WHITE PAPER

CBRS LTE Technology for the Enterprise: Network Implementation and Design

aruba
a Hewlett Packard Enterprise company

TABLE OF CONTENTS

OVERVIEW

The launch of the CBRS band in 2019 and 2020 marks an opportunity for enterprises to build private networks using LTE technology that was previously only available for public cellular networks.

The new band covers 150 MHz of spectrum from 3550 to 3700 MHz. The US regulator has taken a new approach to spectrum sharing with CBRS, allowing unlicensed use but requiring users to be authorized by spectrum access systems to avoid interference with incumbent and high-priority, paid users.

While CBRS spectrum is not allocated to a particular technology, it is a natural fit for LTE (and later 5G), as it requires little modification to existing equipment. Mobile devices and IoT sensors are already available, as is radio and core network infrastructure.

As enterprises and other organizations start to build private LTE networks in the CBRS band to supplement Wi-Fi coverage, networking engineers who are already expert in LAN and WLAN face a steep learning curve to master the complexities of cellular technology in the enterprise.

This series of papers is intended for Wi-Fi engineers wishing to learn about CBRS equipment in enterprise networks. It explains those parts of LTE that are most relevant to CBRS equipment, and enables Wi-Fi engineers to successfully evaluate and operate enterprise CBRS networks.

This is the third and last paper in the series. Earlier papers cover 'The Radio' and 'Signaling & Control'.

INTRODUCTION TO CBRS

A recent development in the wireless landscape, CBRS (Citizens Broadband Radio Service) is a US regulatory initiative for shared spectrum which allows new technology to be introduced into enterprise networks.

The FCC[1] (Federal Communications Commission) laid the groundwork for this spectrum, 150 MHz from 3.55 to 3.70 GHz, from 2015 to 2019, and the first commercial networks were deployed in 2020. CBRS pioneers a three-tier licensing structure with opportunities for enterprises to purchase licenses allowing dedicated access to spectrum, or to operate wireless networks opportunistically in locations or RF channels where other licensees are not active. CBRS is intended to allow enterprises and organizations to set up private base stations for cellphones and other client devices.

High-powered outdoor transmitters can be operated, as well as low-powered indoor small cells, enabling both long-range and in-building communications. The sharing and allocation of spectrum is coordinated by national SAS (Spectrum Access System) servers, which allocate RF channels dynamically in order to avoid interference and protect incumbents. The FCC does not specify a wireless technology for CBRS, but the industry has adopted LTE (Long Term Evolution) standards to allow infrastructure equipment, devices and client modules used in the 4G cellular network to be adapted for use in the CBRS band.

However, traditional LTE infrastructure is typically optimized for the needs of cellular operators fielding very large networks, with skilled operations teams. Operators' requirements for macro network equipment differ from enterprise needs where small, indoor radio units that can be wall-mounted and Ethernet-connected are preferred. Another drawback of traditional LTE infrastructure is the complex cellular core network software for management, control, and data functions. Enterprises do not have the scale of a national cellular operator; they require a small footprint, simpler structures and above all, simpler management interfaces.

Thus, it is unlikely that traditional cellular equipment vendors will find it easy to re-purpose their existing product lines for the emerging enterprise CBRS market. A contrasting approach can be found with startups built on 4G/5G expertise, better able to re-factor complex core software into simpler, cloud software and container-based platforms, and to interface these with small radio units more suited for enterprise use.

This paper concentrates on the technical aspects of LTE technology as it relates to CBRS networks. There is considerable complexity inherent in the standards, but many options in LTE standards do not apply to CBRS and are omitted. The paper also deals only with LTE, or 4G technology; although 5G equipment will eventually be adapted for CBRS, all networks installed through 2021 will be LTE-based.

[1] FCC 3.5 GHz band overview

Technical terms and jargon are unavoidable in LTE. This paper simplifies where possible, sometimes at the expense of precisely accurate terminology. The first LTE terms to learn are 'UE' (User Equipment), the client device on the network, and 'base station' which is sometimes known in the market as CBSD (Citizens Broadband Radio Service Device), eNodeB (an LTE term) or even 'access point'. Other important abbreviations are DL (Downlink, from base station to UE) and UL (Uplink, from UE to base station). LTE standards are developed by the 3GPP [2] (3rd-Generation Partnership Project) standards development organization.

The 3GPP is a standards body for the cellular industry, similar to the role that the IEEE [3] (Institute of Electrical and Electronics Engineers) plays for the enterprise network industry. Ethernet and Wi-Fi are IEEE standards, whereas LTE and 5G are 3GPP standards. 3GPP cellular standards are known as "Releases" and are numbered. For example, the first 4G standard was initially introduced in 3GPP Release 8 in 2008. The 5G standards were initially introduced in Release 15 (2018) and are being extended in Releases 16 (2020) and 17 (due late 2022).

The CBRS Alliance [4] plays a similar role in the enterprise cellular ecosystem as the Wi-Fi Alliance. The Alliance has adopted 'OnGo' as the brand name of their certification for CBRS products.

The Wireless Innovation Forum [5] (WinnForum) produces standards and certifications related to CBRS, including base station and professional installer certifications.

OVERVIEW OF THIS PAPER

This series of papers is organized to give enterprise networking engineers a view of the technology underlying CBRS networks.

This, the third paper, deals with some higher-level concepts. These concepts have parallels in Wi-Fi technology, and where relevant, introductions compare Wi-Fi technology with LTE to help those with a Wi-Fi background. The paper sequences through four broad themes.

First, we consider network performance and QoS. The scheduler has been a key part of 3GPP architecture for some years and is responsible for ensuring service level agreements are achieved. We explain the inputs and outputs of the scheduler, and trace how it provides an important link in end-to-end QoS in a CBRS network.

The second section deals with cell selection and SIM authentication, branching into security and encryption across both wireless and wired links of the CBRS network. The end-to-end security architecture of an LTE network is quite different from Wi-Fi, and its implementation in CBRS networks will be of importance to enterprise networking engineers.

This leads to an explanation of inter-cell handover, where, unlike Wi-Fi, decision making is centered on the base station rather than the UE. Even though CBRS cells cover larger areas and handovers will be less frequent than Wi-Fi, this area will be important for larger networks.

The fourth section moves into link budgets, and some rough models are presented, predicting how a CBRS network can be expected to perform. Regarding multi-cell networks, LTE allows for single-channel and multi-channel plans: while we expect the latter to be typical of CBRS networks, both are discussed, along with the mechanisms LTE provides to mitigate cell to cell interference.

Other topics of interest to network design include base station synchronization and performance for high-velocity clients at highway speeds and above.

Finally, appendices briefly explain how CBRS regulations affect network design and frequency allocation.

[2] 3GPP
[3] IEEE Standards Association
[4] CBRS Alliance
[5] WinnForum CBRS activity

THE SCHEDULER AND QUALITY OF SERVICE

Introduction for WI-FI experts

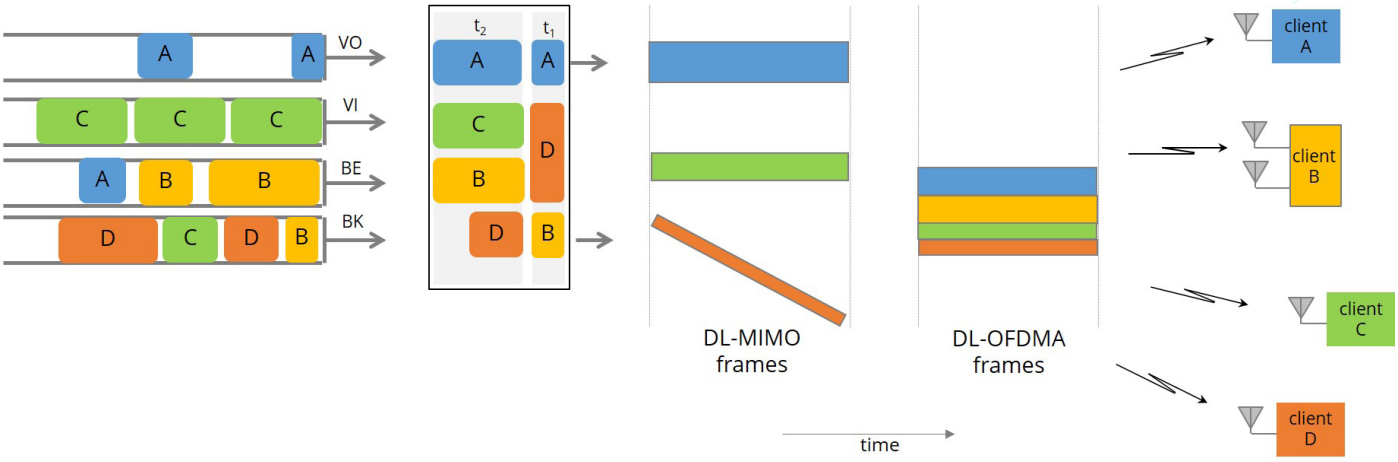

Figure 1. Wi-Fi 6 AP scheduling for downlink multi-user modes

Wi-Fi introduced central scheduling in Wi-Fi 6, allowing the access point's scheduler function to control OFDMA transmissions in both uplink and downlink directions. LTE has used a similar scheduler function for many years, and in many base station implementations it has grown to be quite sophisticated. The scheduler looks ahead by considering traffic buffered for transmission, and, considering priority, latency, error rate and other QoS requirements, MIMO conditions and supportable rates for each client device, schedules transmit slots and modulation rates for the optimum match between offered traffic and available capacity. The LTE architecture differs from Wi-Fi 6 in detail, but the overall effect is the same.

Before Wi-Fi 6, QoS was effected in Wi-Fi networks by mapping incoming layer 2 or 3 priority tags to one of the 4 WMM (Wi-Fi Multimedia) access categories. Each access category was queued separately, and higher-priority traffic was given preferential access to air time in the network, using algorithms for loosely-coupled nodes. With the advent of Wi-Fi 6, the scheduler on the access point is able to exert centralized control over all wireless transmissions, enabling finer-grained SLA (Service Level Agreement) templates per-flow. The mapping to LAN and WAN QoS priorities is streamlined in an all-IP system. LTE has a more complicated end-to-end path, as the data plane is processed by the EPC as well as the base station, but the external interfaces and internal decision architecture are similar.

The LTE scheduler

Earlier papers in this series explained Resource Blocks, how they are allocated to UEs and how the level of modulation and coding affects data rates and error rates. This section explains how this allocation is controlled.

All the information needed to make optimal scheduling decisions, for both DL and UL, is brought to one place: the scheduler function running in the base station. Inputs include administrator configuration, current traffic levels, channel quality indicators, and QoS class configuration.

Administrator configuration indicates the desired QoS or SLA for various classes of traffic (email, voice, video…). This may include limits for throughput, packet loss, error rate, latency and sometimes jitter, as shown below, and a relative priority rating for the traffic type, and the user generating the traffic.

Note that traditional LTE systems can distinguish between voice, data and video as explicitly requested services, but have not implemented separating of email from file transfer or Web browsing, as we expect in modern IP networks and WLANs. This is an area where newer CBRS equipment can adopt enterprise networking techniques for traffic monitoring and classification.

