SiGNALS Ahead
Redefining Research
October 23, 2013, Vol. 9 No. 8

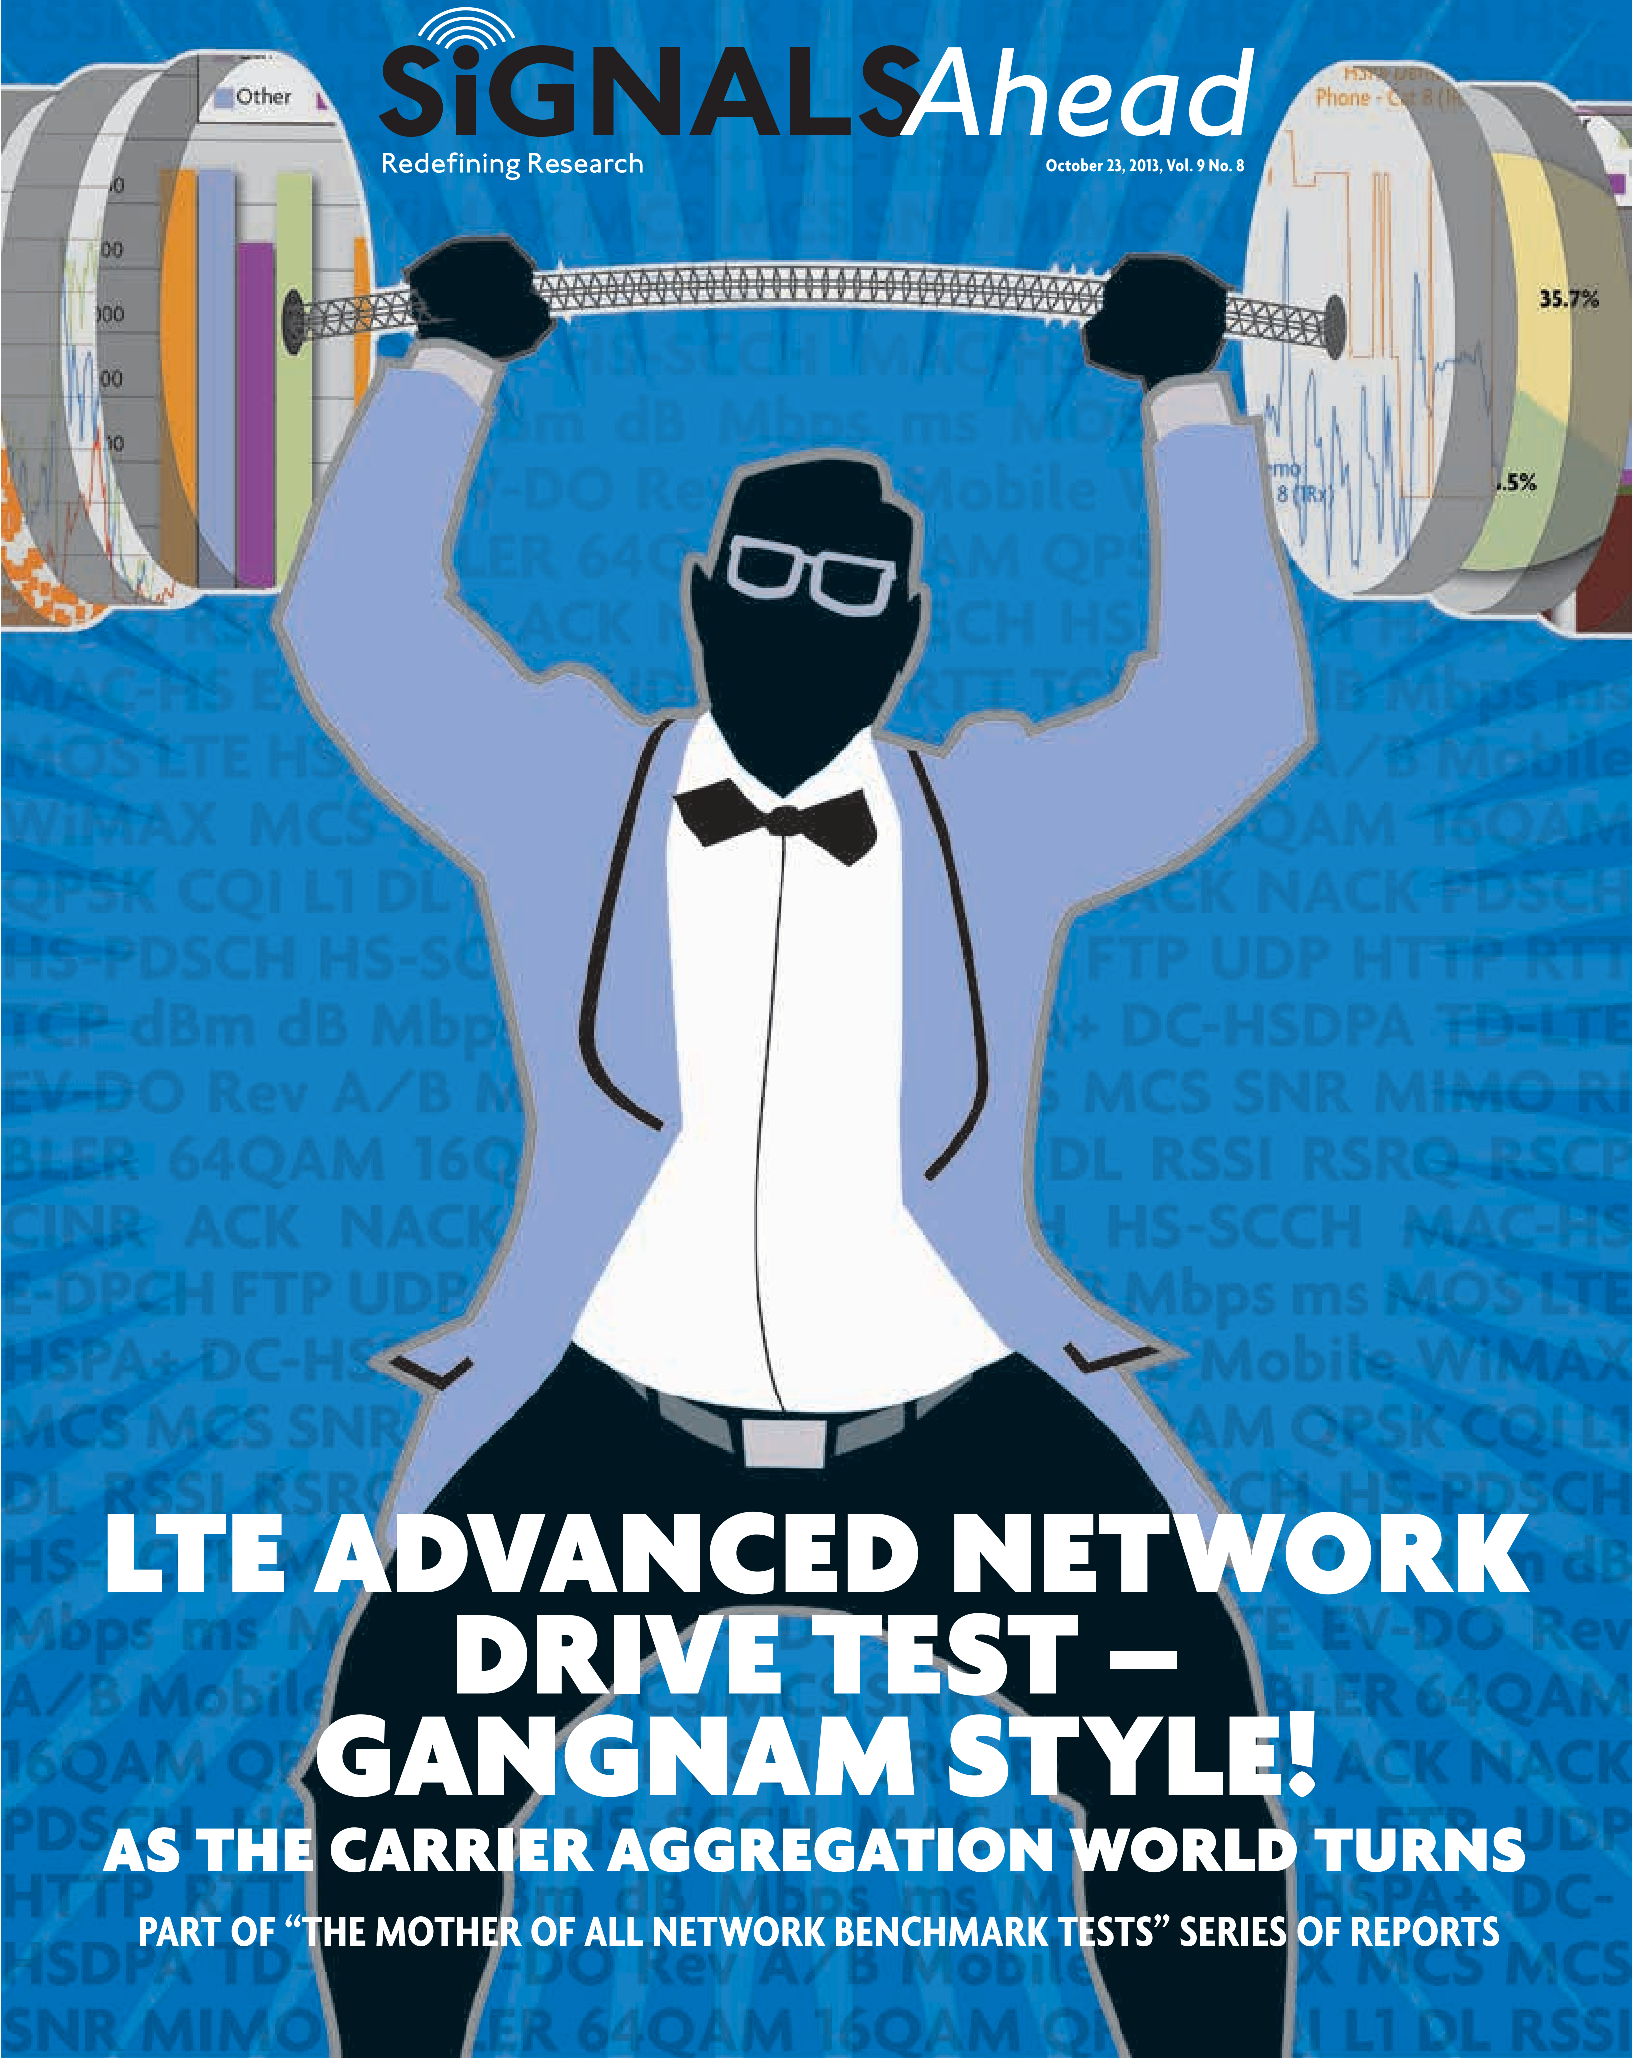

LTE ADVANCED NETWORK DRIVE TEST – GANGNAM STYLE!

AS THE CARRIER AGGREGATION WORLD TURNS

PART OF "THE MOTHER OF ALL NETWORK BENCHMARK TESTS" SERIES OF REPORTS

YOUR ATTENTION PLEASE

1.0 Executive Summary

Signals Research Group conducted what we believe is the first in-depth independent analysis of LTE Advanced Carrier Aggregation. This effort would not have been possible without the support of Accuver, who provided us with access to its XCAL data collection tool and its XCAP post-processing software. We have used the solution numerous times over the last several years so we are very accustomed to using it, although we do stumble upon new capabilities and features each time we use it. In our most recent benchmark studies, including LTE TDD in Tokyo and LTE Advanced Carrier Aggregation in Seoul, the solution's ability to support recently introduced technology features, including Category 4 chipsets and Carrier Aggregation, proved to be invaluable.

