Optimizing and Protecting the Value of Unlicensed Spectrum

A Technical Paper prepared for SCTE/ISBE

Narayan Menon, CTO/EVP of Engineering, XCellAir

Angelo Cuffaro, Senior Director R&D, XCellAir

Jean-Louis Gauvreau, Director R&D, XCellAir

Todd Mersch, EVP of Sales & Marketing, XCellAir

Table of Contents

List of Figures

Introduction

The primary wireless access technology at the disposal of cable operators today is Wi-Fi. The ubiquitous proliferation, low device cost and lack of spectrum license has made Wi-Fi the clear choice. However, as cable wireless services transition from best effort to managed broadband, the ability to deliver reliable high-performance access using unlicensed spectrum is critical. This paper describes the challenges, techniques and results obtained from intelligently optimizing unlicensed spectrum. It also discusses ways to enable harmonious coexistence of Wi-Fi with new technologies such as LTE License Assisted Access (LTE-LAA), LTE – Wi-Fi Aggregation (LWA) and MuLTEfire (standalone LTE), exploiting techniques being leveraged in today's Wi-Fi optimization landscape.

Unlicensed Spectrum – the Congestion Problem

Cable operators have been aggressively expanding their Wi-Fi networks over the last couple of years. This has been driven by goals to expand the customer base, reduce churn and deliver a variety of value-adding services with a high-quality user experience. These could include managed voice, video, gaming and IoT applications.

As Wi-Fi networks expand and densify, radio resource optimization becomes critical to delivering a high quality of experience. Wi-Fi works on a Listen Before Talk (LBT) basis, where devices contend politely for access to the medium (channels). If a channel is in use by a device, other devices wait for the channel to become free before accessing it.

In a dense Wi-Fi network, the large number of contending devices and access points, and a high level of user activity, can combine to cause heavy contention, resulting in congestion. This is true even in a really "well-behaved" Wi-Fi cluster, where devices and access points can all see one another, and defer to each other gracefully when contentions occur. Excessive medium contention results in long wait times for packet transmission opportunities, high latencies and low throughputs.

This issue is likely to be exacerbated when other technologies, such as LTE, start using unlicensed spectrum. LTE-LAA devices will increase contention levels on Wi-Fi channels. Wi-Fi devices will have to also compete with LTE-LAA terminals and access points, and LTE-LAA devices will be contending with Wi-Fi and other LTE-LAA endpoints.

Addressing congestion and interference issues requires the use of radio resource management techniques. Techniques for intelligent allocation of Wi-Fi channels, steering of client devices to a less congested Wi-Fi band, power control and device mobility between access points in a multi-AP environment become critical in making Wi-Fi work well despite the increased contention from Wi-Fi and other technologies using unlicensed spectrum.

Automated management and optimization solutions must include powerful radio resource management (RRM) tools to dynamically provision and tune radio resources. Running RRM algorithms on a cloud server, and interacting with the AP to set and change channels, bands, power levels and other parameters provides scalability, maintainability and a centralized view to deliver a network level solution. The RRM schemes must work proactively to avoid congestion and interference, and steer client devices between

APs to improve coverage. The end results are the optimal usage of available Wi-Fi capacity, dramatically improved latency/ jitter/ throughput performance, and a significantly enhanced quality of experience.

Environmental Observations

XCellAir carried out a study to characterize channel usage in a real-life Wi-Fi environment. The observations were done in downtown Montreal, a busy urban environment with multi-storied office buildings and dense Wi-Fi deployment. At any given time, anywhere from 100 – 250 access points (APs) are visible in the immediate vicinity of a Wi-Fi network.

The objective of the study was to observe Wi-Fi channel availability, i.e. to see how much bandwidth headroom was available at a given time in a typically busy deployment scenario. The study gathered per-channel utilization levels over several periods of time. This was done for both 2.4 and 5GHz bands.