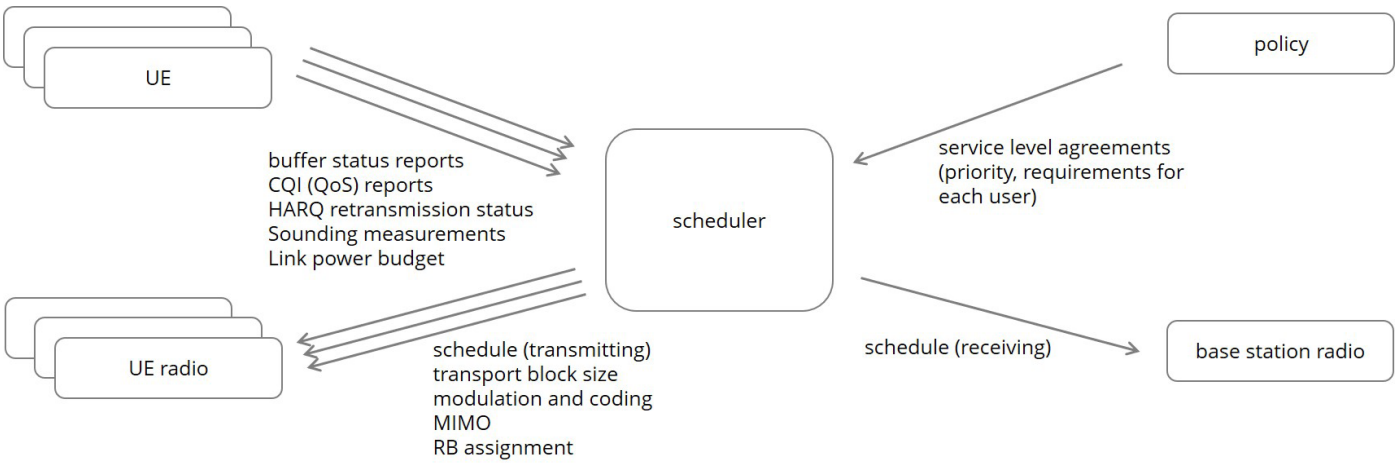

Figure 2. Scheduler inputs and outputs for the uplink

The system is scheduled dynamically; each UE with traffic to send requests service based on its buffered transmit queue state by sending a BSR (Buffer Status Report) on the PUSCH (Physical Uplink Shared Channel). The base station, seeing its own downlink buffers, is able to make its own calculations about offered traffic in that direction.

Traffic allocation is scheduled every TTI (Transmission Time Interval), usually a subframe. Scheduler inputs and outputs are handled differently for uplink and downlink, as it is much easier to get information to and from the base station radio in the same physical unit than UEs across the radio link.

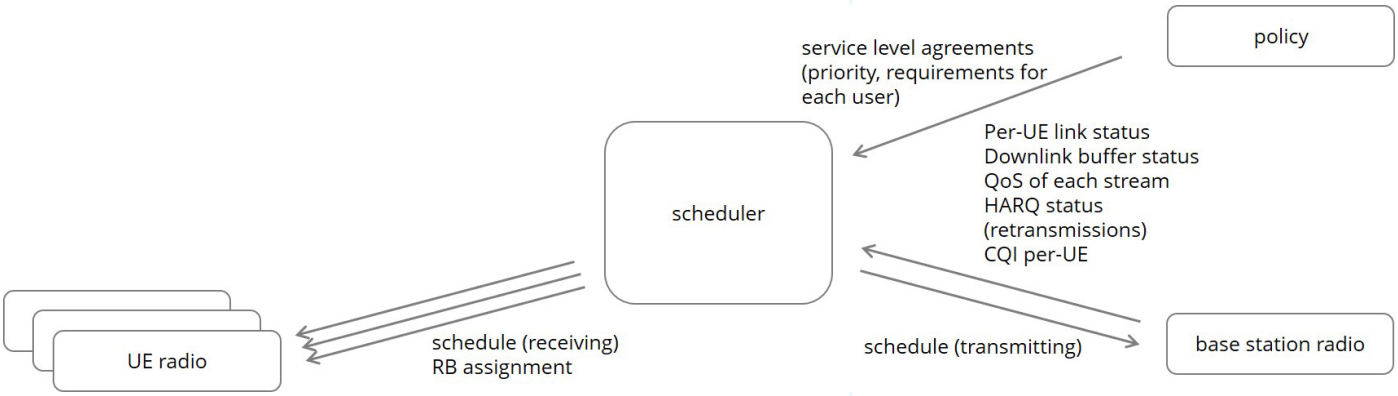

Figure 3. Scheduler inputs and outputs for the downlink

Once the traffic type has been identified and matched against desired QoS characteristics, the scheduler can map its resource grid of available PRBs to DL and UL flows associated with different UEs.

LTE QOS CLASS IDENTIFIERS*					
QCI	**bit rate type**	**Packet error rate**	**Delay budget (msec)**	**QCI priority**	**Example services**
1	guaranteed	10-2	100	2	Bidirectional voice
2	guaranteed	10-3	150	4	Bidirectional video
3	guaranteed	10-3	50	3	Gaming
4	guaranteed	10-6	300	5	Broadcast video
65	guaranteed	10-2	75	0.7	Mission critical PTT
66	guaranteed	10-2	100	2	Non-mission critical PTT
5	non-guaranteed	10-6	100	1	Voice video signaling
6	non-guaranteed	10-6	300	6	Web, email...
7	non-guaranteed	10-3	100	7	Latency-sensitive traffic
8	non-guaranteed	10-6	300	8	Web, email...
9	non-guaranteed	10-6	300	9	Web, email...
69	non-guaranteed	10-6	60	0.5	Mission critical signaling
70	non-guaranteed	10-6	200	5.5	Mission critical data

***others are added for 5G, e.g. V2X, Intelligent Transport Systems, Discrete Automation**

However, the bit rate associated with a symbol or PRB is not constant. As UEs become more distant from the base station, modulation and coding rates must be reduced to meet packet error rate targets, and more PRBs are required for a given flow. The scheduler takes the CQI (Channel Quality Indicator) sent by the UE and calculates the required PRB allocation to sustain these values.

Further, in overload conditions there will be more traffic than PRBs available, and the scheduler determines how to allocate scarce bandwidth resources, and which lower-priority traffic may not be delivered. This is a familiar situation for dynamic traffic loading on constrained channels, but not always an easy calculation.

When the scheduler has optimized the PRB allocation for the next TTI, normally a subframe, it transmits information about the UL to UEs, so they are ready to transmit as required by the schedule. This information is carried on the PDCCH (Physical Downlink Control Channel) and is known as DCI (Downlink Control Information). The scheduler must also send information about the DL traffic, so UEs know when to listen for their data.

The goal is to optimize the use of over-the-air bandwidth, while maintaining the required service levels. There are many options available to enhance schedulers, but in simple terms they match the offered data pattern to the available PRBs, while considering channel conditions and administrator-configured rules. Schedulers are key software components of LTE and CBRS networks.

Over the LAN and WAN QoS

Quality of service is an end-to-end phenomenon. Superior over-the-air performance can be negated by congestion or loss of QoS identifiers elsewhere in the network, for example on the WAN.

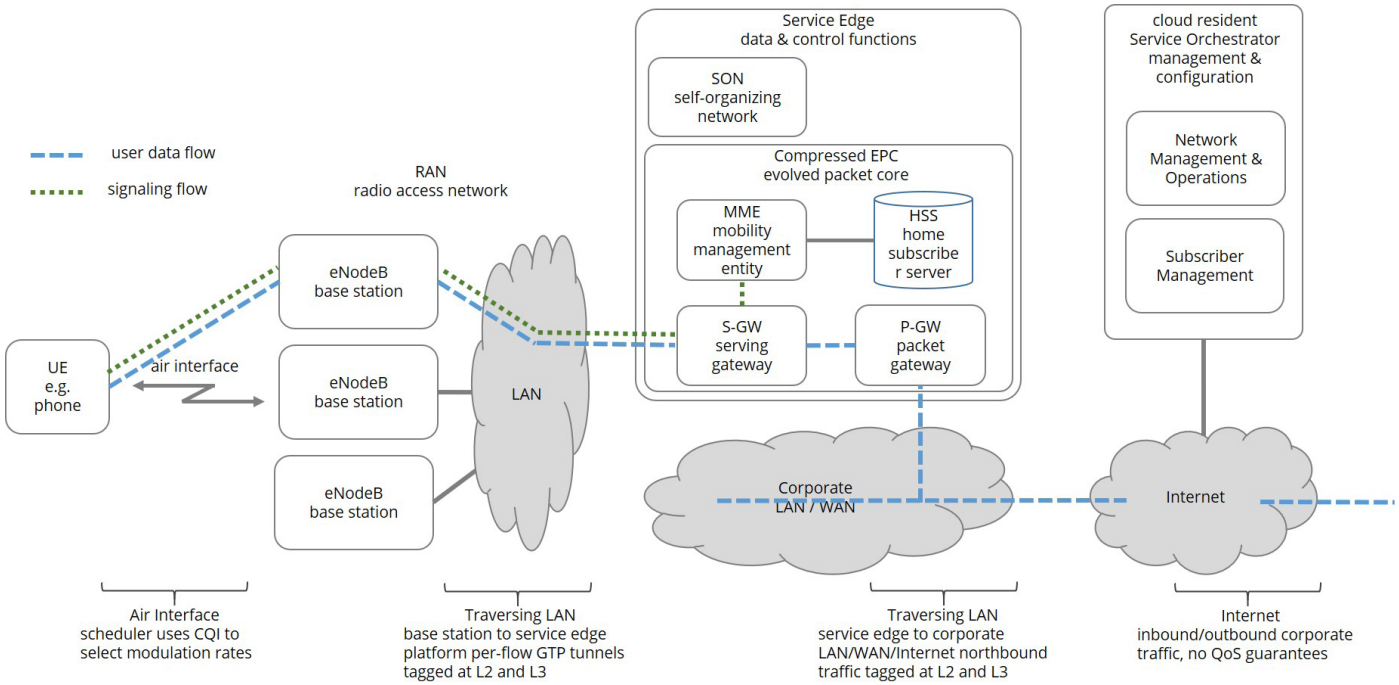

Figure 4. CBRS private network architecture with QoS paths

LTE systems are able to extend QoS priorities by translating internal QoS SLAs to appropriate L2 and L3 tags, ensuring these network segments are able to see the desired priority. Unfortunately, the Internet is generally not QoS-aware, but enterprise LANs and increasingly SD-WANs are capable of supporting good QoS with granular priority levels.

For example, user data between CBRS base stations and the S-GW is carried in GTP (GPRS Tunneling Protocol) tunnels, traversing the enterprise LAN. Each flow has its own tunnel, and the base station and S-GW add appropriate tags to the outside of the tunnels.

CELL SELECTION, AUTHENTICATION AND HANDOVER

This section is concerned with how a UE discovers a suitable network, attaches to a base station, and authenticates, how inter-base station handover is accomplished and the end-to-end encryption architecture.

Introduction for Wi-Fi experts

In Wi-Fi, a network is identified by its SSID, and an individual access point by BSSID. LTE has a similar addressing scheme except that the PLMN-ID is broader; it is globally unique and, in the case of CBRS, a common identifier for all US CBRS networks. CBRS uses the Tracking Area Identifier to identify a network, discussed later in this section. The next level of CBRS identification is the cell ID, identifying an individual base station.

Whereas a Wi-Fi access point advertises its addresses and capabilities in a beacon broadcast as a periodically repeating frame, a CBRS base station establishes its frame structure with synchronization sequences and information blocks. A UE needs to synchronize to decode information about the resource grid of a cell. LTE has no equivalent of the Wi-Fi probe request mechanism; base station discovery is passive.

For authentication, enterprise WLANs use WPA3-enterprise, where the 802.1X framework allows authentication of a client device to an identity store using passwords or certificates and resulting in a handshake where the client and access point negotiate keys for subsequent encryption. The LTE architecture is similar, in that authentication is between the UE and an HSS identity store, based on pre-shared secrets embedded in the SIM card and HSS. Although 3GPP standards now recognize non-SIM authentication, SIMs must be used for authentication in all current CBRS UEs.

Following successful authentication, encryption keys are distributed to the UE, base station and MME element, allowing encryption over the air from UE to base station, and from UE to MME to protect mobility management traffic. The control and data planes in LTE pass through more elements than for Wi-Fi, necessitating a more complex security architecture, although this may be simplified by collocating

or combining elements in a CBRS system.

Inter-base station handover in LTE depends on the same underlying data as for Wi-Fi: measuring signal strength of the current base station while scanning for available handover targets. However, while Wi-Fi clients are responsible for their own decisions on when and where to handover, in LTE control is with the base station – the UE reports measurements, and the base station issues handover commands. (Wi-Fi now allows more centralized handover control while LTE specifications allow some autonomy for UEs, but the models above are generally followed today.)

Network identification, cell acquisition and handover

When a UE is switched on or moves to a new area, it starts a search process by scanning for LTE signals.

