [16] <http://www.cisco.com/en/US/products/hw/wireless/ps5279/>
- [17] <http://www.winncom.com/moreinfo/item/ETHERANT-3-20/index.html>
- [18] IEEE Std 802.1F-1993. IEEE Standards for Local and Metropolitan Area Networks: Common Definitions and Procedures for IEEE 802 Management Information, 1993.
- [19] Digital enhanced cordless telecommunications (DECT). ETSI EN 300 175, 1988. available from <http://www.etsi.org/>

- [20] Broadband Radio Access Networks (BRAN); Requirements and architectures for broadband fixed radio access networks (HIPERACCESS). ETSI Technical Report TR 101 177 V1.1.1, May 1998.
- [21] Anurag Kumar, D.Manjunath and Joy Kuri. *Communication Networking, An Analytical Approach*. Morgan Kaufman Networking Services, Elsevier, 2004.
- [22] Atilla Eryilmaz, R.Srikant and R.Perkins. Stable Scheduling Policies for Wireless Fading Channels. *IEEE/ACM Transactions on Networking*, 2005.
- [23] Arash Behzad and Izhak Rubin. On the Performance of Graph Based Scheduling Algorithms for Packet Radio Networks. *GLOBECOM*, 2003.
- [24] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 2001.
- [25] Sem Borst, Ken Clarkson, John Graybeal, Harish Viswanathan, and Phil Whiting. User-Level QoS and Traffic Engineering for 3G Wireless 1xEV-DO systems. *Bell Labs Technical Journal*, 8(2):33–47, 2003.
- [26] Harish V. Shetiya. Efficient Routing and Scheduling for IEEE 802.16 Mesh network. Master's thesis, Electrical Communication Engineering, IISc, 2005.
- [27] Leandros Tassiulas and Anthony Ephremides. Stability properties of constrained queuing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control*, 1992.
- [28] Bhaskar Raman and Kameswari Chebrolu. Revisiting MAC Design for an 802.11-based Mesh Network. *Third Workshop on Hot Topics in Networks*, San Diego, Nov 2004.
- [29] E. Brewer, M. Demmer, et. al. The case for Technology in Developing Regions. *IEEE Computer*, June 2005, pp. 25-38.
- [30] M.J. Lee, J.L. Zheng, Y-B. Ko and D.M. Shreshta. Emerging standards for wireless mesh technologies. *IEEE Wireless Communications*, April 2006, pp. 56-63.