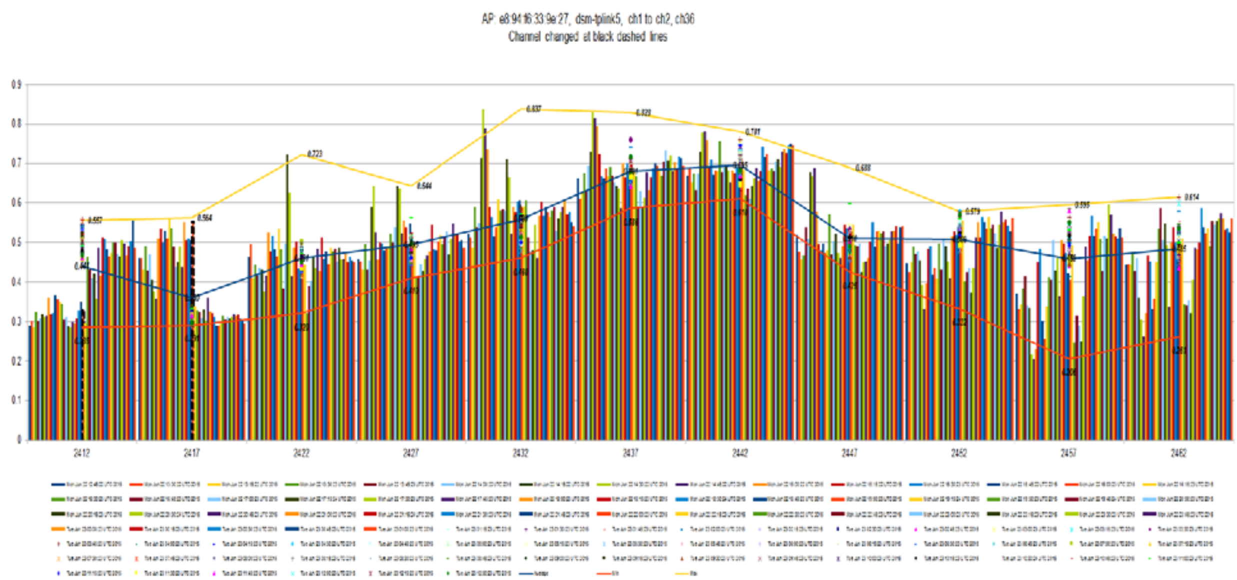

Figure 1 - Per-Channel Utilization Levels in 2.4GHz Band

Figure 1 shows the channel utilization pattern observed for the 2.4 GHz channel over a period of several hours. On the x-axis are the 2.4 GHz channels – 1 to 11. Channel utilization levels (0 – 0.9 or 90%) are shown on the y-axis. Channel utilization reflects the degree (in percentage terms) to which a channel has been occupied by Wi-Fi devices over a measurement period. Each vertical spike on the graph reflects utilization data for the given channel over a 15-minute time period. The horizontal trend lines running through the graph indicate the minimum (red line), average (blue) and maximum (yellow) utilization levels per channel.

Some interesting conclusions can be made from Figure 1:

- Spikes in channel occupancy are visible on several channels at different points in time. These reflect inflection points at which contention and congestion levels increase, and service quality starts to degrade, e.g. high packet error rates, high latencies and jitter, low throughputs etc.
- However, based on average utilization levels, there is bandwidth headroom available at any given time – more on some channels than others. Not all channels typically experience high occupancy at the same time.
- Considering an allocation pool of even five of the channels (e.g. channels 1, 4, 6, 8 and 11), up to two channels' worth of aggregated bandwidth is available on average (considering a maximum channel occupancy of 80 – 85%).
- A dynamic channel allocation algorithm can unlock this bandwidth by moving access points from heavily loaded channels to less occupied ones.