First, due to the tightly scheduled nature of the LTE air interface, it needs to synchronize to the signal in the time and frequency dimensions, allowing it to demodulate symbols.

Next, it starts to read the frame structure, in particular the Primary Synchronization Sequence (PSS) and Secondary Synchronization Sequence (SSS). These include cell identity, cyclic prefix length and TDD or FDD mode configuration. These sequences are embedded at known positions in the frame.

After PSS and SSS, the UE can start to decode the Physical Broadcast Channel (PBCH) and read the Master Information Block with more cell information and the configuration of the Physical Control Format Indicator Channel (PCFICH). This allows decoding of the control information channel which includes Secondary Information Blocks (SIB).

The diagram below shows some of the information included in the MIB and various SIBs.

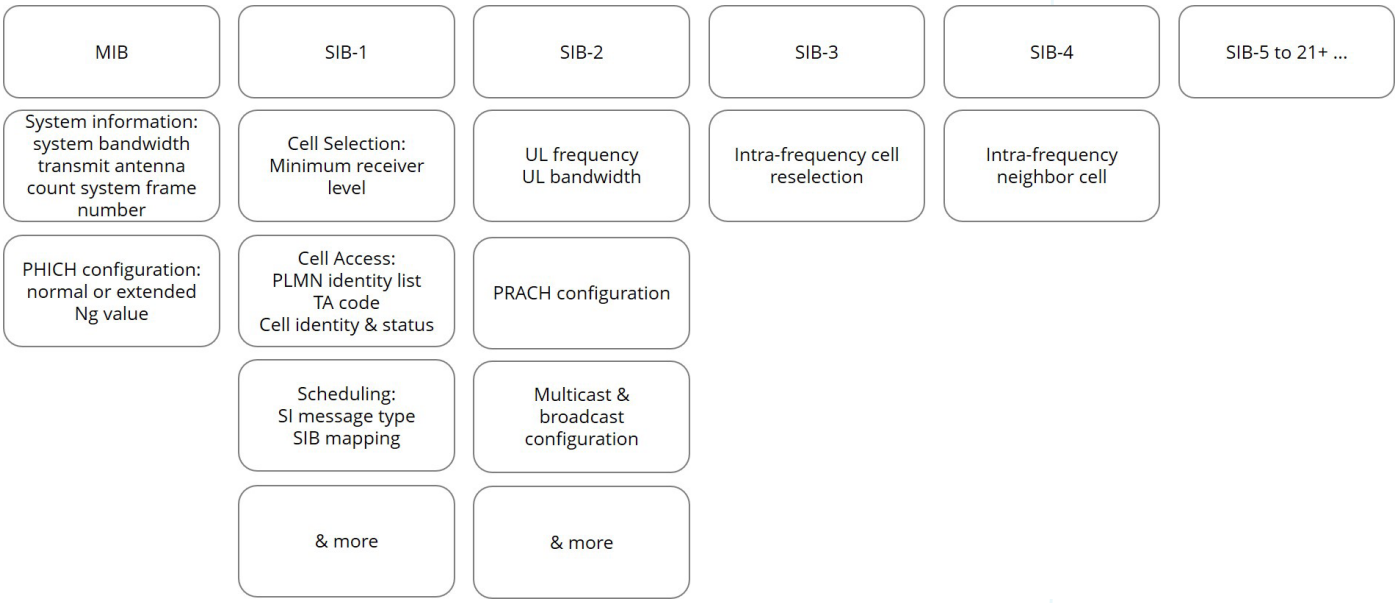

Figure 5. Master and System information blocks

SIB-3 provides information about neighboring base stations on the same RF channel for handover purposes, while SIB-4 covers neighboring base stations on different channels. This will be relevant for most CBRS networks, as spectrum availability will allow multi-channel plans, where adjacent base stations use different frequencies.

When a UE powers up, it must find a suitable base station, attach to it, and authenticate. The first stage is to search for signals across available RF channels, then if necessary, to download enough SIB information to identify the networks present and match them against the various lists stored on the SIM card. Most decisions are made by matching the PLMN (Public Land Mobile Network) identifier and access technology, e.g. PLMN 310-260 with LTE for T-Mobile US, or 315-010 with LTE for CBRS.

The UE recognizes networks it can join by matching base station broadcast values against registers in its SIM card.

Registers are programmed with subscriber identity, authentication credentials and network information.

Home PLMN with access technology. An ordered list of PLMNs : access_technology pairs to match as the Home PLMN	Access Control Class. An index that matches base station Access-Barred lists in SIB-2 (e.g. for emergency call only)
Equivalent Home PLMN. An ordered list of PLMNs to match as the Home PLMN	Location Area Identity. Temporary cache of connected network & location & temporary subscriber ID if roaming to another PLMN
User controlled PLMN with access technology. As above but user-configured	Security Context. This includes the KASME master key, derived from K, is used between the UE and MME
Operator controlled PLMN with access technology. As above but operator-configured	Closed Subscriber Group. A list of groups that should match the SIB-1 broadcast from the base station. (Also optional 'home eNodeB' identifier.)

Figure 6. Selected SIM card (USIM) registers (simplified)

In order to select a suitable network, or the best network if more than one is available, the UE first tries the MRU (Most Recently Used) network held in memory, then runs through a series of steps comparing scanned networks with information in the SIM card.

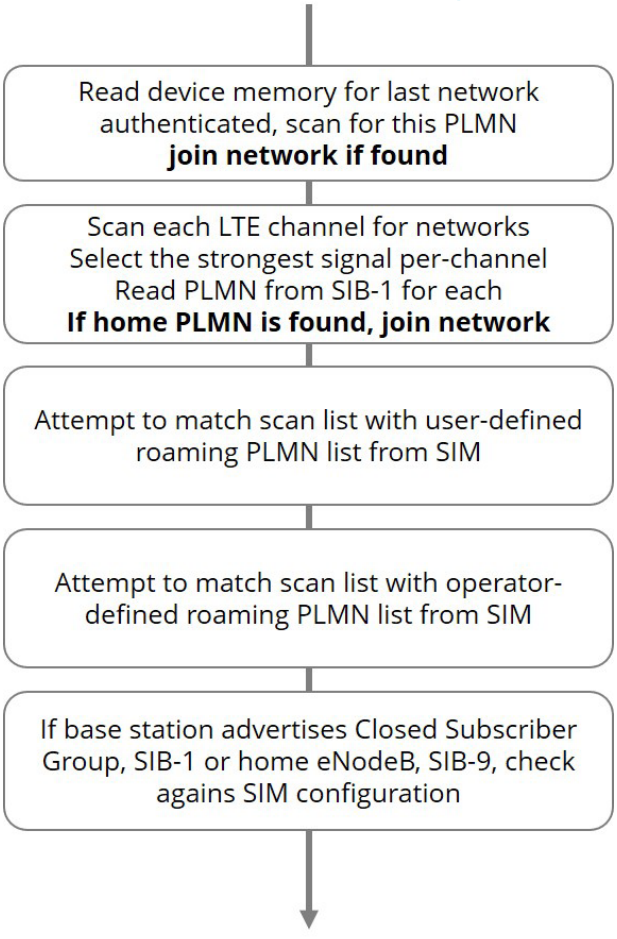

Figure 7. UE power-on network selection (simplified)

Once a suitable PLMN is found, there are additional tests for barring, where a base station protects itself from congestion by accepting only certain types of UE or traffic, and for some enterprise small cells and residential femtocells, a CSG (Closed Subscriber Group) match on group and 'home eNodeB'.

CBRS uses TAIs (Tracking Area Identifiers) to identify particular networks. The TAI concatenates the CBRS-specific MCC-MNC (315-010) and a network-specific, operator-assigned TAC (Tracking Area Code). A UE with a CBRS SIM will attempt to authenticate where it sees the CBRS MCC-MNC, and blacklist TAIs and base stations where its authentication is rejected with a specific reject code. Thus, the UE selects the correct base station even if in an area of overlapping CBRS networks. See the appendix on 'Designated identifiers for CBRS networks' for more information about PLMN and other identifiers specific to CBRS.

SIM authentication and integration with enterprise networks

LTE communications security procedures are rooted in the SIM (Subscriber Identity Module) card, properly known as UICC (Universal Integrated Circuit Card) in LTE and sometimes as USIM (Universal SIM) from 3G. The SIM is a tamper-proof electronic computer on a card that identifies the subscriber to the network and also provides network authentication.

It is an old security maxim that authentication can depend on something you have, something you know or something you are. The SIM card is something you have – other forms of identity, such as username-passwords and X.509 certificates are not supported in LTE (although phones can have access passwords).

The SIM contains a unique subscriber identifier, the IMSI (International Mobile Subscriber Identity) which is a 15-digit number, and a root key, 'K' in LTE terminology, for that subscription. When the SIM card is programmed, K is written into its file system, and also into an authentication server in the operator's network, associated with the HSS (Home Subscriber Server) function. All authentication and subsequent authorization in LTE depend on the IMSI and K values in the UE and HSS.

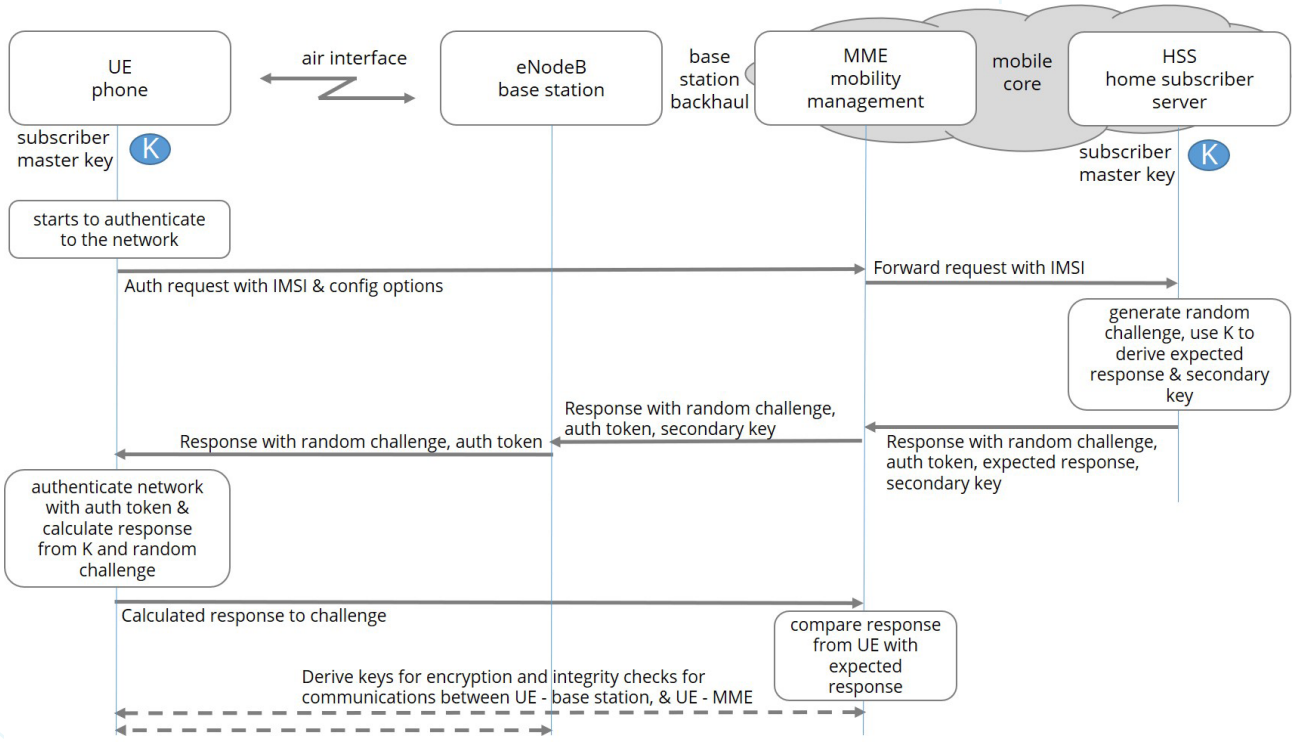

Figure 8. LTE SIM (AKA) authentication sequence (simplified)

The diagram above shows the authentication sequence in simplified form. The important network elements are the HSS, MME (Mobility Management Entity) and base station.