For the LTE Advanced testing we also used Spirent Communications' Datum solution, which it inherited when it acquired Metrico Wireless. We used Datum for some of the user experience tests that we conducted. We've used the tool in the past, including for a multi-network benchmark study that was commissioned by a major operator – Datum was the operator's preferred solution for the study.

After spending five days in Seoul testing LTE Advanced, our expectations for what we consider to be great network performance have been raised to an unattainable level.

After spending five days in Seoul testing LTE Advanced we are forever tainted and our expectations for what we consider to be great network performance have been raised to an unattainable level. To put things into perspective, the average downlink throughput during all of the testing was more than three orders of magnitude higher than what operators advertised a little more than a decade ago. The uplink throughput for a 10 MHz radio carrier was equally impressive, or roughly 33% higher than the best performance that we have observed in the past – AT&T's pre-commercial LTE network in Houston.

The great performance that we observed will be hard to replicate, largely because the South Korean operators have deployed very dense networks in Seoul. Separate from the published numbers of deployed base stations and remote radio heads in the country, the quality/density of the network can be observed in the stellar RSRP values, the high uplink throughput with modest transmit power levels, and the fact that we observed throughput greater than 100 Mbps during rush hour traffic for sustained periods of time.

Further, not all operators will deploy 10 MHz + 10 MHz implementations of carrier aggregation since at least in the near term some operators lack 10 MHz of contiguous spectrum in two suitable frequency bands. AT&T, for example, only has 5 MHz of 700 MHz spectrum in Chicago and Miami so the best it will be able to do in these markets is 5 MHz + 5 MHz or 5 MHz + 10 MHz. Other operators, including TeliaSonera and Verizon Wireless, may also slow down their rollout of carrier aggregation since with current chipset limitations they will not be able to use a full 2 x 20

MHz LTE carrier in Band 4 (VZW) or Band 7 (TeliaSonera). We hope to return to South Korea late next year when carrier aggregation with 10 MHz + 20 MHz channels is ready for primetime.

Despite the strong performance and the clear indication that carrier aggregation is already a fairly mature technology, we believe that the results could have been even better.

Despite the strong performance and the clear indication that carrier aggregation is already a fairly mature technology, we believe that the results could have been even better. In particular, we observed that the radio carrier in the lower frequency band had a much lower SINR than we would have expected and when the SINR dropped to a certain threshold, the throughput from the secondary carrier stopped. The issue, which can be resolved through additional network optimization, resulted in the LTE Advanced network reverting to LTE Release 8 more often than we would have liked, and the area where it occurred wasn't limited to the edge of the cell. We expect that this issue will largely resolve itself over time, but operators who are deploying carrier aggregation will need to play close attention to the matter.

Interestingly, we observed that the primary carrier wasn't always the radio carrier in the lower frequency band as one might expect. In fact, we have multiple log files in which the primary and secondary carriers switched frequency bands during a drive test. Operators in other parts of the world may not have the luxury of allowing either frequency band to be the primary/secondary carrier. This outcome is due to the intermodulation issues that can exist when combining certain frequency bands.

The beauty of carrier aggregation is in the little things that frequently go unnoticed.

The industry sometimes gets enamored in the peak data rates that carrier aggregation can deliver, but we believe the beauty of carrier aggregation is in the little things that frequently go unnoticed. For example, in our smartphone web browser tests we observed the concurrent use of both radio carriers, but due to the small data payload the frequency of using both carriers wasn't high and when both carriers were being used the throughput wasn't necessarily higher than what it could have been with a single 10 MHz radio carrier. Instead, what impressed us was that the network scheduler was smart enough to use either carrier by itself (i.e., we observed the use of the secondary carrier when the primary carrier wasn't being used), based on the underlying RF conditions in both bands. This functionality benefits the operator in terms of higher network capacity and it benefits the consumer in terms of a better user experience.

We also analyzed the performance of a Category 4 device and whether or not it provided any performance benefits relative to a Category 3 device. Depending on the time of day that we tested, we calculated that the Category 4 functionality was being used for 10 – 15% of the time – higher during the nighttime and lower during rush hour. When the Category 4 functionality was impacting the throughput, the average incremental increase in throughput was as much as 20% compared with a Category 3 device.

The user experience testing included web browsing, video telephony, VoLTE, Skype Voice/Video, 1080p video, and Google Play.

Chapter 2 contains the key observations and conclusions from our study. Chapter 3 provides our analysis of the downlink throughput, including the Category 4 device analysis and the results from testing the concurrent use of a Release 10 and Release 8 device. Chapter 4 focuses on the uplink performance, including a comparison with the LTE FDD and LTE TDD networks in Japan. Chapter 5 provides the results from some user experience testing, including web browsing, video telephony, VoLTE, Skype Voice/Video, 1080p video, and Google Play. In this chapter we also demonstrate why some popular third-party test measurement websites can report erroneously low throughput values. Chapter 6 contains our Test Methodology and Chapter 7 provides some very short closing remarks. All this and more in this issue of *Signals Ahead.*

Contents

Index of Figures

LIMITED FREE CARRIER PASSES AVAILABLE!

DON'T MISS OUT... CLAIM YOURS TODAY!

6th Annual

21st-22nd November 2013
Westin Galleria, Dallas, Texas, USA

lte™ north america 2013

(LTE IS A TRADEMARK OF ETSI)

Co-located with:
SIGNALING FOCUS DAY
connected cars 13 north america
MVNOs INDUSTRY SUMMIT USA 2013

Strategic Partner:
3GPP A GLOBAL INITIATIVE

Officially Endorsed by:
4G americas

Previous Event Host Carriers:
verizon
at&t
metroPCS

The *ONLY* 4G event in North America!

1250+	150+	75+	40+
High-level attendees	Distinguished Speakers	Exhibitors and meeting rooms	Analysts, Press & Influential Bloggers!

VISIONARY SPEAKERS ALREADY CONFIRMED FOR 2013:


BILL STONE
Executive Director, Network Strategy
Verizon Wireless


NAROTHUM SAXENA
Senior Director Advanced Technology & Strategy
US Cellular


MASAAKI MAEDA
President
NTT DoCoMo USA


RICK VERGIN
CEO
Mosaic Telecom


JOHN SAW
CTO
Clearwire


DEAN HARDY
CTO, **NetSet Communications**


EMIL OLBRICH
Public Safety Demonstration Network Project Lead Engineer, **NIST**


BERNARD ST-LAURENT
Project Manager, 700MHz Broadband Spectrum,
Public Safety Canada


SOLYMAN ASHRAFI
Vice President, Product Management
Metro PCS


GREGORY CASEY
CEO
Texas Energy Network


FEZA BUYUKDURA
Lead Principle Technical Architect, **AT&T**

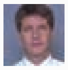
GLENN WELLBROCK
Director of Backbone Network Design
Verizon

CAMERON DUNN
Director International, Alliances & Integrations,
AT&T Mobility


BEJOY PANKAJAKSHAN
Principal Member Technical Staff, **Metro PCS**


ERIK NEITZEL
DMTS, Technology Development Group
US Cellular

BONNIE CHEN
Research Scientist
Sprint

LTE WORLD SERIES SUPPORTED BY

Alcatel·Lucent
CISCO
ERICSSON
HUAWEI
NEC
Nokia Siemens Networks
SAMSUNG
ZTE中兴

2013 SPONSORS ALREADY CONFIRMED:

Allot communications
bics
Citrix ByteMobile
COMMSCOPE
iBasis
kpn
OPENET
Panasonic
pt
radisys
SAP
Tekelec
TELMAR NETWORK TECHNOLOGY A JABIL COMPANY
Ulticom
WebRadar

Register today! Visit www.lteconference.com/northamerica

informa telecoms & media

IN CASE YOU MISSED IT: *SIGNALS AHEAD* BACK ISSUES

- **9/23/13 “124.2 GB in a LTE TDD Network - been there, done that, bought the [Helly Kitty] t-shirt”** Based on extensive testing in Softbank’s LTE TDD and LTE FDD networks in Tokyo, Japan, we provide the results from the industry’s first extensive drive test of the two duplex options. In addition to looking at the basic KPIs, like downlink and uplink throughput, we analyze more important KPIs, such as RSRP, transmit power and SINR to determine the important differences that inherently exist when the two networks are deployed in different frequency bands, with different cell site densities, and with different channel bandwidths. We also look at the incremental benefit of Category 4 devices with 20 MHz of LTE TDD spectrum and the performance of LTE TDD and LTE FDD with applications, such as Skype Video and Skype Voice. This report is a must read for any operator considering a dual network strategy with the two duplex schemes.

- **8/12/13 “Fifty Shades of MIMO (Quantifying the Impact of MIMO in Commercial LTE Networks)”** We provide test results and analysis that looks at the incremental impact of Transmission Mode 3 (Open Loop MIMO) versus Transmission Mode 2 (Receive Diversity) based on testing that we did in specially-configured commercial LTE networks. The results that we provide quantify how MIMO influences the downlink data rates as a function of various KPIs, including RSRP and SINR. In summary, we demonstrate that while MIMO can double the data rate, the actual benefits are far more modest, and a negative benefit is even possible. Further, we show that MIMO doesn’t necessarily improve the user experience in all cases, but there is still a benefit to the operator in terms of increased network efficiency.

- **5/28/13 “What’s the PSC, Kenneth? (Quantifying the need and benefits of interference cancellation solutions in a 3G network)”** We provide insight into the amount of interference that exists in a 3G network, its potential impact on data rates and network efficiency, and how an advanced equalizer can be used to maximize performance when these challenging conditions exist. For purposes of this report, we used AT&T’s HSPA+ network in San Francisco and the surrounding vicinity. This report was done in collaboration with Accuver who provided us with its XCAL and XCAP drive test solutions.

- **4/25/13 “Everything under the SON”** We discuss the background of SON, including discussions of work within NGMN, 3GPP and the SOCRATES/SEMAFOUR projects. We also cover the basics of SON including the laundry list of SON-like features, explain how they work, and what they mean for operators and vendors. We then move on to discuss the present and future requirements of SON, including what may be in store with Release 12 and beyond. Finally, we discuss the motivations and challenges of SON, including multi-vendor integration, vaguely-defined use cases, OSS limitations, 3G SON, and centralized versus decentralized architectures.

- **3/22/13 “Rich Communication Services - reinventing voice and messaging”** In this issue of Signals Ahead we provide a detailed analysis of RCS. In addition to providing the history of RCS since its introduction in 2008, we examine why operators have not yet fully adopted it, the capabilities by release, the inherent challenges that exist, the business relationships that exist or at least should exist, and the opportunities that could allow operators to beat the OTT providers at their own game. 2/25/13 “Chips and Salsa XVI: Sweet 16 and never been benchmarked” This report provides performance benchmark analysis of 8 LTE baseband chipsets, including Altair, GCT, Intel, NVIDIA, Qualcomm, Renesas Mobile, Samsung, Sequans. This benchmark study marks the 8th time that we have collaborated with Spirent Communications to leverage its 8100 test system and engineering support. All chipsets performed well under less challenging conditions but with the more challenging conditions there was a wide variance in the results with more than a 20% difference between the top- and bottom-performing chipsets. Three chipsets vied for top honors but ultimately we had to declare one the winner.