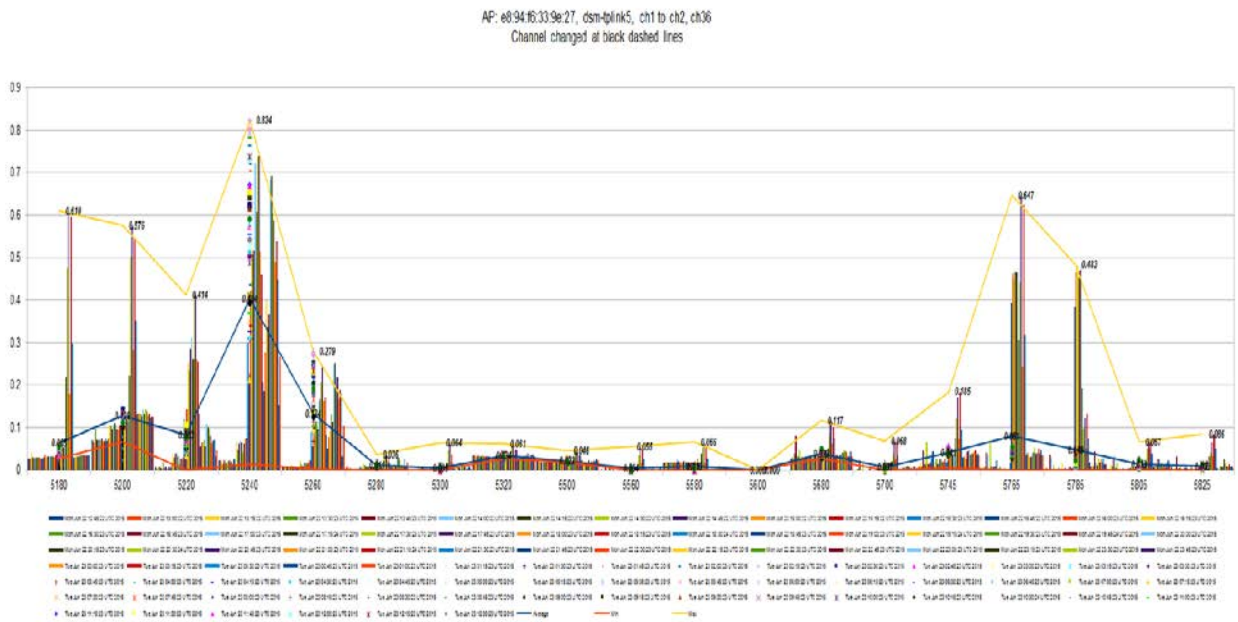

Figure 2 - Per-Channel Occupancy in 5GHz Band

Figure 2 shows a similar channel occupancy view observed for the 5 GHz band. In general, this band is less loaded today than the 2.4 GHz band. A band steering feature (i.e. the ability to move dual-band capable devices on a tightly loaded AP from the 2.4 GHz channel to a less loaded 5GHz channel) takes advantage of this loading imbalance by moving active devices from 2.4 to 5 GHz to leverage the clearer channels and higher bandwidth in the 5 GHz band.

A separate study conducted by XCellAir looked at access point products deployed in the Montreal downtown neighborhood, and observed channel change behavior exhibited by the APs. The conclusion was interesting – less than 8% of the surveyed APs changed channels, once they powered up. Without RRM capabilities, APs stay rooted to a channel regardless of how high the congestion is, and are unable to leverage headroom available on other channels.

RRM Benchmarking & Impact on Quality of Experience

This section discusses a representative sample of test results highlighting impacts of congestion / interference on service quality, and the performance-enhancing influence of two categories of optimization tools – dynamic channel management and band steering. While these results are for Wi-Fi, they are relevant to multi-technology coexistence scenarios in the unlicensed band, e.g. LTE-LAA coexisting with Wi-Fi.

1. Dynamic Channel Management

Dynamic channel management algorithms can mitigate Wi-Fi channel congestion and interference. They detect developing congestion; at a configurable trigger point, if the algorithms decide to change the AP's operating channel, they select a cleaner (less congested) channel and switch the AP and its clients over to the selected channel. The goal here is to maintain good service quality, by not allowing key service quality KPIs to degrade to poor levels as a result of congestion and interference.

Test Setup & Methodology

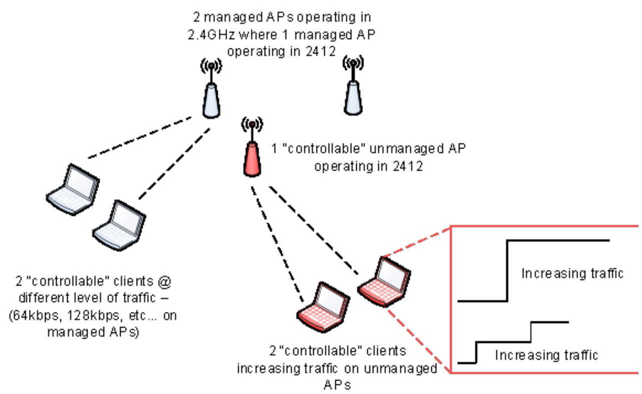

Figure 3 - Example Lab Test Setup

The test setup outlined in Figure 3 was used to create channel congestion and observe impacts on key operational metrics, e.g. latency, jitter, throughput, etc. A "target" AP was used as an observed target, with multiple client devices connected to it. Clients were distributed at distances of 3 – 10 meters from the access point. A mix of Voice over Wi-Fi (VoWiFi) and video traffic was run through the target AP. End-to-end metrics such as jitter, latency and throughput for the Voice over Wi-Fi (VoWiFi) and video traffic were measured using the IxChariot tool.