- The UE transmits its IMSI to the MME, which forwards it to the HSS.
- The HSS generates a random number, which it then uses with the K key to calculate an expected response value. It also derives a secondary key for use later, and an authorization token. It returns this set of values, or tuple, to the MME.
- The MME forwards all but the secondary key to the UE.
- The UE compares the authentication token from the HSS with a value it derives locally from the K key. If they match, the UE knows this is a valid network, rather than a rogue, so it proceeds.
- At the UE, the random number from the HSS via the MME is combined with the local K key to calculate a response value, which is returned to the MME.
- The MME compares the UE's response with the HSS's expected response. If they match, authentication succeeds.
- Following authentication, a multi-way exchange results in a number of keys, ultimately derived from the K key, at the UE, the base station and MME. These keys are used to encrypt and integrity-check over-the-air transmissions, signaling and control traffic.

The various keys used by different entities are shown below.

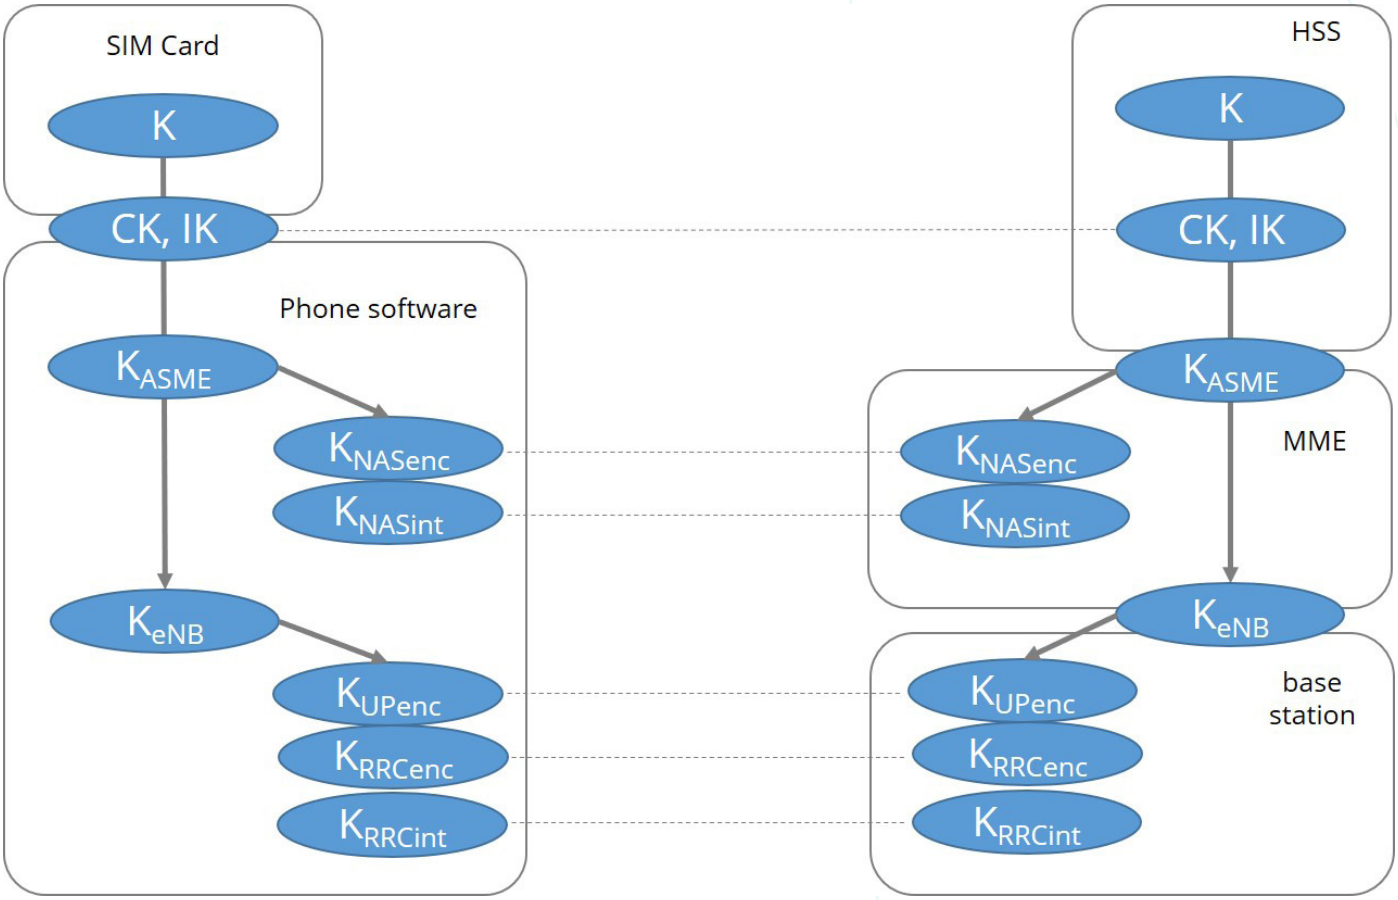

Figure 9. LTE SIM (AKA) key derivation

CBRS equipment makes use of all the authentication functions above. SIM cards must be programmed against the CBRS cloud resident HSS and inserted in UE devices. In practice this means that the HSS owner (or CBRS equipment provider) programs the SIM with IMSI, master key secret and other values, and communicates that information to the HSS to set up the trust relationship between the SIM and network. In private CBRS networks, these SIMs allow authentication only to their parent network – they will not work on public cellular networks, and SIMs from public networks cannot roam onto a private CBRS network using currently-deployed equipment, for lack of commercial agreements for neutral-host arrangements.

LTE and CBRS networks implement an authentication, key distribution and security standard defined by the 3GPP and using AES encryption[6] over the air. There are no widely known snooping attacks; the majority of publicized vulnerabilities[7] involve fake base stations tricking the UE into down-selecting to 2G or 3G algorithms. CBRS networks do not offer the end customer any configuration options for security algorithms.

SIM card vulnerabilities have been widely reported in recent years[8], partly because most two-factor authentication methods for over-the-Internet authentication use a one-time password delivered via SMS, so misusing SIM technology to divert SMS messages to a criminal's device can be lucrative. The most common exploit is not based on the SIM card itself, but SIM-swapping, where a hacker persuades a service provider to switch the target's service to a new phone and SIM card.

It is also possible to clone a SIM card. This normally requires the criminal taking possession of the SIM to be cloned and using a SIM card reader to extract information and program a new SIM. It can take some time, as the master key is not directly accessible and must be discovered by repeated queries, but it is feasible. However, when the new card is used in a network, the HSS will generate alarms if it sees the same IMSI in use on two devices simultaneously. More sophisticated variations of this exploit are possible.

For situations where SIM card physical security may be an issue, the HSS can bind the SIM identity to the physical identity of the UE, normally the IMEI (International Mobile Equipment Identity) of the desired UE.

LTE requires encryption for all data and control plane traffic that may be exposed to interception.

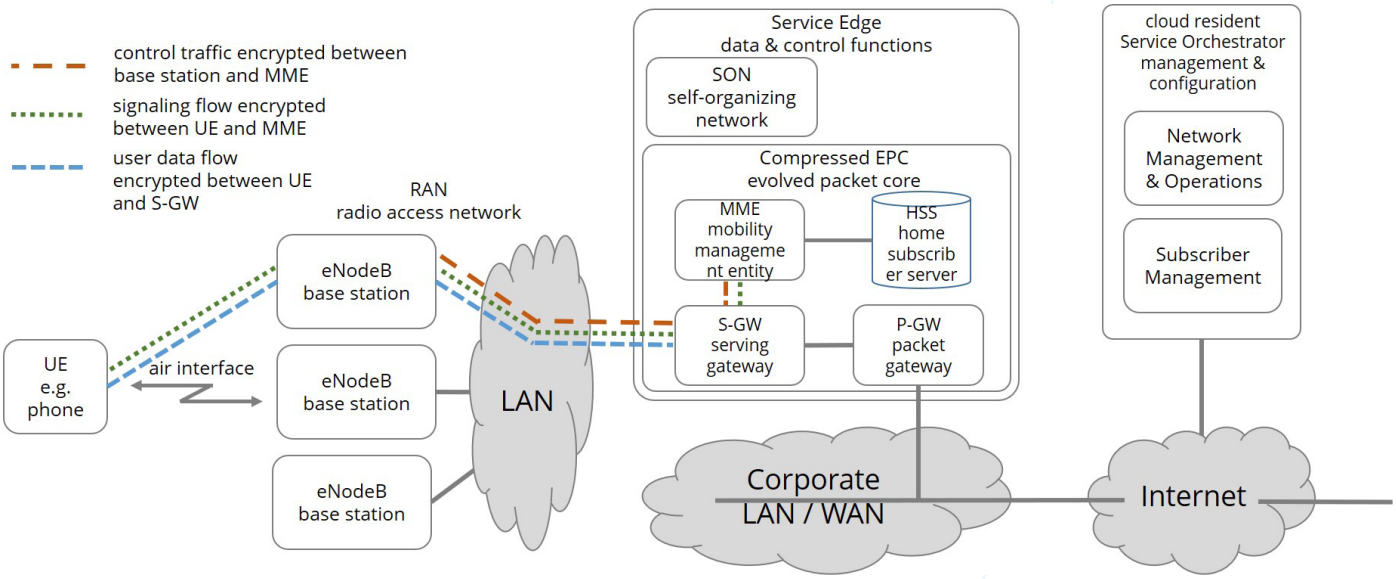

Figure 10. CBRS private network architecture with encryption paths

[6]ITU, 4G and 5G networks security techniques and algorithms

[7]NIST, LTE Security – how good is it?

[8]FCC Cell Phone Fraud consumer guide

Base station traffic is backhauled across the LAN to the CBRS service edge platform. Data, signaling and control traffic between these nodes, and over the air, is encrypted per-flow with keys known only to the endpoints. The diagram above shows encryption paths.

SIM provisioning

For enterprise customers, a CBRS network vendor will normally provide SIM cards pre-programmed with keys matching a virtual HSS that it maintains in the cloud. This means the enterprise will not be required to program SIMs, but may need to contact the CBRS network vendor to activate them.

However, an increasing number of devices now support eSIM functionality, and this allows remote administration, obviating the need for physical SIM cards.

An eSIM, or more properly eUICC, is a tamper-proof memory-processor component built into the device, into which a software SIM profile can be programmed. It is possible for a device to have several programmed eSIM profiles and switch between them, but today's consumer phones more often support one physical SIM slot and one eSIM profile.

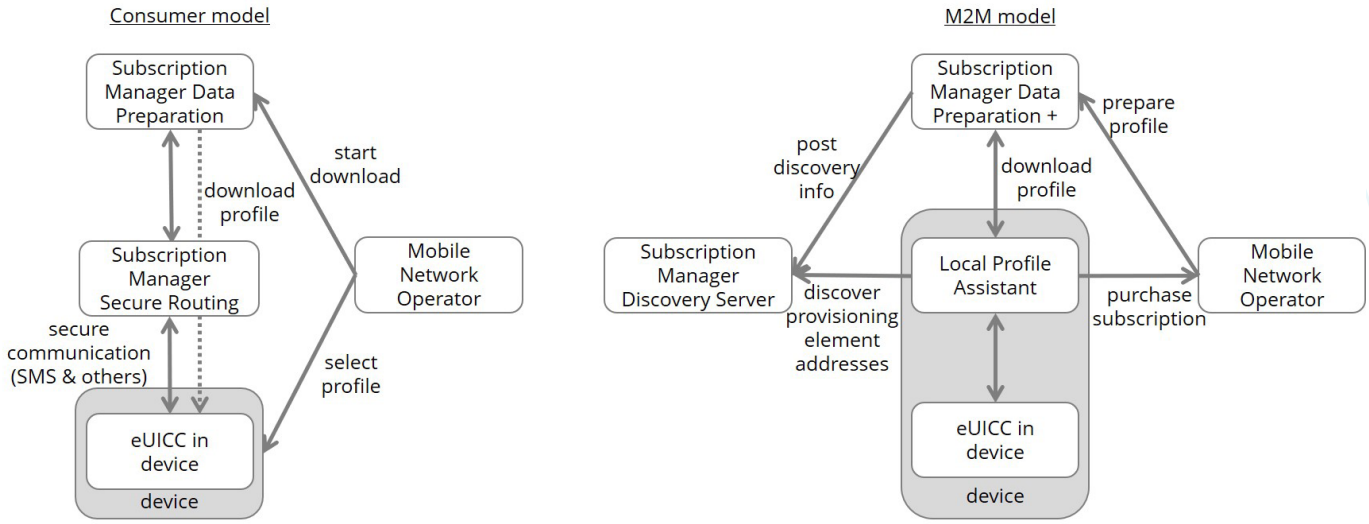

Figure 11. eSIM programming architecture (simplified)

The specifications support two different architectures for programming eSIM profiles[9], 'consumer' and 'M2M'. The M2M solution is defined for headless devices such as IoT sensors; it is a 'push' architecture where control is exercised by the mobile network operator managing the device. In contrast, the consumer solution is a 'pull' architecture for mainstream consumer phones where the user is in control of downloading and activating new eSIM profiles.