- **01/23/13 “The Mother of all Network Benchmark Tests - On the Inside Looking Out: evaluating the in-building performance capabilities of commercial LTE networks (Band 4, Band 7, Band 13, and Band 17)”** With the continued support of Accuver, we leveraged its XCAL-M drive test solution and its enhanced support for in-building testing to evaluate the performance of four LTE networks at Band 4, Band 7, Band 13 and Band 17. In this report we quantify the amount of LTE network traffic that we observed in the outdoor macro network and how it compares with our in-building testing. We also demonstrate that 700 MHz isn’t a panacea for in-building coverage, that potential coverage problems are being masked by ample capacity, and that some in-building networks may not scale to support future traffic demands. Finally, we compare and contrast the performance of the VZW and AT&T LTE networks.

- **12/5/12 “LTE Band 7 versus LTE Band 4 - GAME ON!”** With the support of Accuver, we used its XCAL-M and XCAP drive test solutions to conduct a network benchmark study of LTE Band 7 and LTE Band 4. This benchmark study leveraged the Rogers Wireless network in Vancouver, Canada where they have deployed both frequency bands in virtually every single cell site. In addition to looking at basic throughput, we include a host of other device-reported KPIs to analyze the downlink and uplink performance characteristics of the two frequency bands under identical network conditions, including edge-of-of cell and in-building.

- **11/6/12 “M2M – toward the Internet of things”** We analyze the M2M landscape and some of the key players involved in realizing this vision. The business models for M2M are still in flux and eventually multiple business models will have to be implemented. We look at the new business models being explored by mobile operators and MVNOs. The global connectivity requirements of M2M services make it natural fit for cloud services so there will need to be new cloud platforms in both the operator networks and enterprises to support M2M services. We also analyze the requirements and vendors for such platforms. More importantly, the radio and core networks will require enhancements to support the deluge of new M2M connections. We discuss some of the major issues and how the 3GPP standards body and operators are planning to address these issues.

2.0 Key Observations and Conclusions

Signals Research Group conducted what we believe is the first in-depth independent analysis of LTE Advanced Carrier Aggregation.

Signals Research Group conducted what we believe is the first in-depth independent analysis of LTE Advanced Carrier Aggregation. This effort would not have been possible without the support of Accuver, who provided us with access to its XCAL data collection tool and its XCAP post-processing software. We have used the solution numerous times over the last several years so we are very accustomed to using it, although we do stumble upon new capabilities and features each time we use it. In our most recent benchmark studies, including LTE TDD in Tokyo and LTE Advanced Carrier Aggregation in Seoul, its ability to support recently introduced technology features, including Category 4 chipsets and Carrier Aggregation, proved to be invaluable.

For this report, we also leveraged Spirent Communications' Datum solution, which it inherited as part of the Metrico Wireless acquisition. We used the Datum application to do some of the user experience testing and to do the subsequent analysis of the data. We hope to leverage the solution in the future in conjunction with a more thorough benchmark study of how carrier aggregation impacts web browsing as well as the overall network + device performance of VoLTE, including NB-AMR and WB-AMR.

SRG takes full responsibility for the conclusions and commentary included in this report. Based on our analysis of the data, we offer the following conclusions and observations.

Even in the very early days of its commercial availability, LTE Advanced with carrier aggregation performance was stellar. It is always a bit precarious to test a new technology in the very early days of its commercial availability. More often than not, the results are underwhelming and there are vendor limitations which degrade the performance from what it otherwise could have been. In the early days of UMTS, the handsets overheated and the peak data speeds seldom exceeded a few hundred kilobits per second. The supposed killer application, namely video telephony, was a joke unless you took pleasure in watching a pixelated and choppy image that frequently froze until the cell phone battery died a short time later.

LTE changed things and when we tested the first commercial networks a few months after they launched the performance was very impressive, but things weren't perfect. There were some network performance issues that appeared, one of the networks was limited to 10 MHz, and the only mobile device that existed was a USB dongle.

The introduction of LTE Advanced Carrier Aggregation has gone even better and the most significant performance issue that we identified (see the last observation in this chapter) seems to be network specific, albeit indirectly associated with the carrier aggregation functionality. And despite paying an arm and a leg for the LG G2 smartphone, it was one heck of a device. The video quality during the video telephony test was also phenomenal and a far cry from the 64 kbps video calls that UMTS tried to support.

The maturity of LTE Advanced stems from the maturity of LTE Release 8.

Like DC-HSDPA, carrier aggregation builds on the capabilities of an existing technology so to some extent the maturity of LTE Advanced stems from the maturity of LTE Release 8. There are technical nuances that get introduced into the standard with carrier aggregation but the overall LTE Release 8 protocols remain largely in place. We suspect there are, or will be, vendor-specific performance attributes that exist, just as they have existed with all previous technologies, and as we get the chance to evaluate other networks/vendors these differences will come to light. The impact of these differences on network performance remains to be seen.

However, we believe the biggest hurdles that operators will have to solve when it comes to deploying carrier aggregation are more basic and not necessarily technology related. Operators will need to deploy two LTE carriers at every single cell site and then optimize both carriers so that they have the same, or at least very similar, coverage footprint. There will need to be sufficient backhaul at all cell sites, otherwise the peak throughput of carrier aggregation will not be realized. Finally, operators will need to deploy carrier aggregation in those band combinations that the industry supports. This last requirement is part market-oriented – is there enough demand to justify solutions that support the specific pairing of frequencies – and part technical – does the proposed combination of frequencies introduce potential technical challenges due to interference.

Figure 1 provides a high-level summary of the downlink throughput results. We consider this figure to be “shareware” meaning that it can be used and shared externally. We do ask that you notify us if you wish to use the figure and that you give us full credit. This is the only figure in the report that can be used in this manner.

Figure 1. LTE Advanced Carrier Aggregation Downlink Throughput

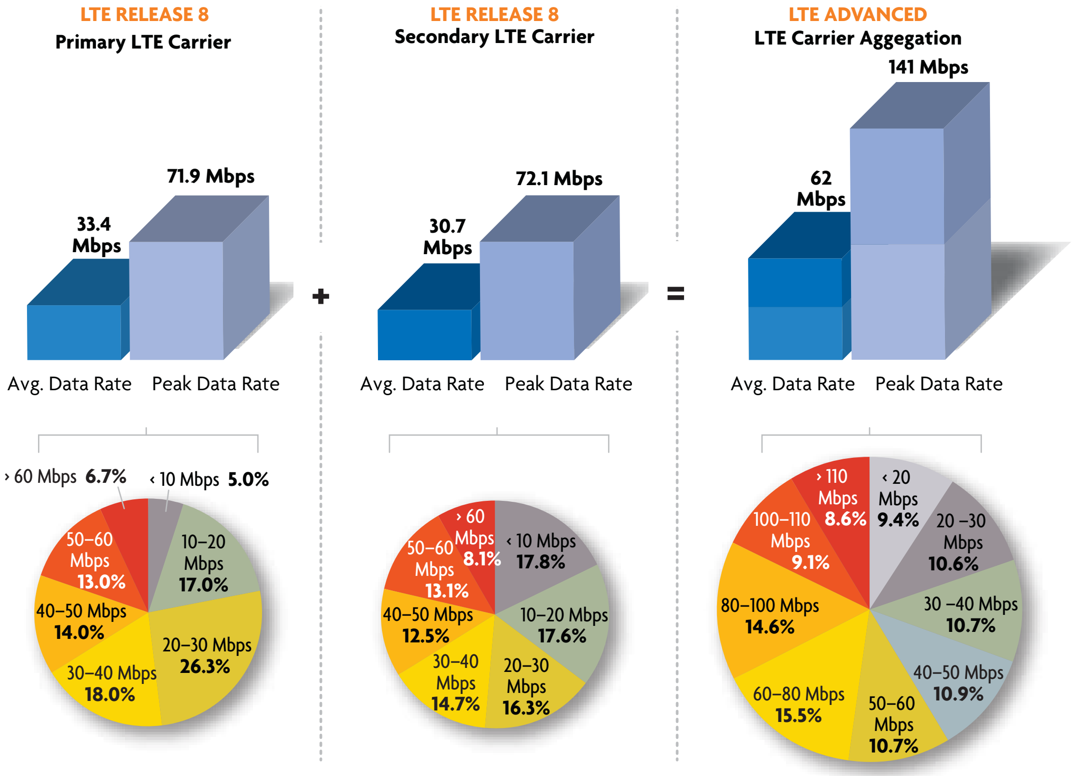

Source: Signals Research Group

The throughput was higher than 20 Mbps for 74.5% of the time.

THE UPLINK PERFORMANCE WAS SECOND TO NONE. Carrier Aggregation is limited to the downlink but we still wanted to collect at least some uplink throughput performance data. We’re glad we did because we observed what was arguably the best uplink throughput that we’ve seen in a 10 MHz radio channel. The average uplink throughput during a 20.9 minute mid-morning drive test was a remarkable 20.21 Mbps. Even more noteworthy, the throughput was higher than 20 Mbps for 74.5% of the time. For comparison purposes, the best we’ve seen in the past was AT&T’s pre-commercial LTE network in Houston where the average uplink throughput was 15.2 Mbps.

In the chapter on uplink performance (Chapter 4), we analyze the uplink throughput as a function of the transmit power and the number of allocated PUSCH resource blocks. We also revisit the results from Tokyo (LTE FDD and LTE TDD) to compare and contrast the network performance, including an analysis of the average transmit power per resource block and how it translated into a corresponding uplink throughput value. Based on this analysis the LTE FDD 2 x 5 MHz network in Tokyo takes the prize, followed by the same operator’s LTE FDD 2 x 10 MHz network.

RESULTS IN OTHER MARKETS WILL LIKELY BE MATERIALLY DIFFERENT. It would be nice to claim that the carrier aggregation results that we observed in Seoul will be replicated in other operators' networks, but that simply will not be the case. The primary reason is that while carrier aggregation can make a decent network look good by throwing more spectrum/bandwidth at the problem, there must be a great network in place (sans carrier aggregation) to get to the level of performance that we witnessed in South Korea. You have to be able to run before you can fly, and some operators are still taking baby steps when it comes to deploying a "Grade A" 3G/LTE network.

The relatively low cell site density in other markets is the primary reason why the South Korean carrier aggregation results will be difficult to replicate elsewhere. For example, the average RSRP that we observed in Seoul was in the low -70s, including the radio carrier at 2115 MHz. Throughout all of the early AM tests that we did, the average RSRP at 889 MHz was an almost unbelievable -67.07 dBm. Further, throughout all of the drive tests that we conducted, the RSRP seldom dropped below -90 dBm and it was only below -100 dBm for less than 1% of the time. Granted, we only tested outdoors from a moving vehicle but we doubt that there are networks in other parts of the world where we would see these kinds of results.

Not all operators will be deploying 10 MHz + 10 MHz implementations of carrier aggregation.

Another reason why we believe these results will be difficult to replicate, at least in the near-term, is that not all operators will be deploying 10 MHz + 10 MHz implementations of carrier aggregation. AT&T is probably the poster child for this claim since in some markets it only has 5 MHz LTE carriers at 700 MHz – Chicago and Miami come to mind. For AT&T, carrier aggregation means 5 MHz + 10 MHz – perhaps even 5 MHz + 5 MHz if it isn't able to clear 10 MHz of spectrum in the higher band. This deployment scenario doesn't negate the value of carrier aggregation. In fact, it increases its importance since it allows the operator to more effectively compete with its peers. Read the next key observation for another reason.