Separate "aggressor" APs were set up on the same channel and loaded with multiple clients sending iPerf and Youtube data. Identical tests were run with RRM disabled and enabled. The objective was to load up the channel to a congested level and: (a) observe the deteriorative impacts of rising congestion on service quality on the target AP (without RRM); and (b) observe the congestion mitigation and resultant quality improvement when RRM was enabled.

Test Results

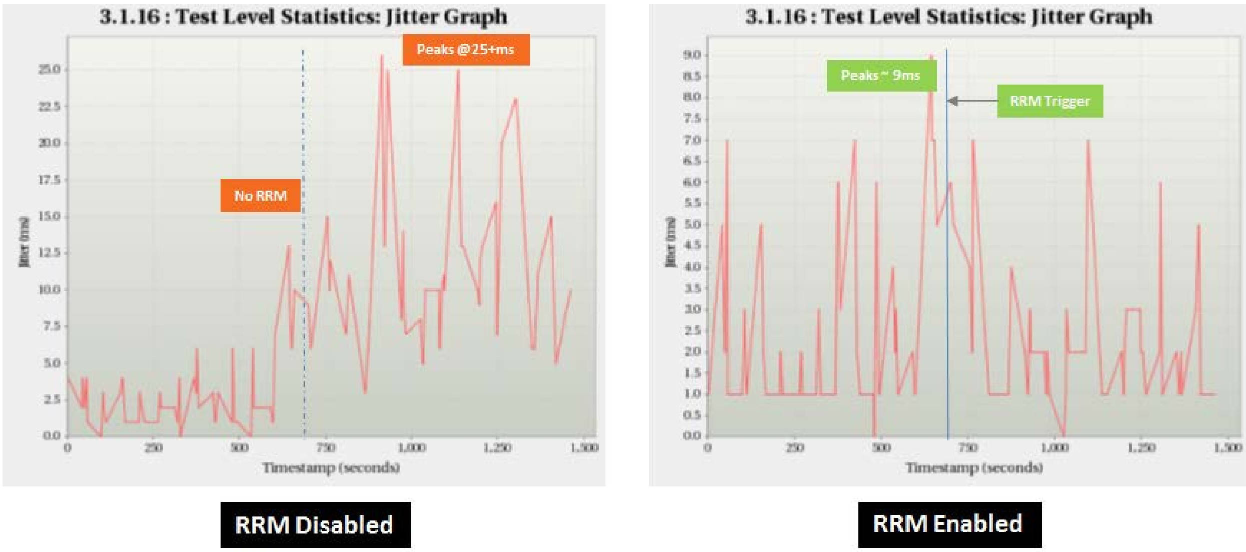

Figure 4 - Impacts on Jitter

Figure 4 depicts the impacts on jitter for VoIP, with RRM turned off and on respectively. With RRM disabled, jitter reaches a peak of 25 ms, and stays in a high range subsequently, hitting high peaks frequently. With RRM enabled, jitter is allowed to reach a peak of only 9 ms before RRM kicks in and brings congestion under control. Jitter subsequently settles into a 1 – 7 ms range, compared to the 5 – 25 ms range demonstrated without RRM.