INTRA-NETWORK HANDOVER BETWEEN BASE STATIONS

For UEs that are already attached to a RAN that are approaching a cell edge, LTE base station to base station handover is a subset of mobility management and incorporates a multitude of sophisticated and potentially complex options. Many of these options can be ignored in a CBRS network, as its architecture is a lightweight form of the public cellular infrastructure. To take a few examples, inter-base station signaling is handled as inter-process communication in software (the RAN is virtualized); these networks sometimes operate in a single RF channel; and each radio unit corresponds to a single cell and sector. Also, in this paper we ignore handover between CBRS and the public cellular network. But the core properties of LTE mobility apply to inter-base station handovers in CBRS.

The simplest form of handover is when the UE is in the idle state, with no active connections. It is important that as it moves through the network, it registers with base stations as it goes, so the infrastructure knows where to direct incoming connections. The important factors in this case are maintaining a low-power state for good battery life, and long breaks in connectivity are acceptable. Handover under these conditions does not need to be high-performance.

Our focus in this section is on UEs with active connections, where packet loss, latency and data rate fluctuation are key performance parameters. This type of inter-base station handover in LTE is a multi-step process under control of the

[9]GSMA eSIM white paper

network – the UE has no autonomy – with preparatory steps before the move to a new base station, and redirection of internal connections afterwards. The actual move can be very fast, with only a few milliseconds' gap in active data connections.

The first step, a measurement configuration phase, can occur well before a handover is necessary. The serving base station sends messages to the UE defining a set of measurements it should make as it approaches the cell edge.

The messages are quite sophisticated, setting receive signal strength (RSRP) and optionally quality (RSRQ) thresholds that should trigger the UE to start measurements, after a delay. (Current CBRS equipment uses RSRP only.) Messages may also define a high receive signal strength threshold where measurements can stop, should the UE move back towards the center of the cell. As shown below, we expect that as the UE moves towards the cell edge, receive signal strength falls until it reaches the threshold. If it continues to stay below the threshold for a defined time period, the UE will start making measurements of the signal strength of other base stations in the vicinity. It makes these measurements every few seconds, transmitting them to its serving ('old') base station. Thus, the network can monitor the situation of the UE in real-time, and hence make an informed decision about when and where to trigger a handover.

LTE defines a concept of 'measurement gaps' for the case where a UE has to go to another RF channel to find an adjacent base station. A gap is a periodic time interval of a few milliseconds when the old base station commits it will not schedule uplink or downlink traffic for the UE, so it knows that it can go off-channel during a gap and not miss any traffic opportunities. Where channel allocation allows, CBRS networks will set adjacent cells to different frequencies, so this mechanism is likely to be used.

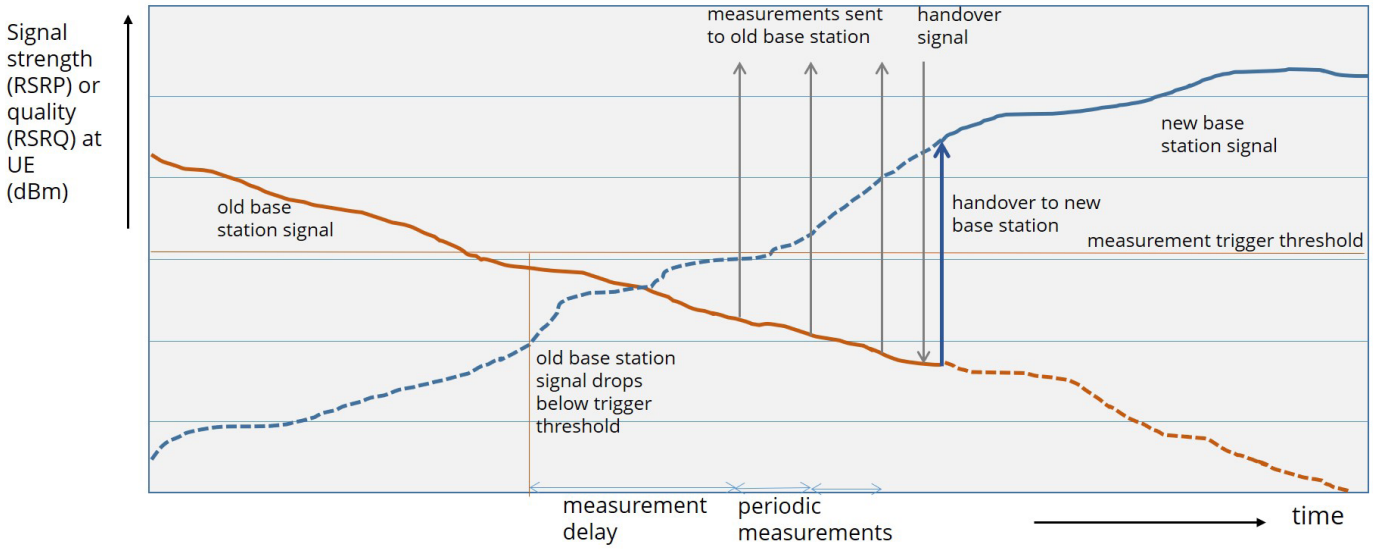

Figure 12. Inter-base station handover

At some point in time, the old base station will receive measurements from the UE that show the adjacent (new) base station's signal is significantly above its own signal at the UE, and it will initiate a handover. This threshold may be configurable and tunable per CBRS network, depending on implementation details; the decision will use RSRP values, and may be more complex than a simple power ratio. The next step is for the old base station to contact the new base station, inform it of the UE's current connection characteristics and request an admission control calculation to ensure the new base station has sufficient capacity to handle the new UE.

Assuming success, the new base station prepares a connection reconfiguration message for the UE informing it of its new control and data channel configurations. When this message arrives at the old base station, it unpacks it and forwards it to the UE.

Now the UE is ready to make the move to the new base station. It detaches from the old one and uses the channel reconfiguration information to connect to the new base station and resume transmissions.

In conventional network equipment with traditional base station architecture, these messages would use a standard ("X2") interface between base stations. In CBRS equipment, with a virtualized RAN, the base station functions run in software containers and the messaging will be inter-process between containers. The diagram below treats base station hardware as transparent to control signals, while all computation and decision-making is abstracted into the cloud or devolved to the Service Edge platform.

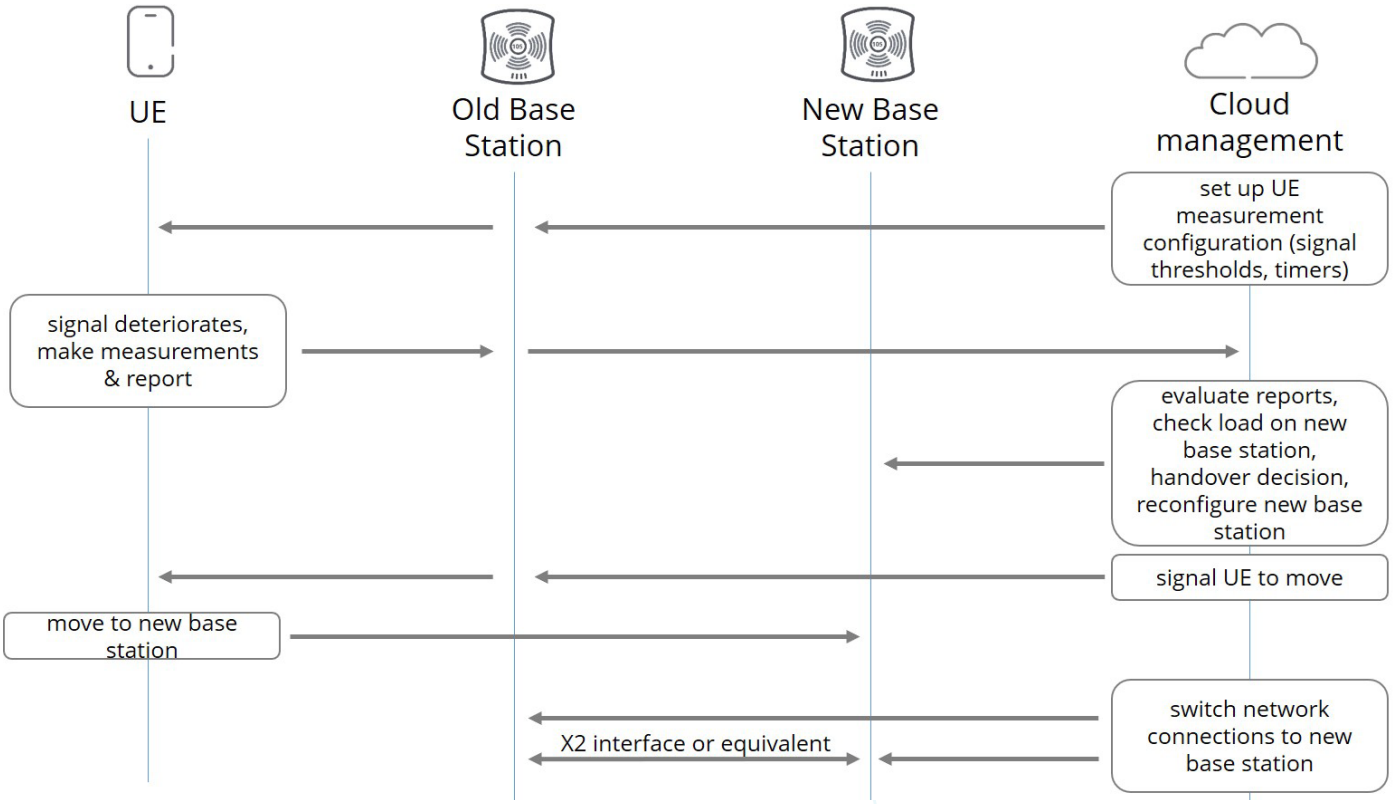

Figure 13. Handover control in the CBRS architecture

Following the UE handover, or simultaneously, the old base station transfers state to the new base station. This would normally include any outstanding downlink or uplink packets from its buffers. This ensures the shortest possible user data interruption.

Finally, the core network may need to run some cleanup operations, re-homing the GTP tunnels for the UE traffic between the serving gateway and the new base station and tearing down the tunnels to the old base station. This involves the S-GW and MME functions in LTE terminology.

Some UEs are able to trigger handover requests autonomously based on channel conditions and base station-configured parameters, but this is not prevalent in CBRS networks.

NETWORK DESIGN CONSIDERATIONS

This section of the paper moves beyond the architecture and protocols used in CBRS networks to more practical information of use to the network designer.

It presents specimen RF link budgets and radio parameters, suggested configuration of multi-base station networks and synchronization requirements.

The section concludes with a discussion of CBRS networking for high-speed clients, and typical measured performance figures.

Introduction for Wi-Fi experts

Engineers with a Wi-Fi background will be familiar with the concepts in this section.

CBRS spectrum, at 3.5 GHz, is between the two established Wi-Fi bands and has similar propagation characteristics. Path loss, wall penetration and multipath behavior is similar.