Operators who are currently supporting 20 MHz LTE channel bandwidths will be less aggressive when it comes to deploying carrier aggregation.

As a side note, with current chipset limitations we believe that operators who are currently supporting 20 MHz LTE channel bandwidths will be less aggressive when it comes to deploying carrier aggregation. For example, if an operator had 20 MHz LTE channels deployed in Band 7, it would have to first split the 20 MHz channel into two separate 10 MHz channels so that it could then combine one of the newly created 10 MHz channels in Band 7 with a secondary channel at a lower frequency band. This strategy wouldn't make sense since it would take away the higher data rate potential that all Release 8 mobile devices support while only offering the higher data rate to the limited number of Release 10 mobile devices. Since Verizon Wireless has 20 MHz of spectrum in Band 4 (AWS) in many of its markets, we believe that its initial carrier aggregation activities will be fairly limited. Once devices/chipsets are available sometime late next year that support a total channel bandwidth of at least 30 MHz these operators will take a more aggressive stance toward the Release 10 feature.

Some operators won't have the luxury of being able to assign the primary carrier to either the high or low frequency bands as was the case in Seoul. Combining 2100 MHz and 800 MHz is considered easy since there are not any intermodulation issues. In other words, the mobile device doesn't create self-interference when it receives data in both bands and simultaneously transmits in either the high band or the low band. For other combinations of frequency bands there is a high likelihood that self-interference could exist, and for this reason operators could be forced to use a specific frequency for the primary carrier. We addressed this issue in an earlier *Signals Ahead* report (reference *SA* 05/23/12, "Improve your RF Front-End in Seven Easy Steps").

THE BEAUTY OF CARRIER AGGREGATION IS IN THE LITTLE THINGS THAT FREQUENTLY GO UNOBSERVED. Truth be told, while it is still somewhat exciting to observe data speeds in excess of 100 Mbps – we recall getting excited when we saw 100 kbps toward the beginning of the last decade – there are not a lot of applications that benefit from these speeds. For example, with web browsing via a smartphone the TBS (transport block size) is hardly ever higher than 75,000 bits, indicating that a single 10 MHz channel could deliver the payload. In fact, a majority of the time the TBS while web browsing with a smartphone was less than 25,000 bits in the tests that we conducted. Even watching a 1080p video on YouTube didn't trigger the need for uber-high speeds, and, in fact, both radio carriers were seldom used when we watched a mind-boggling high-quality Planet Earth video on YouTube. With Skype video on the smartphone and the operator's video telephony service, carrier aggregation was non-existent. The simultaneous use of two radio carriers to download Angry Birds from Google Play was definitely nice and it did shorten the download time (peak data rate = 102.3 Mbps), but we still spent a lot of time navigating around the Play Store site before we clicked on the download button.

It was during the web browsing tests that the beauty of carrier aggregation was realized. During the smartphone web browser test, the simultaneous use of both radio carriers occurred for 31.31% of the time. Consistent with what one would expect, during web page browsing the use of MIMO limited the amount of time that both carriers were used while when MIMO wasn't being used, there was a higher dependence on using both radio carriers to deliver the relatively small payload.

The secondary carrier was used by itself for a measurable period of time, thus providing a tremendous benefit even in the absence of high data rates.

But what really struck us was that the secondary carrier was used by itself (e.g., no primary carrier) for a measurable part of the test, or 5.7% of the time. What this information tells us is that there were time intervals during the test when the eNodeB scheduler recognized that while it didn't have enough data to transmit over two radio carriers, it was more efficient to send the data over the secondary carrier than the primary carrier. Further, in more congested networks we would expect this percentage to become far more significant. This level of sophistication is good for the operator [increased network capacity] and it is good for the consumer [a better user experience], even if the marketing data rates associated with carrier aggregation never materialize.

Carrier aggregation seamlessly takes as much capacity as possible from both radio carriers and provides it to the mobile data user, while also being fair to other users in the network.

We observed a similar phenomenon during the downlink throughput tests. In those tests both radio carriers were contributing throughput but the dominant provider of the throughput varied. Sometimes the primary carrier contributed the most throughput and sometimes the secondary carrier provided the most throughput. From the perspective of the mobile data user, this level of information is hidden and it doesn't even matter since the only thing that matters is the overall throughput, regardless of how it is achieved. Carrier aggregation seamlessly takes as much capacity as possible from both radio carriers and provides it to the mobile data user, while also being fair to other users in the network. In the absence of carrier aggregation the mobile device would have to constantly switch between radio carriers and this action would be far less efficient, plus it would generate a tremendous amount of overhead.

THE BENEFITS OF USING A CATEGORY 4 DEVICE WERE FAIRLY EVIDENT IN THE DATA. Like we did in the LTE TDD tests, we analyzed the results to determine the incremental benefits of a Category 4 device versus a Category 3 device. The benefits of a Category 4 device are only realized with wider channel bandwidths so with a 10 MHz channel there wouldn't be any benefit. In Japan, we observed that the Category 4 LTE TDD device had some impact on the throughput but the benefit wasn't all that meaningful. It was a different situation in the LTE Advanced network.

The Category 4 functionality was being used 10-15% of the time, resulting in as much as a 20% increase in the average throughput compared with a Category 3 device.

Depending on the time of day that we tested, we calculated that the Category 4 functionality was being used for 10–15% of the time – higher during the nighttime and lower during rush hour. When the Category 4 functionality was impacting the throughput, the average incremental increase in throughput was as much as 20% compared with a Category 3 device. Averaged over the entire duration of the drive test, the incremental increase in throughput was barely measurable or a very low single-digit percentage. On a per TTI/subframe basis we frequently observed a peak TBS value of 146,784 bits – in one nighttime drive test we observed this value for 5.2% of the time. This TBS value was obtained when both carriers were active, MIMO was enabled in both carriers, and the TBS value per code word was 36,696 bits.

Category 4 functionality will not always provide a benefit, but it will be opportunistic and provide increased network efficiency and a better user experience whenever possible.

Like MIMO and 64 QAM, Category 4 functionality will not always provide a benefit, but it will be opportunistic and provide increased network efficiency and a better user experience whenever possible. A stationary user sitting somewhere within the inner circle of a cell could witness its benefits far more frequently than we did from our moving vehicle. From a network infrastructure perspective, Category 4 functionality comes for free, and, in fact, the networks have supported the higher data rates that we can only now observe from the first day the LTE networks were turned on. There is a cost associated with deploying Category 4 functionality, and this cost is in the form of more expensive devices/chipsets due to the additional processing and memory requirements. We don't have the cost analysis to make the claim that the incremental cost is worth it, but we have provided useful information on how this recently introduced device category can impact throughput.

The Category 4 device analysis appears in the first three sections of Chapter 3. Due to the tremendous amount of data generated when doing this analysis we limited the Category 4 analysis to only three representative drive tests – very early morning, mid-morning, and rush hour.

Additional network and/or device optimization should meaningfully improve the results that we observed. It is hard to complain about a new technology when out of the box it delivered an average throughput that was more than three orders of magnitude higher than the average throughput that operators advertised, but didn't necessarily deliver, a little more than a decade ago. Then again, we would be remiss if we didn't identify a few areas of opportunity for improvement.

In the network that we tested, the primary carrier was frequently [but not always] at the higher frequency band, or 2115 MHz, and the secondary carrier was at 889 MHz. As we understand it, with this combination of frequencies, either frequency could support the primary carrier with the actual selection based on which frequency the mobile device was camping on at the time the network activated carrier aggregation. The mobile operator might steer the mobile device to a particular band, else the mobile device/network will use the radio channel that offered the best network conditions.

There was a meaningful difference in the reported SINR between the two frequency bands with the 2115 MHz band offering the much better signal quality.

Throughout all of the drive tests, there was a meaningful difference in the reported SINR between the two frequency bands with the 2115 MHz band offering the better signal quality, albeit with a lower RSRP as one would expect. There was generally a 2–3 dB difference in the average SINR between the two bands during most drive tests, plus with the 889 MHz carrier the SINR was below 0 dB for a measurable period of time (e.g., 12.4% of the time during all of the nighttime testing that we conducted). When the SINR dropped to a lower value the throughput from the secondary carrier at 889 MHz typically disappeared. In effect, carrier aggregation reverted back to a Release 8 network. Once the SINR level improved the throughput from the secondary carrier would typically resume. Sometimes this sequence of events would take tens of seconds.

We also show in this report that the low SINR regions could extend throughout the entire cell and that the low SINR regions were not limited to the edge of a cell. During a lengthy drive test it wouldn't be uncommon for the secondary carrier to be "missing" for as much as 20% of the drive test. Given our test methodology and the use of high bandwidth servers, we believe there was enough data in the eNodeB scheduler's memory buffer to justify the use of carrier aggregation so the absence of the secondary carrier was seemingly due to the poor RF conditions.

Operators who deploy carrier aggregation will need to spend sufficient time optimizing their networks at some point during the pre-launch or initial launch phases.

Operators who deploy carrier aggregation will need to spend sufficient time optimizing their networks at some point during the pre-launch or initial launch phases. The advantage of a 10 MHz + 10 MHz carrier aggregation deployment relative to a single 2 x 20 MHz carrier is that the former will likely leverage a lower frequency band, which helps improve the coverage. Generally, operators don't have enough contiguous spectrum in the lower frequencies to support a single 20 MHz carrier so they have to use higher frequency bands that have unfavorable RF propagation characteristics. The disadvantage of using a lower frequency band and a higher frequency band with carrier aggregation is that radio carriers for both frequency bands must exist at every single cell site. Given the big differences in RF propagation between the high and low frequency bands it will be easy for the lower band to bleed into adjacent cells. Alternatively, if the lower frequency band serves as the primary carrier and the cell grid was designed to support the lower frequency band, then the secondary carrier, which is deployed at the higher frequency band, may not cover the entire cell. It is a problem that can be addressed but it takes time, and for Korean operators it doesn't help that they seemingly have a different cell located in every single urban block.

We also observed the loss of the secondary carrier due to what we believe was a higher protocol layer issue.

We also observed the loss of the secondary carrier due to what we believe was a higher layer protocol issue, and potentially a problem that resided within the mobile device/chipset. As part of our test methodology, we established multiple FTP and a UDP data session in order to sufficiently load the data pipe. Due to the relatively small size of the files on these servers we had to queue up multiple files so that when one of the file transfers finished another file transfer would start in its place.

Quite frequently, when a request for a new file transfer was sent it would disrupt the throughput from the secondary carrier for a brief period of time. This situation would happen even though there were multiple FTP sessions running in parallel – we ran two concurrent FileZilla applications plus a UDP application. The problem was also far more prevalent with UDP than it was with FTP. During the last two days of testing we stopped using UDP to send data and the problem was greatly minimized but it still occurred. We show an example in this report. Since the network scheduler should be indifferent to the transport protocol, we suspect that the problem was device/chipset related. In terms of "normal usage" the problem could appear when playing a YouTube video and synching email, for example. Whether or not it would actually impact the user experience is unclear, but it is a matter that is worth investigating.

In Chapter 3 we present the detailed analysis of the downlink throughput and how each radio carrier contributed to the throughput. We illustrate how the low SINR in the 889 MHz radio carrier impacted the throughput and how the carrier aggregation algorithm made adjustments to compensate for varying RF conditions.