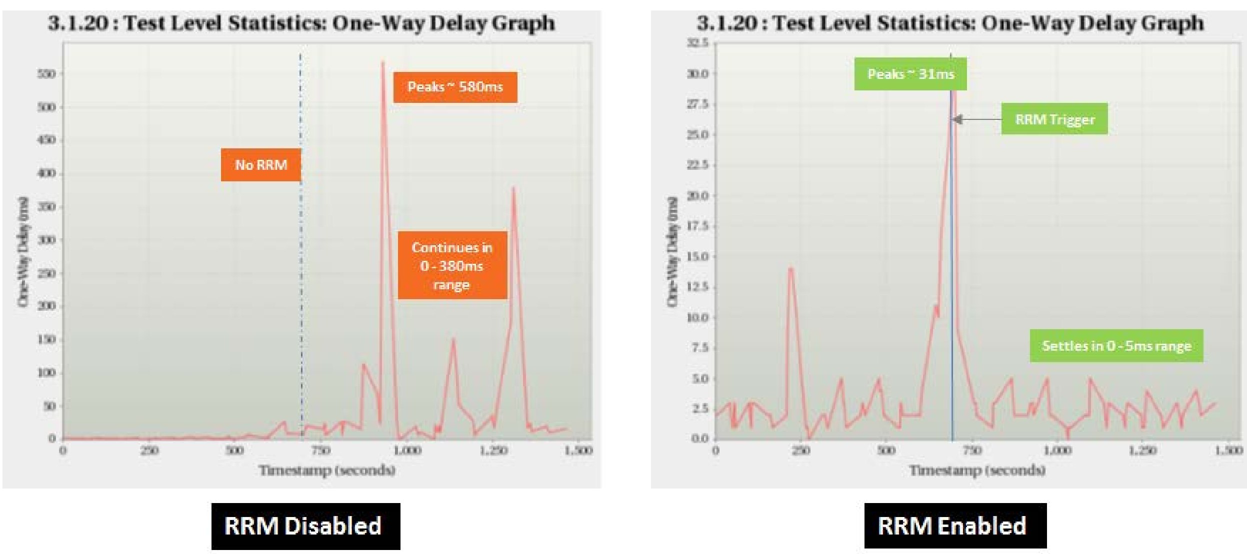

Figure 5 - Impact on VoWifi Latency

Figure 5 illustrates the impacts of channel congestion and RRM on one-way latency. Latency increases by an order of magnitude as congestion builds up, and without RRM, peaks at around 580 ms and persists within the 0 – 350 ms range. Latency levels above 100 ms result in noticeable interactivity and echo issues with voice calls. With RRM enabled, the congestion situation is corrected before latency escalates; in this case, latency settles down in an excellent 0 – 5 ms range, with a much lower and far more acceptable peak value.

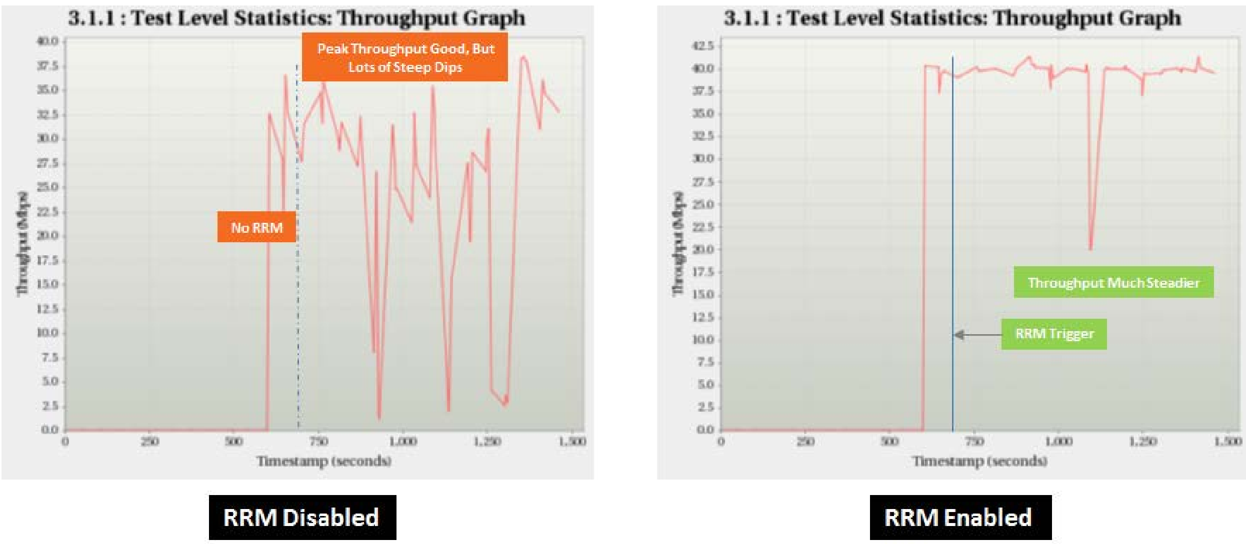

Figure 6 - Impact on Throughput

Figure 6 and Figure 7 respectively show the impacts on throughput and Mean Opinion Score (MOS) for VoWiFi. RRM keeps throughput much steadier and minimizes dips in throughput levels. The MOS drops to unacceptable levels without RRM, but hovers between excellent and acceptable levels with RRM enabled.