The range of CBRS connections is generally greater than Wi-Fi, due to two effects. First, transmit power, particularly for outdoor base stations can be higher. Second, the receive sensitivity extends to lower SINR values, chiefly due to lower modulation and hence data rates. A rule of thumb is that indoor CBRS deployments should require ¼ the number of base stations compared to Wi-Fi access points covering the same area.

The higher base station transmit power, and differences in chip design between UEs and base stations drive asymmetric data rates for most connections, where the DL achieves a higher rate than the UL. This differs from WLANs where connections tend to be symmetrical in rate.

Another difference between Wi-Fi and LTE is in RF planning. Whereas Wi-Fi always uses a multi-channel plan, LTE in the public cellular network setting usually has limited bandwidth, and all base stations in a network use a single channel. This can cause considerable mutual interference at cell edge, so LTE has developed some sophisticated mitigation mechanisms relying on synchronization and coordination between adjacent base stations. However, CBRS bandwidth allocation for GAA networks is expected to be sufficient to allow multi-channel topologies, which will resemble Wi-Fi networks. PAL licenses may be only one or two 10 MHz channels, so single-channel plans may be more prevalent if network managers wish to stay within their priority allocation.

Base station synchronization is an issue that does not exist in Wi-Fi but is critical in LTE. The cellular network makes extensive use of GPS receivers for accurate timing at base stations, but indoor CBRS base stations find it difficult to acquire GPS signals and will need network timing sources. Synchronization techniques for CBRS continue to evolve.

The range characteristics of CBRS are particularly significant outdoors, where regulatory EIRP limits are high. This has generated interest in covering large outdoor areas with road traffic, where CBRS can provide connections to vehicles at highway speeds and beyond.

RF link budgets for CBRS systems

Rate-range characteristics of LTE systems are well-established, although CBRS equipment operates in a new frequency band with different propagation characteristics. This section includes general information and expected performance levels from a theoretical viewpoint: network designers should seek practical advice to supplement these predictions.

The LTE link differs from some other radio systems in that it is often asymmetric; downlink rates are generally higher than uplink rates for two reasons: base station transmit power is greater than UE power (especially outdoors), and UEs, being cost constrained and battery powered, traditionally support lower transmit modulation rates than base stations (in CBRS, 64-QAM vs 256-QAM).

The CBRS base station is, like other LTE base stations, limited in transmit power by regulation. For CBRS, two classes of base station are defined.

CBRS TRANSMITTER REGULATIONS, MAXIMUM POWER		
Class	**EIRP (dBm / 10 MHz)**	**Deployment**
end user device	23	in/outdoor
base station A	30	indoor
base station B	47	outdoor

This section will consider indoor and outdoor CBRS networks separately. For indoor deployments, the typical base station may use an omni antenna with 2 dBi gain, for an EIRP of around 26 dBm. Outdoors, a beefy radio unit with directional antennas will be able to radiate the full 47 dBm EIRP if required, but typically less if an omni antenna is used.

The simplified link budget depends on the sum of tx power, tx antenna gain, path loss, rx antenna gain and rx sensitivity. A fade margin is subtracted, resulting in the link margin.

In practice, modern radio technologies adjust the modulation rate, and hence receive sensitivity to fit the offered link margin: if conditions are good, data rates are increased until error rates are barely acceptable, and vice versa, as conditions deteriorate, data rates are decreased to keep the error rate acceptable. Of course, there are limits to both the high and low data rates.

Propagation in the 3.5 GHz CBRS band is similar to the 2.4 and 5 GHz unlicensed bands used for Wi-Fi. The band (3.55 to 3.70 GHz) does not have any special atmospheric absorption peaks but is attenuated by building materials, trees, office furniture and more.

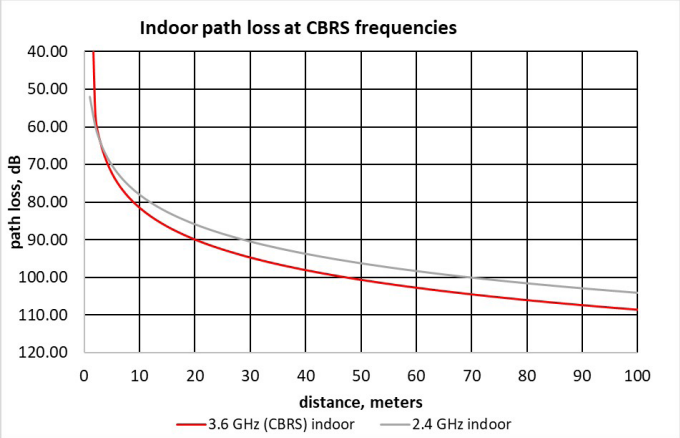

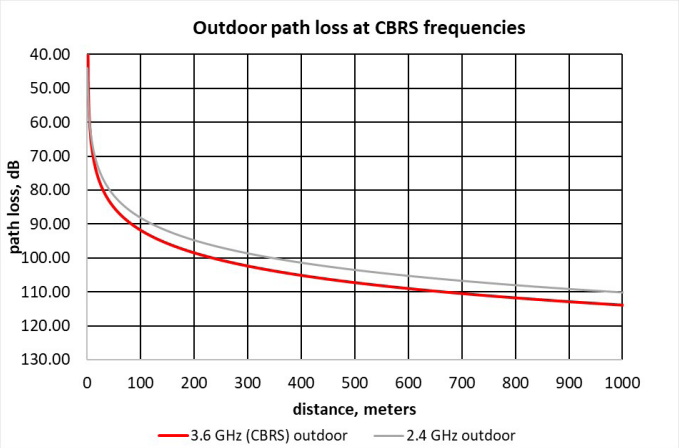

(The charts above use a path loss exponent of 2.6 for indoor and 2.2 for outdoor propagation).

One interesting aspect of LTE systems is that low data rates extend into very low SINR regions. The required SINR for the lowest data rate, QCI 1 (QPSK, 1/8 coding) is negative, at -5 dB and in a CBRS system operating in a 10 MHz channel with 2x2 MIMO, this can support a rate of 2 Mbps (although to do so requires the whole capacity of the cell). A representative link budget is below.

LINK BUDGET FOR CBRS CONNECTIONS OPERATING AT ~2 MBPS IN A 10 MHZ, 2X2 MIMO SYSTEM					
Factor	**unit**	**Indoor DL**	**(omni) UL**	**Outdoor DL**	**(omni) UL**
Conducted tx pwr	dBm	22	23	30	23
Tx antenna gain	dB	2	0	2	0
Rx antenna gain	dB	0	2	0	2
Receiver noise figure	dB	5	5	5	5
Thermal noise floor, 9 MHz	dBm	-104.5	-104.5	-104.5	-104.5
Required SNR for QCI 1, ~2 Mbps	dB	-5	-5	-5	-5
Fade margin	dB	10	10	10	10
Allowed path loss	dB	118.5	119.5	126.5	119.5
Equivalent distance	metres	230	250	3500	1700

Indoor systems are relatively balanced DL/UL, but the outdoor system supports higher rates/ranges due to the base station transmit power.

If a higher user data rate is desired, the range is considerably reduced, as shown below.

LINK BUDGET FOR CBRS CONNECTIONS OPERATING AT ~60 MBPS IN A 10 MHZ, 2X2 MIMO SYSTEM					
Factor	**unit**	**Indoor DL**	**(omni) UL**	**Outdoor DL**	**(omni) UL**
Conducted tx pwr	dBm	22	23	30	23
Tx antenna gain	dB	2	0	2	0
Rx antenna gain	dB	0	2	0	2
Receiver noise figure	dB	5	5	5	5
Thermal noise floor, 9 MHz	dBm	-104.5	-104.5	-104.5	-104.5
Required SNR for QCI 10 ~60 Mbps	dB	11	11	11	11
Fade margin	dB	10	10	10	10
Allowed path loss	dB	102.5	103.5	110.5	103.5
Equivalent distance	metres	55	60	650	310

And for point to point links, the 30 dBm conducted transmit power of the outdoor radio unit can be boosted to the full 47 dBm EIRP limit by high gain antennas, giving an impressive span (the calculation below makes no allowance for obstructions or Fresnel zone encroachment and is better than would be achieved in the real world).

LINK BUDGET FOR CBRS POINT-POINT LINKS, ~60 MBPS IN 10 MHZ		
Factor	**unit**	**Outdoor DL**
Conducted tx pwr	dBm	30
Tx antenna gain	dB	17
Rx antenna gain	dB	17
Receiver noise figure	dB	5
Thermal noise floor, 9 MHz	dBm	-104.5
Required SNR for QCI 10 ~60 Mbps	dB	11
Fade margin	dB	10
Allowed path loss	dB	142.5
Equivalent distance	metres	~ 25 km

The link budgets can be extrapolated to derive rate-range graphs for CBRS networks using omnidirectional antennas.

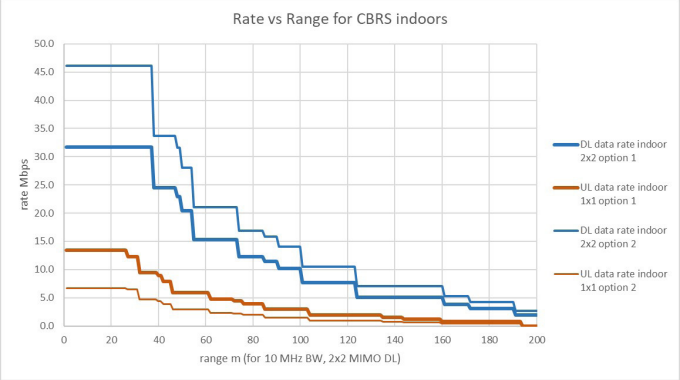

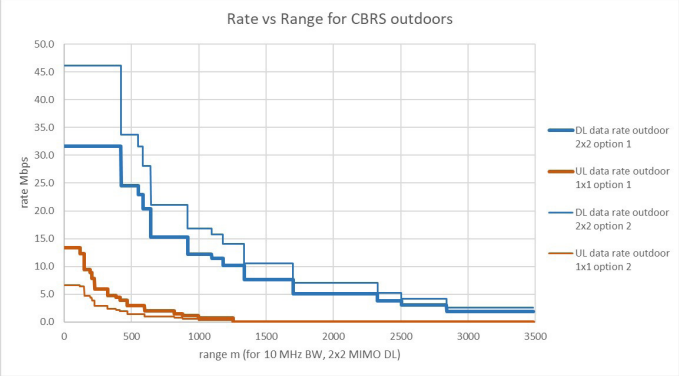

These modeled figures should be taken as indicative, not definitive for any particular planned network. In fact, nearly all practical CBRS networks are designed to capacity limits rather than coverage. This is apparent from the downlink and uplink rates above: these are not only peak per-UE rates, they define the capacity ceiling of the cell. While the indoor radio unit supports 64 simultaneous, active UEs and the outdoor unit 128, these device counts could only be attained at very low average throughput levels.

Consult CBRS equipment and network vendors for specific network design advice. Most offer their own software design tools, backed by more sophisticated modeling.

Multi-cell RF designs

LTE can be deployed in a single-channel or multi-channel frequency plan and includes some sophisticated mechanisms for managing inter-cell interference in the single-frequency configuration.

The CBRS band is expected to allow sufficient bandwidth in most locations that multi-channel plans can be used. However, if conditions require the SAS to restrict the amount of bandwidth available to a particular network, the software managing the CBRS network will adopt a single-channel plan and the mechanisms below will be used to mitigate inter-cell interference.