3.0 Downlink Drive Tests – Detailed Analysis and Commentary

In this chapter we analyze three downlink drive tests in detail. The three drive tests span the range of possible scenarios, including light loading (Early Wednesday AM), heavy loading (1815 hours), and moderate loading (0906 hours). It is worth pointing out that we didn't observe major differences in the network performance between the light loading and heavy loading drive tests. We attribute this outcome to the density of the cell sites in Seoul. In essence, although there may have been a lot of mobile data users in the network, there were also a large number of cell sites so it was entirely possible that we had a cell almost to ourselves for a least a brief period of time, even during rush hour. In fact, we observed the maximum downlink throughput averaged over a one second increment in the 0906 drive test. It is also important to observe that even during the wee hours of the morning the city of Seoul doesn't sleep. *Gangnam Style!*

We observed the maximum downlink throughput averaged over a one second increment in the 0906 drive test.

3.1 Early Wednesday AM Drive Tests

The "Early Wednesday AM" drive tests are comprised of three separate drive tests. The first drive test started at 0313 hours (thanks to jet lag) and the last drive test ended at 0525 hours. During this series of tests we downloaded 59.45 GB of data while driving 29.3 miles. Figure 2 shows the roads that we used during the all outdoor vehicular tests.

Figure 2. Early Wednesday AM Drive Test Routes

Source: Signals Research Group

Figure 3 provides information about the distribution of throughput, including the primary and secondary carriers, as well as the total throughput (e.g., the sum of the two individual radio carriers). The maximum throughput (139.1 Mbps) is based on an average throughput over a one second increment.

Figure 3. Downlink Throughput by Primary and Secondary Carrier (Early AM) – Probability Distribution and Pie Charts

Primary Throughput (Avg) = 34.93 Mbps
Secondary Throughput (Avg) = 32.69 Mbps
Total Throughput (Avg) = 64.81 Mbps

Primary Throughput (Max) = 71.88 Mbps
Secondary Throughput (Max) = 72.05 Mbps
Total Throughput (Max) = 139.1 Mbps

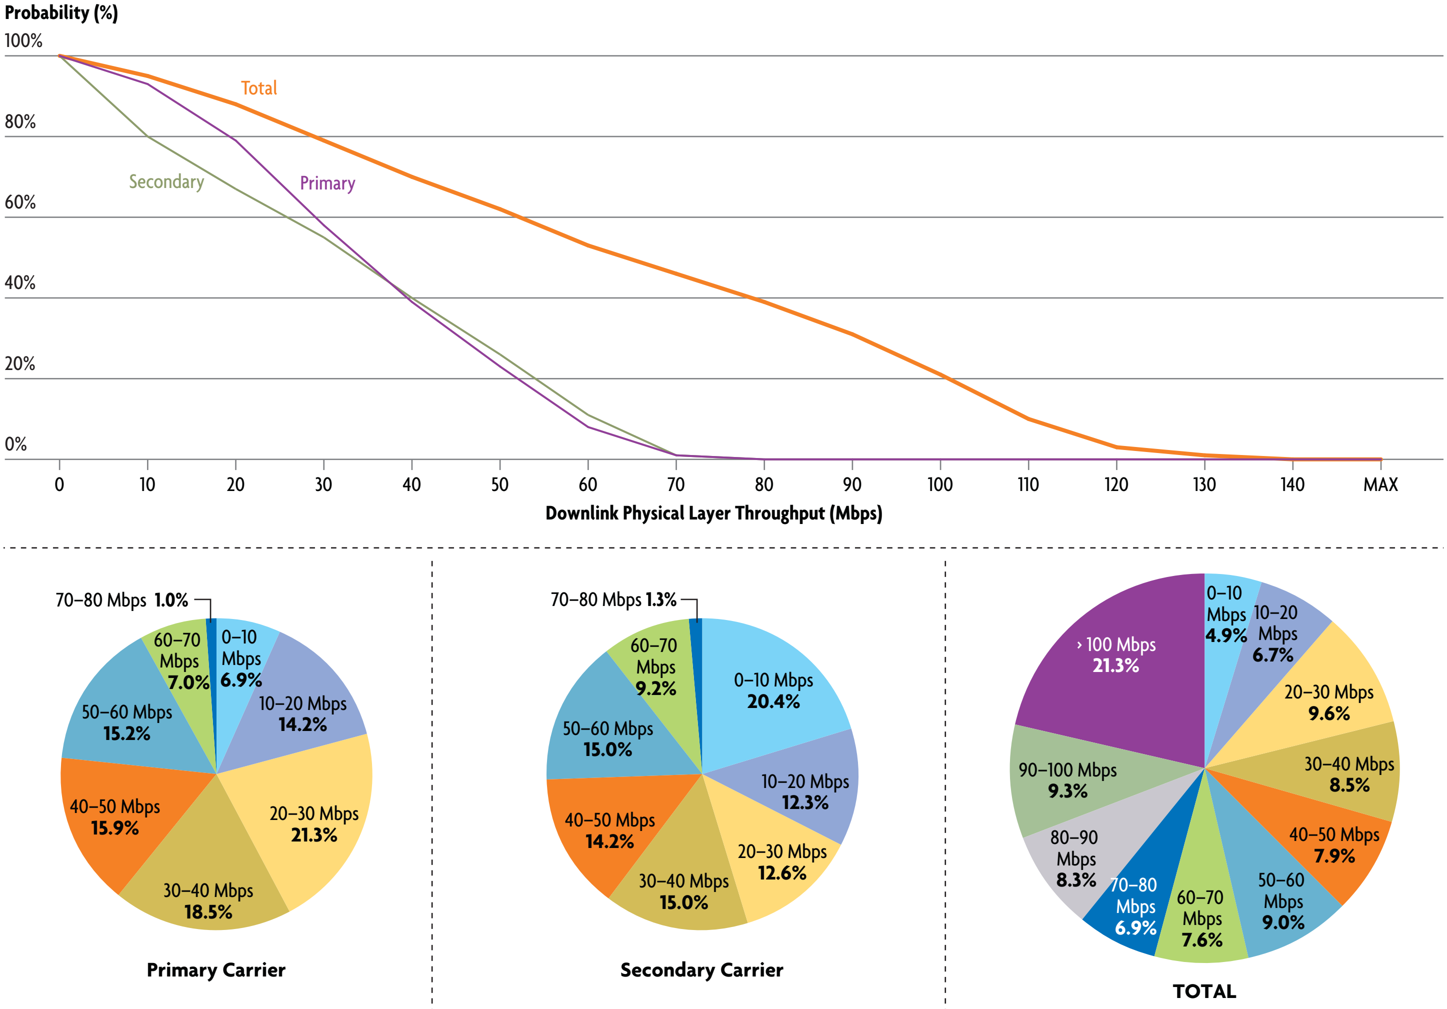

Primary Carrier

Secondary Carrier

TOTAL

Source: Signals Research Group

The use of the higher frequency band as the primary carrier is a bit counterintuitive.

Throughout all of these tests, the primary carrier was always 2115 MHz and the secondary carrier was always 889 MHz. The use of the higher frequency band as the primary carrier is a bit counterintuitive since operators typically want to use the lower band, which offers better coverage, for the primary carrier. However, we believe we understand the rationale as we will explain in a bit. One hint is apparent in the distribution of the throughput for the primary and secondary carriers. In particular, it is evident that at the lower end of the distribution, the primary carrier (higher frequency) had a higher probability of delivering better throughput than the secondary carrier. For example, the secondary carrier throughput was lower than 10 Mbps for 20.4% of the time and for the primary carrier throughput it was lower than 10 Mbps for only 6.9% of the time. The average values are based on the set of throughput values when both the primary and the secondary carriers were active and being used. If the secondary carrier wasn't active and being used then the throughput from the primary carrier was also excluded before doing the calculation. We'll explain the relevance of this point later in the section.

The signal quality (SINR) was much better with the primary carrier than with the secondary carrier.

Figure 4 provides the RSRP probability distribution for the primary and secondary carriers. Figure 5 provides the comparable information for the SINR. As shown in Figure 4, the RSRP favored the secondary carrier (889 MHz) but it is also evident that the signal quality (SINR) was much better with the primary carrier than with the secondary carrier (reference Figure 5). In fact, the secondary carrier SINR was less than 0 dB for 12.4% of the time versus only 0.9% of the time for the primary carrier. Unless the mobile operator steers traffic to a particular band, the mobile device will camp on the radio carrier that offers the best RF conditions. In this case, the LTE carrier at 2115 MHz generally offered the best RF conditions so the mobile device camped on this carrier and it became the primary carrier when the network activated the carrier aggregation functionality.

Figure 4. RSRP by Primary and Secondary Carrier (Early AM) – Probability Distribution

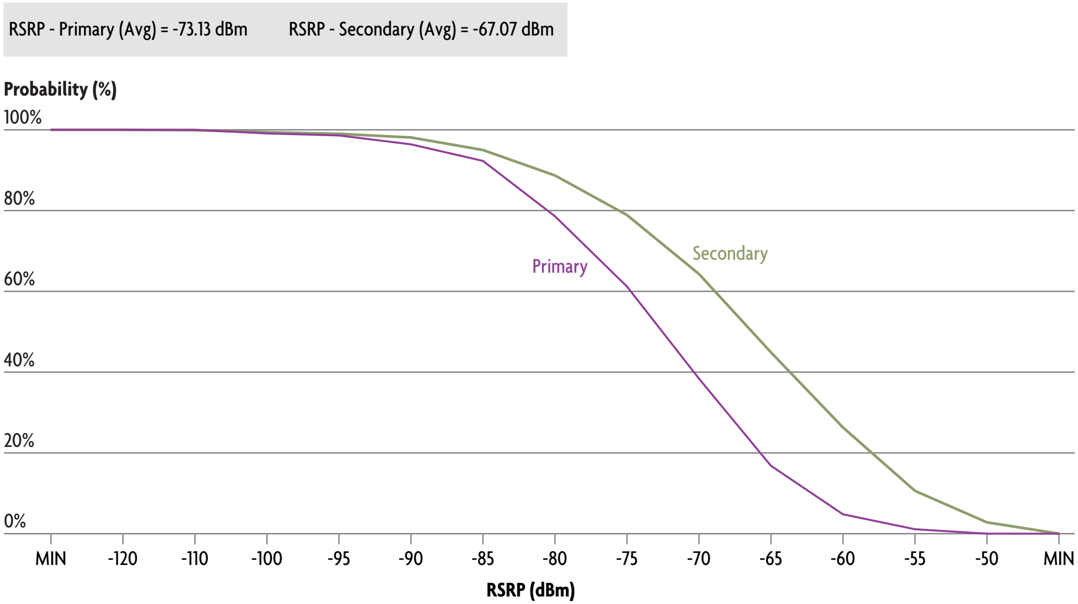

Source: Signals Research Group

With additional network optimization, in particular involving what is now the secondary carrier at 889 MHz, we anticipate that many of the issues that we discuss in this report will vaporize. At that point the 889 MHz radio carrier could also become the default primary carrier.

Nothing that we have observed in other networks comes close to the RSRP values that we observed in Seoul.

In our last report, we highlighted the very favorable RSRP values in Softbank's FDD network – average RSRP = -74.38 dBm. Comparatively speaking, nothing that we have observed in other networks comes close to the RSRP values that we observed in Seoul. An average RSRP of -67.-7 dBm is practically unheard of, not to mention essentially no RSRP values lower than -100 dBm, and we wouldn't have believed it if we hadn't been there to log it. This situation is due primarily to the cell site density.