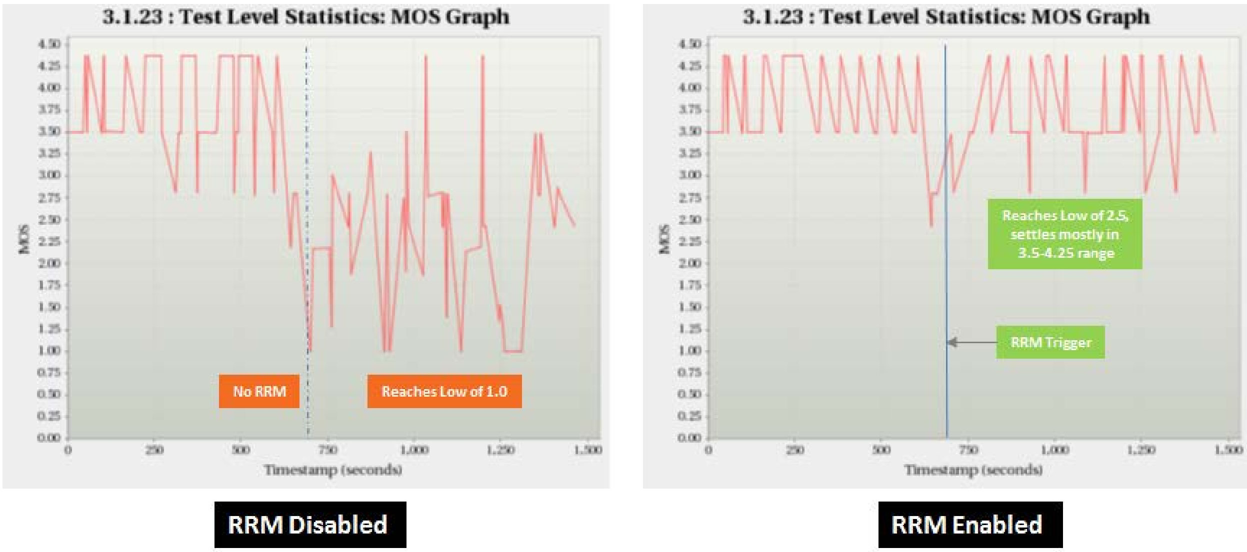

Figure 7 - Impact on MOS score for VoWiFi Call

2. Band Steering

RRM can also support band steering for associated clients, which can be steered to a different radio/frequency band to improve QoE. When an AP or specific radio on the AP is overloaded, and a channel change does not mitigate the situation, band steering functionality moves clients from the 2.4GHz band to the 5GHz band, or vice versa, to mitigate radio overload scenarios and improve QoE.

Test Setup & Methodology

To assess the performance impacts of band steering, a lab test setup was used to create radio overload and observe impacts on key operational metrics, e.g. latency, jitter, throughput, etc. On the target AP, multiple observable traffic flows (VoWiFi, video and data) were run between connected clients on the 2.4 GHz radio. End-to-end metrics such as jitter, latency and throughput for the VoWiFi and video traffic ("target traffic") were measured using the IxChariot tool.

Additional clients were then connected to the same (2.4 GHz) radio to introduce "aggressor" traffic on the radio. When the radio became overloaded, band steering kicked in and moved the "target traffic" flows from the 2.4 GHz radio to the 5 GHz band. The band steering algorithm factors in path loss adjustments for the move from 2.4 GHz to 5 GHz, i.e. given that range reduces when a client moves to 5 GHz, the algorithm only moves clients that have strong enough signal on 2.4 GHz to begin with; a calculation is done to ensure there is sufficient "coverage buffer" to withstand the path loss on moving to 5 GHz. The Chariot tool measured performance KPIs before and after the steering action.