Traditional re-use 3 frequency plan

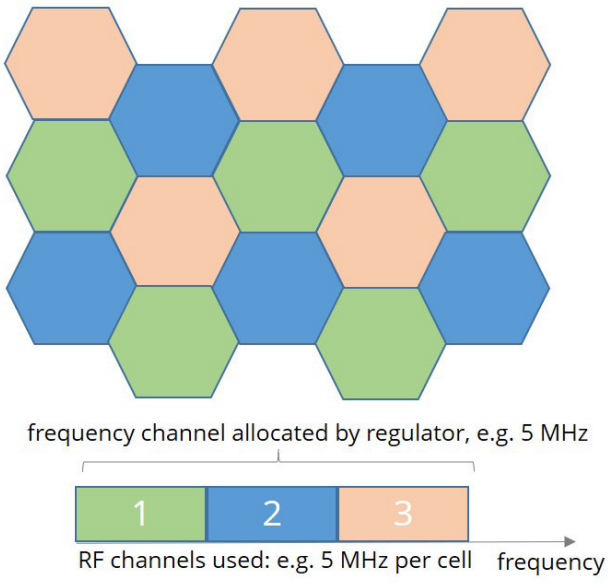

LTE network with re-use factor 1 and ICIC

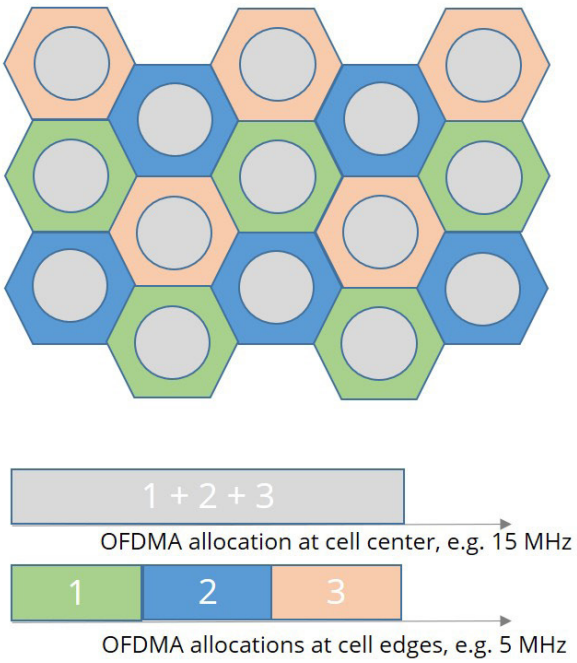

LTE was developed with a frequency re-use factor of 1 in mind. Although it can be used in a multi-channel configuration, most public cellular network designs configure each base station to transmit on the full allocated spectrum.

This works well for UEs near the base station, but towards the cell edge, signals from adjacent base stations will interfere with the desired signal, causing low data rates and high error rates. Serious cell-edge interference is a natural consequence of re-use factor 1 designs.

LTE mitigates this situation by coordinating OFDMA use across adjacent base stations. Scheduler functions agree to use different parts of the channel – different PRBs – for communicating with UEs at the edge. As the diagram above shows, base stations are not spatially aware, but by using a multi-channel plan across the network for cell edge UEs, they can avoid co-channel interference caused by simultaneous transmissions in neighboring cells.

This protocol for frequency-based inter-base station cooperation, a form of FFR (Fractional Frequency Re-use) is known as ICIC (Inter Cell Interference Coordination) in 3GPP Release 8. ICIC was enhanced in Release 10 as eICIC to include both time and frequency-based interference control.

LTE architecture defines the coordination between base stations required for ICIC and eICIC to be over the X2 interface, also used for inter-base station handovers. In CBRS implementations, it is likely to be inter-process communication in the Service Edge platform.

ICIC is not without drawbacks. If cell edge UEs require a large number of PRBs, it may not be possible to avoid frequency overlap. In these cases, cell edge performance can degrade; transmissions from adjacent cells raise the noise floor and lower the SINR for the desired signals.

LTE-TDD uses synchronization of DL and UL transmissions as a secondary mechanism to avoid cell-edge interference. All base stations in a network will use the same TDD option for UL/DL subframe patterns and synchronize their frame start times. This ensures that DL and UL transmissions do not overlap and avoids a cell edge UE suffering interference from an adjacent UE if one were to transmit and the other to receive simultaneously. However, while user data is carried in dynamic PRBs, control and signaling channels have fixed locations in the resource grid and, in a synchronized network, will inevitably coincide in time and frequency and suffer impaired performance at cell edges in a frequency re-use 1 network.

While many network designs involving small cells must consider interaction at the network edge with public cellular networks, CBRS is isolated from the macro environment through use of dedicated spectrum. Interference from neighboring CBRS networks should not be a problem unless they are geographically close, although it is important to remember that SAS systems allocating CBRS spectrum are required to protect incumbents and PAL licensees from lower-tier interference, but not to protect GAA networks from mutual interference.

Time and frequency synchronization across CBRS networks

The previous section showed that for a frequency re-use factor of 1 to be effective, base stations must cooperate in ICIC. This introduces a need for time and frequency synchronization.

Traditional LTE networks based on outdoor radio units make extensive use of GPS receivers to provide accurate time and frequency references, but GPS coverage indoors is not reliable. The alternative method already proven in some areas of the macro network is time synchronization based on IEEE 1588 PTP (Precision Time Protocol), a standard that operates over Ethernet and is capable of accuracy to hundreds of nanoseconds.

Frequency synchronization is required for a number of reasons, including regulatory compliance with spectral masks to limit interference. The requirement for base station frequency accuracy in LTE networks is 16×10^{-9}, translating to 50×10^{-9} over-the-air, which is only achievable with external frequency references such as GPS or PTP.

Time synchronization is especially important in LTE-TDD to maintain accurate DL/UL transitions and timeslot synchronization across base stations, as well as for eICIC functions. The standard requires 1.5 usec accuracy, which translates to less than a cyclic prefix in offset. Some later features, for instance CoMP (Coordinated Multi Point) in r11, require 0.5 usec accuracy.

The use of PTP timing sources in CBRS networks can be complicated and is not yet optimized; synchronization degrades as timing is distributed across network elements, so the placement and management of timing sources affects network performance.

CBRS with high-speed clients

One target scenario for CBRS is vehicular communications, sometimes termed V2I (Vehicle to Infrastructure) or V2X (Vehicle to Everything). It is well-known that vehicle borne UEs can be served by LTE base stations with little service degradation to beyond 150 km/hr, although high-speed trains are able to stress LTE, and aircraft require special treatment. The relevant mechanisms that might degrade service are discussed below.

The underlying mechanism for motion-sensitivity is the Doppler shift, caused by relative movement of transmitter and receiver. This effect is proportional to carrier frequency; the table below lists some relevant values.

DOPPLER SHIFT AT CBRS FREQUENCIES

UE speed		Doppler shift Hz	Coherence time	As fraction of subcarrier spacing
km/hr	m/sec		msec	
50	14	171	2.9	1.1%
100	28	343	1.5	2.3%
150	42	514	1.0	3.4%
200	56	685	0.7	4.6%
250	69	856	0.6	5.7%

Doppler shifts cause subcarrier frequencies to change, causing an offset to the master clock synchronization within the UE's receiver. This can affect the demodulation of signals from adjacent cells, increasing the SINR for the UE. It also affects the base station's uplink demodulation, as the base station sees signals from both moving and stationary UEs with shifted and stationary clocks.

Doppler shift would not otherwise be a significant downlink problem for a UE if it were not associated with multipath and fluctuating multipath effects, but of course this is always the case. The combination gives rise to a number of effects.

When one subcarrier is shifted but adjacent subcarriers are not, signals bleed into each other as subcarrier orthogonality is compromised. This degrades the demodulation process and results in reduced effective SINR. Either error rates increase, or modulation rates decrease, causing lower throughput. The table above shows that, for CBRS signals at highway speeds, Doppler shifts are in the order of 300 – 500 Hz. This is greater than for many LTE systems, as CBRS operates in a relatively high frequency band, but a shift of around 3% will not cause more than a few dB degradation in effective SINR.

Other effects, known as fading, occur as the combination of multipath and Doppler shifts causes fluctuations in signal quality. Fading is a complex topic, sub-categorized as fast-fading and slow-fading. The differentiator is the channel coherence time, approximately the reciprocal of the Doppler frequency shift. The table shows that even at highway speeds, coherence time is in the order of milliseconds and thus much greater than the symbol duration in LTE. Therefore, the operating regime is slow-fading, akin to path shadowing from buildings.

Another consideration when designing for UEs at highway speeds is inter-base station handover. This is due to the small time-window to scan for new cells, identify a base station and execute the handover, during which a vehicle can travel a considerable distance.

Handover scenarios at high-speed require special attention to during network design and commissioning, but generally CBRS will provide good service to vehicle borne UEs. CBRS equipment and network suppliers with experience in designing outdoor systems for mobile clients have rules of thumb for the placement of base stations, antenna configuration and power settings.

Measured performance figures in CBRS systems

Earlier in the paper, a calculation for raw data rate gave theoretical figures for maximum data rate based only on symbols/second and bits/symbol.

Practical figures achieved by test networks are shown below.

PRACTICALLY ATTAINABLE USER DATA RATES IN CBRS SYSTEMS			
Channel width	**TDD frame option**	**Downlink Mbps**	**Uplink Mbps**
10 MHz	1 2	31 23	13 6
20 MHz	1 2	65 94	27 14
40 MHz	1 2	130 189	27 14

Assumes: 2x2 MIMO DL, 1x1 MIMO UL, 64-QAM, single UE

NEXT STEPS

While LTE for the public cellular network is well-understood, applying the technology to enterprise networks with more indoor, micro-cellular focus and a different set of client devices is a relatively new field, and much is yet to be discovered.

Please keep this paper as a reference, and return to it when confronted with questions or problems thrown up by practical issues in your network. If more detail is desired, the bibliography at the end of this paper includes a number of books that we in Aruba have found useful, but beyond this list there are many texts on LTE and 3GPP standards.

In time, CBRS equipment will move from LTE and 4G to the 5G standards, and much of this information will need to be revisited. The functional blocks and protocols may change, but the 3GPP's approach to the problems of moving data across a fundamentally unreliable wireless channel will remain, an interesting counterbalance to the philosophy behind Wi-Fi – even as the two continue to converge.

APPENDIX 1 - LICENSING OF THE CBRS BAND

The CBRS band runs from 3.55 – 3.70 GHz; in LTE terminology it is known as 'band 48'.

LTE FREQUENCY BAND(S) FOR CBRS				
LTE band	**Used for LTE/CBRS in**	**From (MHz)**	**To (MHz)**	**LTE Channels (MHz)**
48	USA	3550	3700	5, 10, 15, 20*
42	Europe, Japan	3400	3600	5, 10, 15, 20
43		3600	3800	5, 10, 15, 20

*** CBRS channels are multiples of 10 MHz**

The band overlaps with two others (band 42, 43) that were designated earlier but not used in the US. All new CBRS equipment is designed for band 48. Channels are identified by beginning and end frequencies, in MHz, e.g. 3600-3610. Channel numbers are not used.

FCC rules govern licensing and sharing of this band[10]. Allocation is in multiples of one to four 10 MHz channels, with 3 tiers of authorized access. Access to specific channel(s) is authorized in real-time by a SAS (Spectrum Access System), a cloud database operated by an independent service provider, certified and overseen by the FCC.

LTE BAND LICENSING TIERS			
			Incumbents
			PAL (Priority Access Licensed)
			GAA (General Authorized Access)
3550 MHz	3600	3650	3700

Spectrum sharing in this new band will allow existing incumbent users and other 'grandfathered' licensed incumbents to be protected, while new installations are allowed where they will not interfere with these protected incumbents. These incumbents occupy the first 'tier' and take precedence over the other two tiers. Tier 1 users are only found in specific locations, typically use only a small portion of the band, and their operations may be temporary or occasional in nature.

In the second tier, private organizations may purchase limited licenses called PALs (Priority Access Licenses) at auction or in the secondary market to operate radios in the lower 70 MHz of this band. Because the minimum geographic area of a PAL is quite large – an entire county – they were primarily purchased by existing MNOs in the initial CBRS spectrum auction in July 2020[11]. PALs are allocated on a county basis, and there are over 3,100 counties in the US. These licensees are endowed with the exclusive right to use the RF channels they purchased after the installation of the base stations operating in the CBRS band, so long as they do not interfere with incumbents in the first tier. A third-tier of priority, GAA (General Authorized Access) allows any other private organization to opportunistically use CBRS channels wherever they will not interfere with incumbents or the PAL users; but they will have to vacate the channel if any higher-priority licensee starts transmitting nearby to particular CBRS access points.