Figure 5. SINR by Primary and Secondary Carrier (Early AM) – Probability Distribution

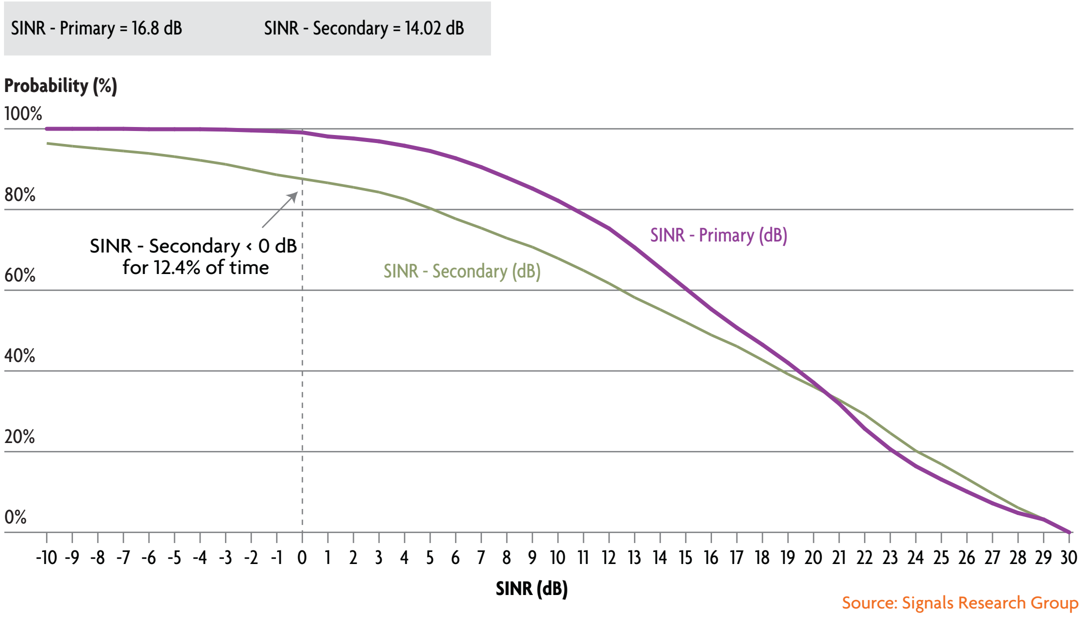

Source: Signals Research Group

Figure 6 provides probability distribution plots for the number of assigned Resource Blocks for both carriers. Interestingly, the plots are a bit "worse" than the results that we obtained during early evening rush hour (reference Figure 19).

Figure 6. Resource Block Allocation by Primary and Secondary Carrier (Early AM) – Probability Distribution

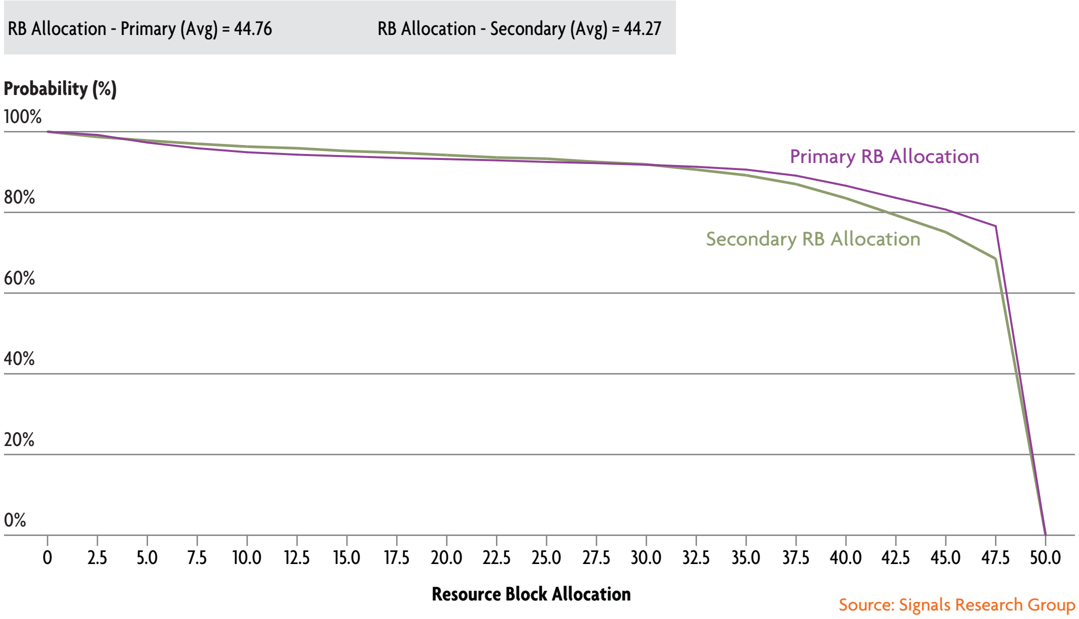

Source: Signals Research Group

Carrier aggregation can seamlessly take advantage of whatever network resources are available from the two discrete radio carriers.

Figure 7 provides a wealth of information. The top figure, which is a bit cumbersome to view, shows the throughput from the primary and secondary carriers as well as the total throughput for the first 1,000 seconds (16.67 minutes) of the first drive test. Readers should note two observations. First, there are numerous periods of time when the secondary carrier is providing higher throughput than the primary carrier. This phenomenon speaks to the real value of carrier aggregation. It isn't the maximum throughput that matters most but the realization that carrier aggregation can seamlessly take advantage of whatever network resources are available from the two discrete radio carriers. Second, there are more than a couple of instances where the secondary carrier doesn't contribute any throughput, meaning that the total throughput was based entirely on the throughput from the primary carrier.

Figure 7. Detailed Analysis of MIMO, Secondary Carrier and Category 4 Device Utilization (Early AM)

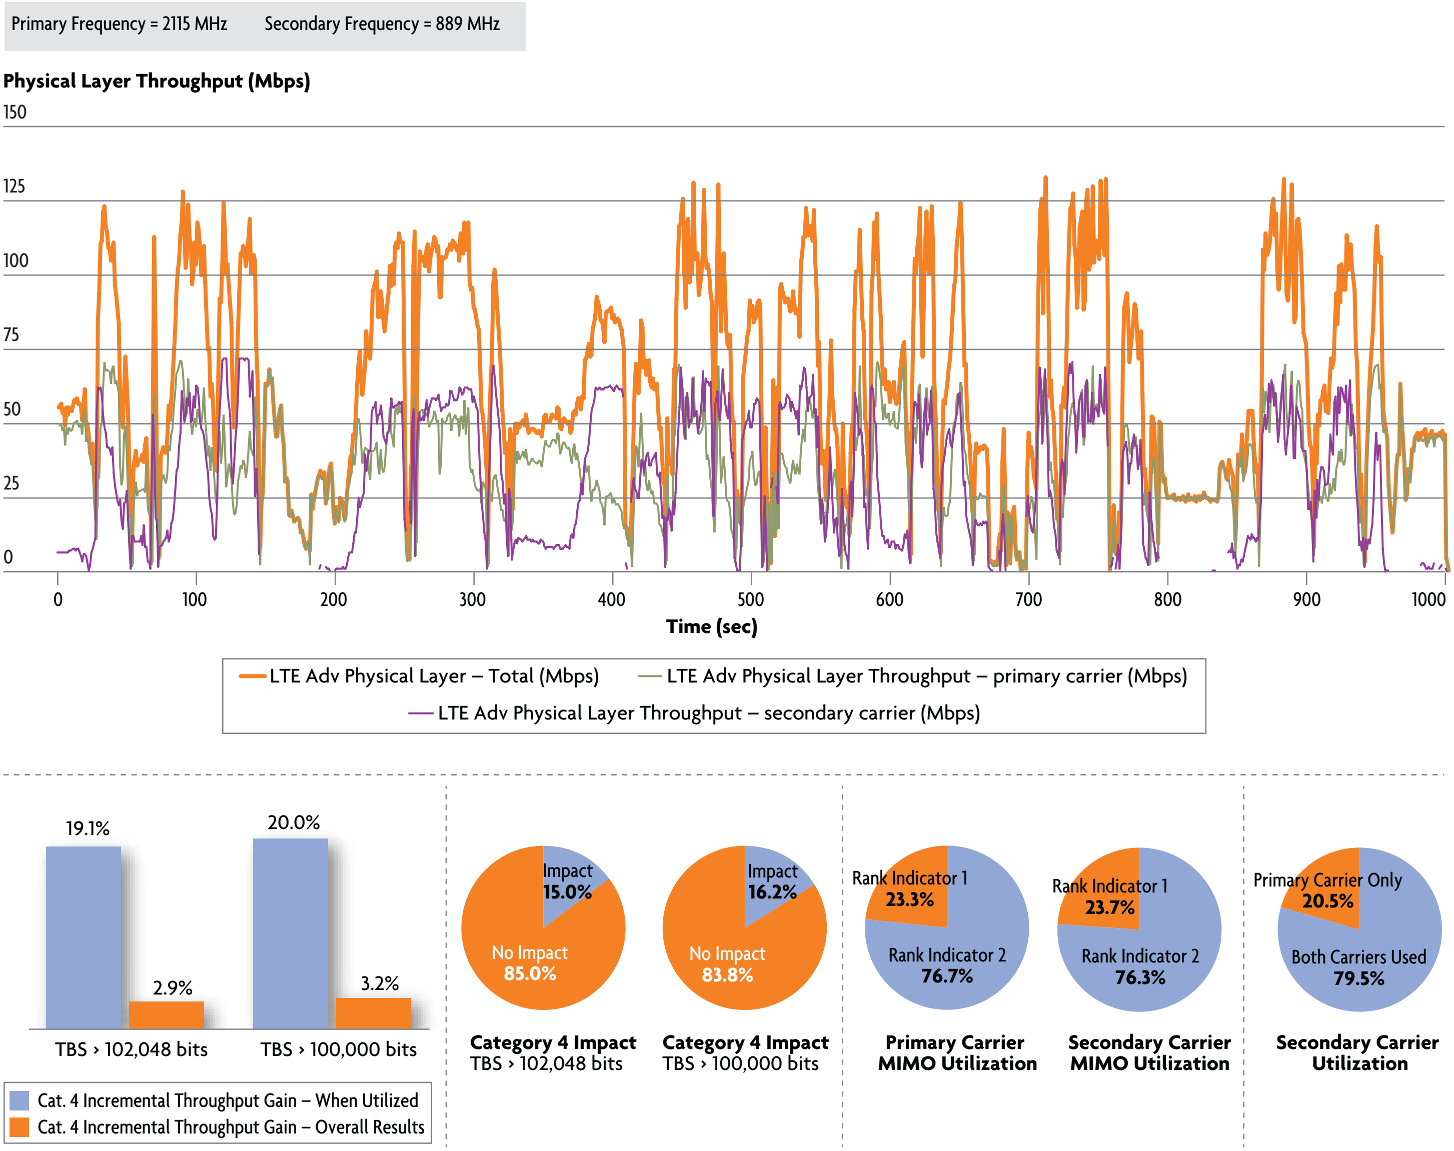

Source: Signals Research Group

The bottom pie charts and the bar chart provide information about the benefits of a Category 4 device, the use of open loop MIMO for each radio carrier, and the concurrent use of both radio carriers. The values are limited to the results from the first drive test (0313 hours) due to the tremendous amount of data that must be analyzed. MIMO utilization and TBS analysis requires looking at values that are reported per one millisecond subframe, meaning a 45 minute drive test could have up to 2.7 million sets of data.

20.5% of the time the total throughput was based solely on the contribution from the primary carrier.