Test Results

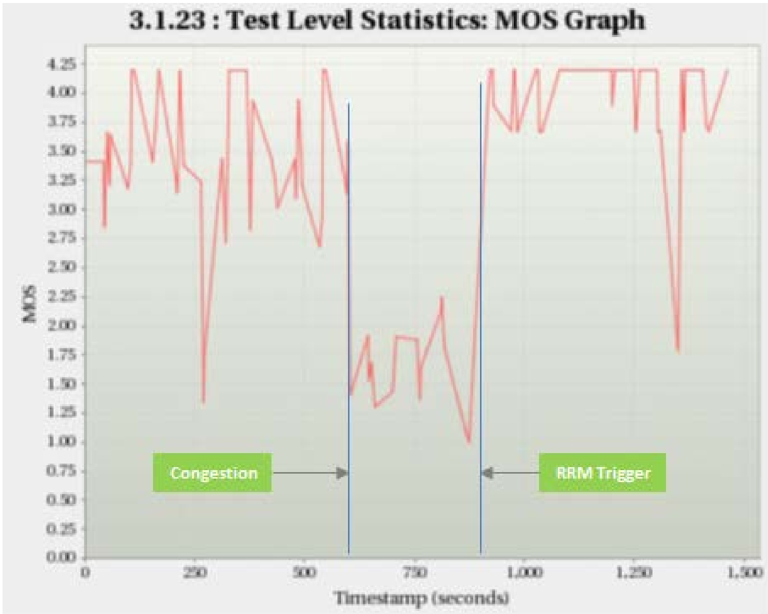

Figure 8 - Band Steering – Impact on MOS Score of VoWiFi

Figure 8 reflects the impacts of developing congestion and the band steering action on the quality of the VoWiFi call in progress (reflected by the MOS score). When the aggressor traffic raised congestion on the 2.4 GHz to high enough levels, the measured MOS deteriorated to unacceptable levels (staying between 1 and 2.25). When band steering kicked in and moved these clients to the much cleaner 5 GHz band, voice quality improved immediately and moved back up to acceptable levels.

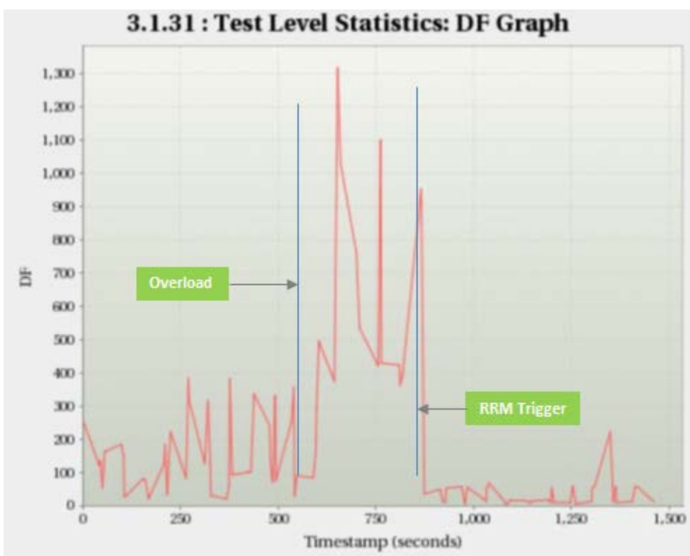

Figure 9 - Band Steering – Impact on Delay Factor

For the same test, Figure 9 shows the impacts of overload and band steering correction on delay factor for the video traffic flow. The delay factor (DF) KPI is measured in milliseconds, and is a time value reflecting the data buffer that has to be maintained to eliminate time distortions (jitter). DF levels reach around 1,300 milliseconds when the radio becomes congested, and settle back down to acceptable levels after the traffic flows are steered to the 5 GHz band. DF values in the 0 – 100 ms range are generally deemed acceptable.

3. Conclusions from Lab Tests

In general, the lab tests demonstrated the following trends:

- RRM brings about order-of-magnitude improvements in latency, jitter, throughput and other key service quality metrics
- With RRM enabled, we see significantly reduced metric spikes / dips. In other words, the KPI values deteriorate far less. Also, with RRM kicking in, we see much better settled metric levels. RRM ensures that congestion is quickly reduced, and that the KPIs that control service quality do not get out of hand.
- Without RRM, we see that the issues persist and escalate, resulting in much greater deterioration, drastic impacts to service quality and possibly eventual loss of service.
- These tests measured the impacts of congestion and RRM on real traffic flows, and it is worth noting that the results from the tests are fully applicable to similar traffic running in real world environment.

Enabling Wi-Fi Coexistence with New Technologies in Unlicensed Spectrum

While this paper has discussed RRM techniques and performance results in the context of Wi-Fi, they are equally applicable to cross-technology coexistence in the unlicensed band. Spectrum sharing and IoT are two areas that will introduce additional technology types into the unlicensed band. Flavors of LTE are poised to share spectrum with Wi-Fi; LTE License Assisted Access (LTE-LAA) and MulteFire will seek to share the 5 GHz band with Wi-Fi. IoT will stimulate increased use of technologies like ZigBee and Bluetooth, which also operate in the 2.4 GHz band.