The protection of tier 1 incumbents – generally 5-10 MHz at a time per location – comprising military users (mostly mobile, long-range ship-borne radars), satellite base stations and others, employs regions called dynamic protection areas (DPAs) that cover around 40% of the US population. GAA deployments within 50 – 100 miles of major coastal cities including Boston, New York and Los Angeles are within range of these incumbents. Since incumbents are not required to inform SAS providers of their operations, sensing networks called Environmental Sensing Systems (ESCs) are deployed on the ground to identify and protect usage by tier 1 incumbents in DPAs. Each SAS operator has its own network of ESC sensors.

[10] FCC 3.5 GHz Band Overview

[11] FCC Auction 105, July 23 – August 25, 2020

If activity is detected, both PAL-based and GAA-based CBRS networks nearby may find their spectrum access temporarily switched to another channel by the SAS. Such changes must occur within 5 minutes of detecting tier 1 user activity.

PAL channels in a specific area may also be used by GAA access points in two cases. First, as noted earlier, because the three-tier system operates on a "use it or share it" basis all PAL spectrum is usable for GAA operations unless and until the PAL owner deploys equipment in that location. Second, PAL spectrum can be subleased from its owner. While each PAL covers an entire county, if a PAL licensee only intends to deploy equipment in one location, the rest of the county may be available for subleasing. These transactions are permitted under FCC rules and are known as secondary market licenses. Following the FCC PAL license auction in July 2020, a secondary market in CBRS spectrum allows entities to purchase PAL licenses for the remainder of their 10-year terms.

It is important to note that the SAS only acts to protect incumbents and PAL licensees. When a GAA user requests authorization, it will be granted channels where it will not interfere with these higher-tier users. But the SAS makes no attempt to coordinate channel use between different GAA users, that is left to the users themselves. A GAA user on a particular CBRS channel can be disrupted by a neighboring GAA user who is allowed by the SAS to transmit in the same channel. Common timing synchronization across CBRS networks will play an important part in mitigating this interference.

CBRS has a CPI (Certified Professional Installer) requirement for some installations, specifically the location of all Class B (high power outdoor) radio units must be configured for SAS queries by a certified installer. The CPI requirement also applies if a Class A (low power indoor) radio unit antenna is >6 meters above ground level or is unable to automatically geolocate (i.e. derive location from GPS).

APPENDIX 2 – SAS SIGNALING AND OPERATION

The base station (CBSD in CBRS terms) to SAS interface specification is defined by the WINN Forum[12] and uses JSON over https, with mutually authenticated TLS anchored by a WINN Forum PKI certificate structure, for a defined set of requests and responses. The high-level architecture defines requests originating at the base station or a proxy for several base stations. However, in most CBRS networks, a cloud-resident function operated by the equipment provider will aggregate requests to the SAS and configure base station transmit parameters based on the responses, so apart from configuring base station locations, there will be no direct involvement from the network customer.

In the registration phase, the base station or its proxy sends a number of identifiers to the SAS, along with its radio type, antenna height and location. The SAS acknowledges the registration.

When the base station needs permission to transmit, it first sends a spectrum inquiry request with the frequency ranges it may wish to use. This is, unsurprisingly, to find which channels are available at that time, in that location. The SAS responds with a list of available channels within the ranges requested.

Knowing currently available channels, the base station can now request a grant to transmit on a specific channel, defined by start and end frequencies. The request also includes an EIRP limit, in dBm/MHz. The response is a grant for PAL or GAA confirming the frequency range and EIRP limit, with additional timing values. Each grant has an expiry time: the base station must apply for a new grant before this expires. There is also a heartbeat interval, as the base station must periodically check with the SAS in case a higher-tier user pre-empts the granted frequency range in its location. If the request is not granted, the SAS can provide hints of alternative frequency ranges or EIRP limits that would be acceptable.

The standard includes several options, including measurement reporting where a base station can scan CBRS spectrum and report measurements to the SAS. The SAS determines where operation is likely to interfere with higher-tier users by RF propagation modeling, but augments this with measurements from the ESC sensing network and measurements by CBRS networks, so it can adjust for building attenuation and other real-world effects.

[12]WinnForum WINNF-TS-0016

APPENDIX 3 - ASSIGNED IDENTIFIERS FOR CBRS NETWORKS

The CBRS Alliance has developed several architectural network models for CBRS networks[13]. These include

- Private CBRS networks, not connected to the any public cellular network
- 3GPP Access Mode where the CBRS network operates as a segment of a macro, public cellular network
- Neutral Host Network, operated by a Neutral Host operator that allows roaming of subscribers from multiple public cellular operators
- NHN Access Mode, an architecture based on MulteFire Alliance specifications for neutral host deployments.

This paper describes the private CBRS network model, as initial deployments through 2021 will follow this architecture.

The private network is a standalone network – a closed system - issuing its own SIM cards and providing services only for those SIM cards. Similarly, subscribers to the private CBRS network cannot roam to other LTE networks.

One consequence of this architecture is that private CBRS networks must be allocated certain unique identifiers allowing SIM cards to identify their issuing network. These identifiers must be coordinated to ensure uniqueness[14].

IDENTIFIERS USED IN PRIVATE CBRS NETWORKS

Identifier*		Issuer		Identifies	Example
	IMSI Admin	CBRS-Alliance	CBRS vendor or operator		
IMSI	MCC+MNC (SHNI) + IBN		UIN	Subscription	315010002412345
CBRS-NID	-	CBRS-NID	-	Network	315010
GUMMEI	SHNI	MMEGI	MMEC	MME	
ECGI	SHNI	Macro eNodeB ID	Cell Identity	Cell or sector (eNodeB)	
TAI/TAC	SHNI	-	TAC	Tracking area	

***CBRS Alliance Identifier Administration Guidelines for Shared HNI**

The first important identifier is the MCC-MNC (Mobile Country Code – Mobile Network Code). All US CBRS networks will use 315-010. MCC-MNC is an HNI (Home Network Identifier), referred to by the CBRS Alliance as a SHNI (Shared Home Network Identifier) to accommodate neutral host networks.

The PLMN-ID used for CBRS networks consists of the MCC-MNC, also 315-010. Therefore, this value cannot be used by UEs to determine if they can join a particular network.

The CBRS-NID (Network ID) identifies a single CBRS network. The 27-bit number is assigned by the CBRS Alliance and is used as an LTE CSG (Closed Subscriber Group) identifier. Each CBRS network or enterprise has a globally unique CBRS-NID assigned by the CBRS Alliance.

ASSIGNMENT OF SUB-FIELDS WITHIN CBRS IMSI

Field	3-digits	3-digits	4-digits	5-digits
Data	MCC (Mobile Country Code)	MNC (Mobile Network Code)	IBN (IMSI Block Number)	UIN (User Identification Number)
Value	315	010	vendor-operator	available

***https://www.atis.org/01_committ_forums/ioc/Docs/IMSI-CBRS-Guidelines.pdf**

Example: 315010002312345

[13]CBRS Network Services Stage 2 and 3 Specification CBRSA-TS-1002

[14]CBRS Alliance Identifier Administration Guidelines for Shared HNI

Subscriber or SIM Identifiers, IMSIs[15][16], concatenate the MCC-MNC, a CBRS Alliance registered identifier (IBN) for the equipment or network vendor and a field to identify the particular subscriber. The IBN effectively allows a vendor to reserve a block of 100,000 IMSIs.

In LTE networks, the MME (Mobility Management Entity) must be addressed by the UE, HSS and other entities. It is identified by the GUMMEI (Globally Unique MME Identifier) which is a concatenation of PLMN-ID/CBRS HNI, MMEGI assigned by the CBRS Alliance, and an MME code assigned by the network vendor.

A cell or base station is identified by the ECGI (EUTRAN Cell Global Identifier). This combines the PLMN-ID/CBRS HNI, a macro eNodeB identifier assigned by the CBRS Alliance and a CID (Cell ID) from 0 to 503 which is directly related to the LTE frame structure and is unique within the network. The full ECGI uniquely identifies any cell anywhere in the world.

When UEs are inactive or sleeping, they can move around the network. The network must know where to route incoming or downlink traffic, and to avoid the UE having to register with each base station it passes, LTE defines tracking areas. Each tracking area encompasses a number of base stations and is addressed by a TAI (Tracking Area Identifier). The TAI incorporates the PLMN-ID/CBRS HNI, appending a network-specific TAC (Tracking Area Code) assigned by the network vendor.

Since CBRS networks all use the same PLMN-ID, 315-010, they must add secondary identifiers for uniqueness. UEs recognize the TAI as identifying a particular CBRS network.

While the complexity of the various identifiers and allocation authorities above is daunting, in commercial deployments the CBRS network vendor will take care of all these transactions and configure the CBRS network accordingly.

[15]Iconectiv, CBRS Assignments

[16]ATIS International Mobile Subscriber Identity (IMSI) assignment and management guidelines for shared HNI for CBRS range

ASSIGNED CBRS-I NETWORK IDS (OCTOBER 2020)*		
HNI (Home Network Identity)	**IBN (IMSI Block Number)**	**Entity**
315-010	0000	Unassignable
315-010	0001	ExteNet Systems, Inc.
315-010	0002	Motorola Solutions, Inc.
315-010	0003	Spectrum Wireless Holdings, LLC
315-010	0004	Arvig Enterprises, Inc.
315-010	0005	QuayChain, Inc.
315-010	0006	Landmark Dividend, LLC "DBA" Landmark FlexGrid
315-010	0007	Ruckus Wireless, Inc.
315-010	0008	Velocity Wireless LLC
315-010	0009	ATC,WiFi,LLC
315-010	0010	Boingo Wireless Inc
315-010	0011	JCI US INC
315-010	0012	Ballast LLC
315-010	0013	Welink Communications
315-010	0014	Athonet USA Inc
315-010	0015	Smart Edge
315-010	0016	Windstream
315-010	0017	Frontier Communications Corp
315-010	0018	Mobilite Management
315-010	0019	SportsMedia
315-010	0020	Sony Corporation
315-010	0021	Mercury Wireless, Inc.
315-010	0022	Switch Ltd.
315-010	0023	Celona Inc
315-010	0024	AMG Technology Investment Group, LLC dba Nextlinnk Internet
315-010	0025	Newport Utilities/NU Connect
315-010	0026	RF Connect LLC
315-010	0027	Hargray Communication
315-010	0028	Hudson Valley Wireless
315-010	0029	Xtreme Enterprises
315-010	0030	GenX Communications
315-010	0031	Castleberry Independent School District
315-010	0032	Federated Wireless
315-010	0033	Data Enterprise Systems
315-010	0034	Terranet Communications, LLC
315-010	0035	InfoLink
315-010	0036	Utah Education and Telehealth Network (UETN)
315-010	9999	Reserved for Test

***https://imsiadmin.com/cbrs-assignments**

APPENDIX 4 – BIBLIOGRAPHY

For the reader that wishes to learn more about LTE, CBRS and 5G, here is a list of technical textbooks on various aspects of 4G and 5G systems.

An Introduction to LTE: LTE, LTE-Advanced, SAE and 4G Mobile Communications, Christopher Cox, 360pp

System wide overview of Evolved Packet System from the core to the air interface. In particular, Cox explains LTE MAC layer, framing, error recovery and channelization.

Indoor Radio Planning: A Practical Guide for 2G, 3G and 4G, Morten Tolstrup, 560pp

This is one of the only dedicated textbooks on RF design & capacity planning for indoor small cells. It complements the other two books which take an end-to-end system approach but do not really consider layer 1 design.

5G NR: Architecture, Technology, Implementation, and Operation of 3GPP New Radio Standards, Sassan Ahmadi, 900pp

Reviews all aspects of the 5G system from core to air interface, including virtualization and disaggregation from core to RAN.

Three-Tier Shared Spectrum, Shared Infrastructure, and a Path to 5G, Preston Marshall

This provides a complete historical view of the origins, history and policy objectives of the three-tier architecture that became CBRS. It also considers other spectrum sharing models, including the 2-tier LSA model that exists in Europe.

Femtocells: Technologies and Deployment, Jie Zhang

An overview of small cell technology.