The Secondary Carrier Utilization pie chart indicates that the secondary carrier was only used 79.5% of the time, in other words, 20.5% of the time the total throughput was based solely on the contribution from the primary carrier. Given the test scenario we would have expected this percentage to be at or near 100%. If we calculate the average throughput based only on the times when both radio carriers were being used then the average total throughput increases to 68.88 Mbps versus the 64.81 Mbps that we indicated in Figure 3.

The MIMO utilization pie charts show how frequently the network used open loop MIMO. Both percentages are very good, but it is worth noting that the secondary carrier MIMO utilization values exclude all instances when the secondary carrier wasn't being used.

The benefits of using a Category 4 were evident at least 15% of the time during the drive test.

The benefits of using a Category 4 device were evident quite often and the incremental benefits on throughput were measurable. For this analysis, we assumed that a Category 3 device was limited to a transport block size (TBS) of 102,048 bits. We also ran the analysis assuming that a Category 3 device was limited to a TBS of 100,000 bits. Thus, any time we observed a higher TBS value in the log file, we knew that it could only have been realized with a Category 4 device.

The Category 4 Impact pie chart shows that during the 35.45 minute drive test, the Category 4 functionality was used 15% of the time – 16.2% of the time if we assume the lower TBS threshold. During these periods of time, the incremental throughput gain of the Category 4 device relative to a Category 3 device was 19.1%, or 20% with the lower TBS threshold. The benefits of a Category 4 device were less meaningful when applied across the entire test – the incremental throughput was a low single digit percentage.

Figure 8 provides a time enhanced view of the time series plot in Figure 7. Specifically, it shows the first 300 seconds of the drive test. The figure shows that the cell handovers for the primary and secondary radio carriers occurred at the exact same times and involved the same PCI values. We have also highlighted a time when the throughput from the secondary carrier was no longer present. This particular time occurred during a cell handover.

Figure 8. Downlink Throughput by Primary and Secondary Carrier (Early AM) – Time Series

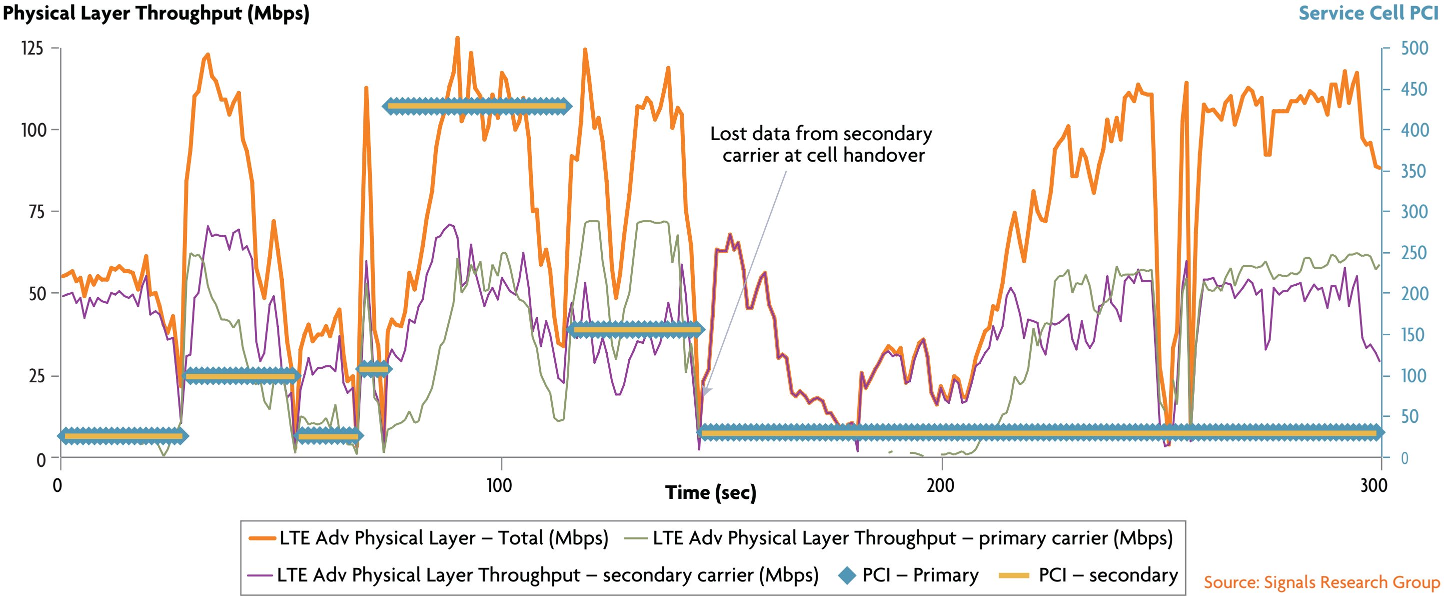

Source: Signals Research Group

Figure 9 and Figure 10 provide some additional insight into the network conditions. Before and after the cell handover, the RSRP from both radio carriers was quite good (reference Figure 9), however, it is evident that the SINR from the secondary carrier became quite poor (e.g., well less than 0 dB) immediately following the handover, thus explaining why there wasn't any throughput from the secondary carrier.

Figure 9. RSRP by Primary and Secondary Carrier versus Serving Cell PCI (Early AM) – Time Series

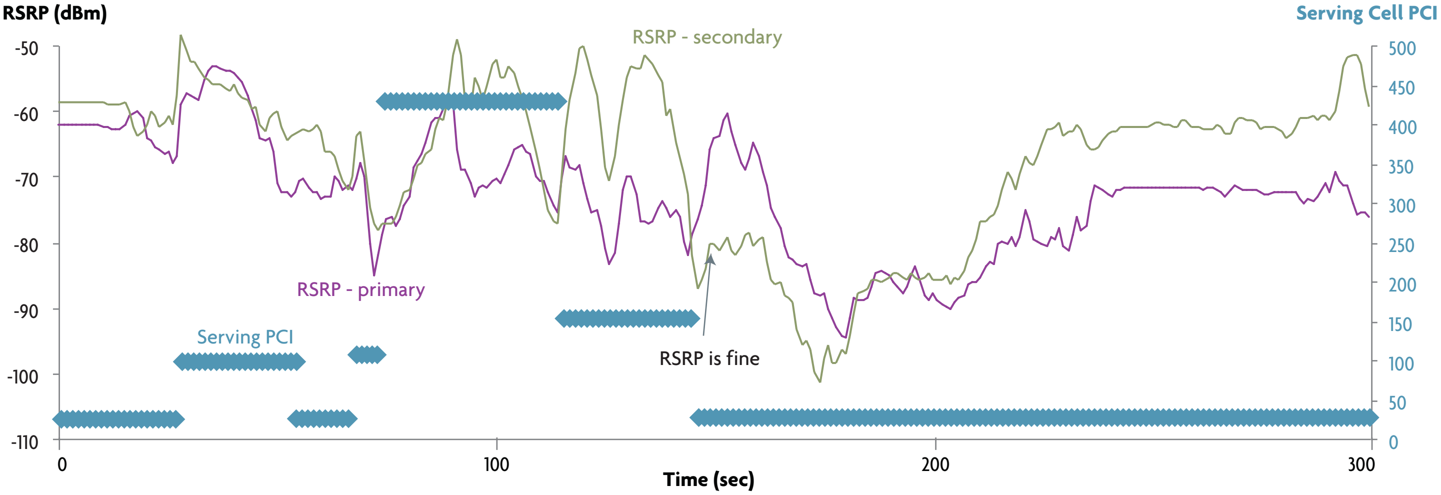

Source: Signals Research Group

Figure 10. SINR by Primary and Secondary Carrier versus Serving Cell PCI (Early AM) – Time Series

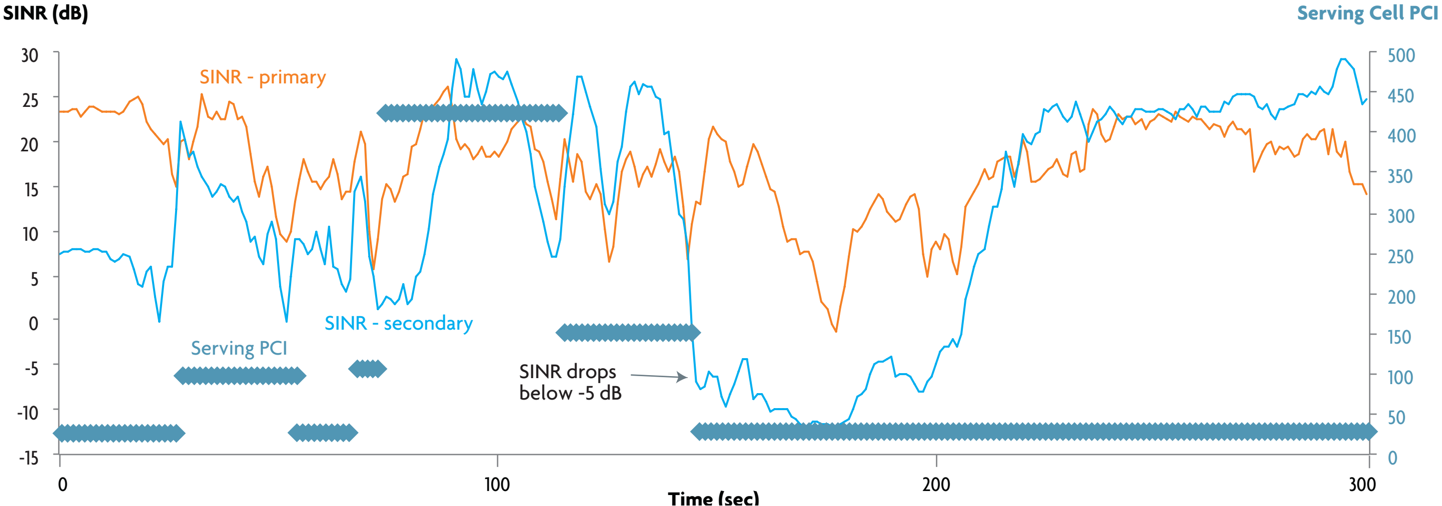

Source: Signals Research Group

Figure 11 and Figure 12 look at the situation somewhat differently. Figure 11 shows a geo plot of the serving cell PCI values for the secondary radio carrier. Figure 12 shows the corresponding SINR values for the secondary carrier. To our eyes there are some apparent correlations between a change in the serving cell PCI value and the measured SINR. In some cases, the SINR looks poor throughout the entire cell and in other instances there is an abrupt change in the SINR when the serving cell PCI value changes.