Congestion and interference will multiply in severity when these new technologies operate in the band and contend with Wi-Fi and each other for the same set of resources. This will be driven by the many more devices and access points now “queueing up” to access the same set of unlicensed band channels. The 5 GHz band is relatively clear today, but can get significantly congested with the arrival of LTE flavors. The already crowded 2.4 GHz band will gain more popularity with IoT operating in it.

Some of the LTE – Wi-Fi spectrum sharing issues will be addressed by the lower layers of the LTE system. The new LTE-based unlicensed band technologies will likely support Wi-Fi –like schemes like Listen Before Talk (LBT), Dynamic Frequency Selection (DFS), Discontinuous Transmission and others, in an effort to minimize interference with Wi-Fi. In effect, these technologies will exhibit Wi-Fi –like behavior, at least with respect to channel access etiquette. These aspects are being standardized within 3GPP.

However, radio resource management schemes will be as important for interference and congestion avoidance, and to enable optimal sharing of unlicensed band resources between devices belonging to different technologies. RRM is a necessary complement to schemes like LBT, and can help reduce channel contention and the amount of time devices end up in "wait mode".

What is also interesting is that centralized optimization schemes (like XCellAir's) that rely on broader system performance metrics such as channel quality, device QoS parameters etc. can assess and mitigate congestion / interference impacts in a multi-technology environment without necessitating any cross-technology interactions. RRM can be used by a Wi-Fi system, for example, to protect itself from other Wi-Fi or LTE systems, without the system having to know of or communicate with the other systems present.

1. Channel Management

Dynamic channel management can be used by any system to protect itself from interference from other devices in the band. For example, a Wi-Fi system can proactively be allocated the best channels to mitigate interference:

- Between Wi-Fi devices in its own network (same operator)
- To / from devices in other Wi-Fi networks (other operators)
- To / from devices in LTE systems operating in the same band (same or other operators)
- To / from devices in Bluetooth or ZigBee systems operating in the same band.

The same schemes can be applied to an LTE-LAA / MulteFire system, or to an IoT system, to coordinate within its own network or protect it from other LTE or Wi-Fi systems. In a scenario where the operator is deploying both Wi-Fi and LTE devices, RRM can facilitate coordinated resource allocation across Wi-Fi and LTE. The performance improvements discussed in this paper are fully obtainable in such a multi-technology scenario.

2. Band Steering

Wi-Fi systems can continue to leverage band steering to keep devices in the best possible band. The 5 GHz band is a lot less congested than 2.4 GHz today, but that is likely to change with LTE devices entering the unlicensed band. Wi-Fi systems can use band steering to use the two bands optimally – move devices out of 5 GHz if high congestion is being caused by LTE, for example. Similarly, devices can be moved out of the 2.4 GHz band if IoT systems cause congestion there.

LTE-LAA and other systems can also use steering techniques to move devices and traffic flows between the licensed and unlicensed bands.

3. Power Control

In high-density deployments, Wi-Fi systems can use power control to minimize interference to other devices (Wi-Fi, LTE, others). Conversely, power can be increased to fill coverage gaps in pockets where deployment is less dense.

Conclusion

Clearly, radio resource management (RRM) is a critical necessity for optimal operation of networks in the unlicensed band. These techniques are readily applicable to Wi-Fi today, especially as cable operators and Internet providers roll out large and dense Wi-Fi networks. As use of the unlicensed band grows to include LTE and IoT devices, congestion and interference issues will worsen significantly, and the use of RRM schemes will become even more critical.

Resource contention, congestion and interference can cause harmful degradations in service quality, as tests described in this paper illustrate. Equally importantly, optimization schemes such as dynamic channel management, band steering and power control can help systems operating in unlicensed bands mitigate congestion and prevent severe service degradation - as quantified in this paper. The end results are the optimal usage of available unlicensed band capacity, dramatically improved KPIs and service performance, and a significantly enhanced quality of experience. These benefits are available to any type of system operating in this band – Wi-Fi, LTE and others.

It is also clear that even in dense network deployments, there is generally some spare headroom that can be leveraged by a smart RRM algorithm. XCellAir's studies have shown that there is, in general, spare channel bandwidth in the system – largely unutilized because most existing systems are unable to dynamically reallocate resources. An intelligent RRM scheme can unlock this free bandwidth for use by access points experiencing congestion or interference.