Figure 11. Secondary Carrier Cell PCI Values – Full and Enhanced Views

Figure 12. Secondary Carrier SINR Values – Full and Enhanced Views

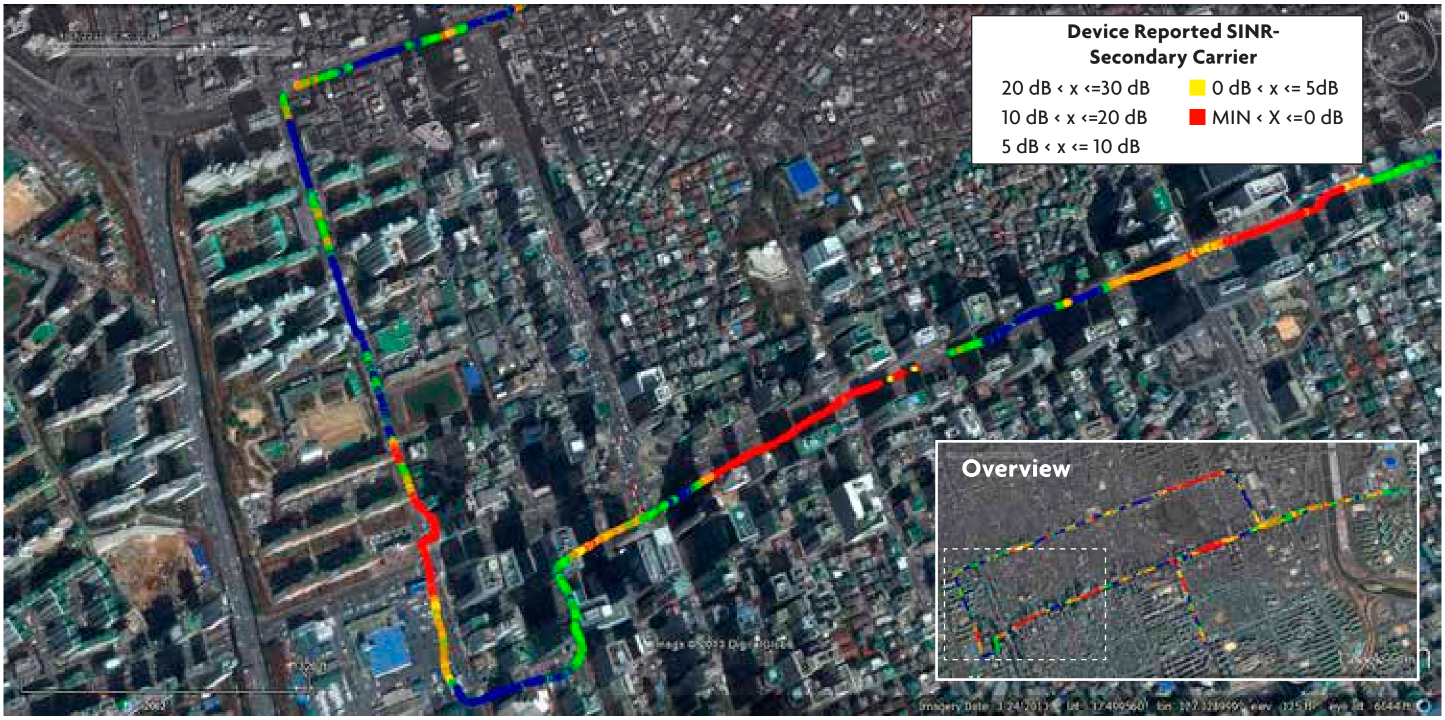

Source: Signals Research Group

Figure 13 provides a plot of the resource block (RB) allocations for both radio carriers as a function of time. We've marked one spot where we believe the sharp drop in the number of assigned RBs was due to another mobile device in the network. It could have also been due to the higher layer protocol issue that we mentioned in Chapter 2. The first sharp drop in the secondary radio carrier RB allocation corresponds to a sudden drop in the SINR while the other sharp drops seemed to occur during cell handovers.

Figure 13. Resource Block Allocation by Primary and Secondary Carrier (Early AM) – Time Series

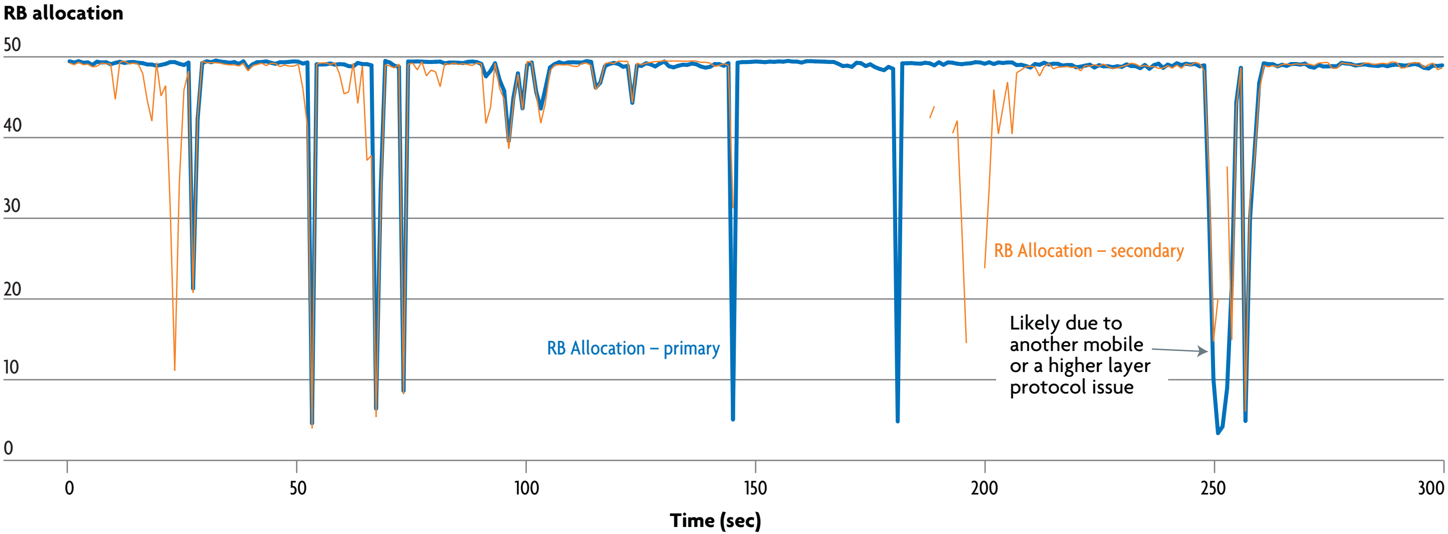

Source: Signals Research Group

Figure 14 provides scatter plots of the SINR versus the downlink physical layer throughput for the two radio carriers. In the lower figure we have highlighted the region where the SINR was low, resulting in low or no throughput.

Figure 14. SINR versus Downlink Throughput – primary and secondary carriers (Early AM) – Scatter Plots

LTE Adv Throughput - primary carrier (Avg) = 34.9 Mbps LTE Adv SINR - primary carrier = 16.80 dB

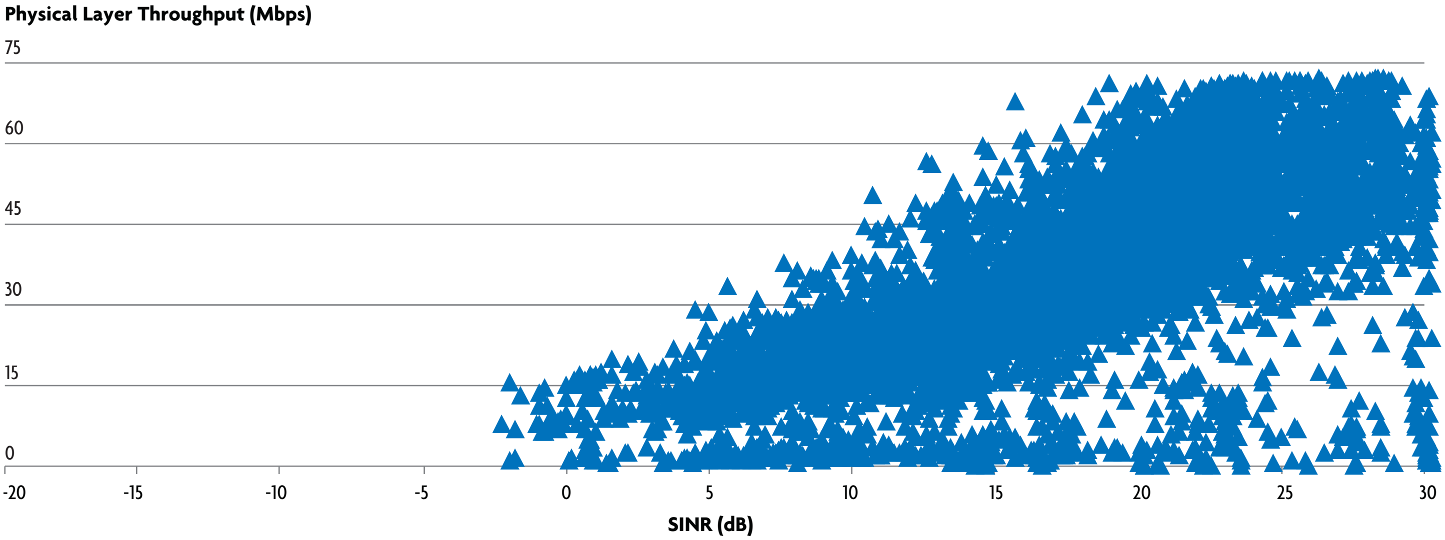

LTE Adv Throughput - secondary carrier (Avg) = 29.97 Mbps
LTE Adv Throughput - secondary carrier (Adjusted Avg) = 32.7 Mbps
LTE Adv SINR - secondary carrier = 14.02 dB

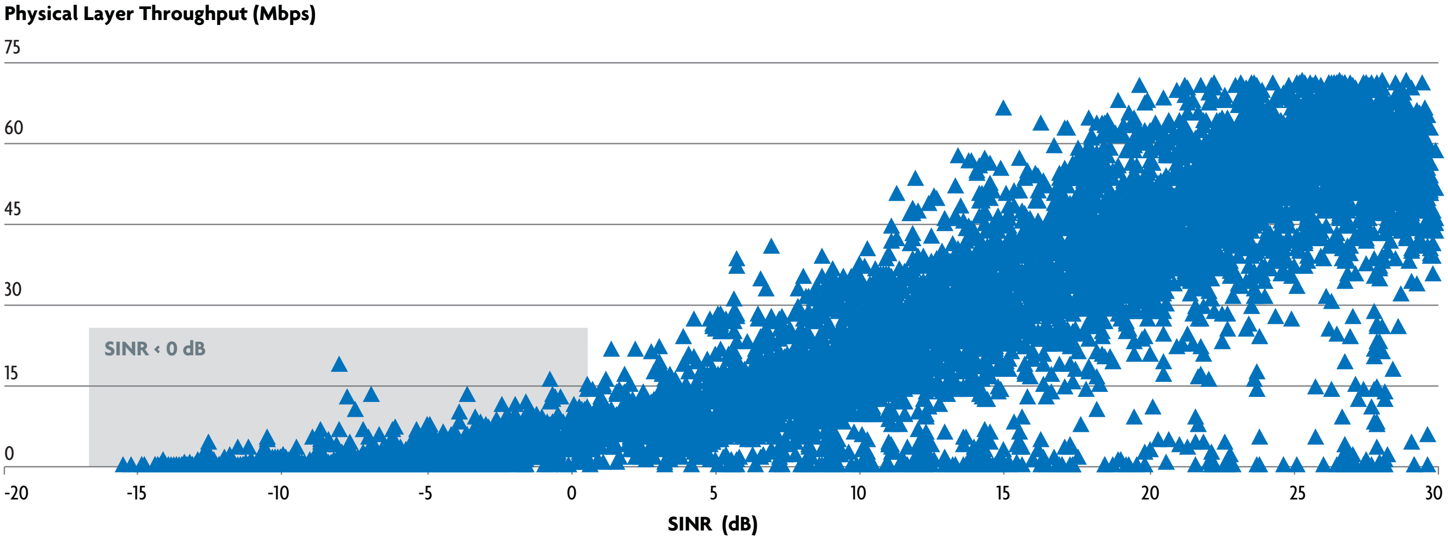

Source: Signals Research Group

3.2 1815 Hours Drive Test

This drive test occurred at 1815 hours on August 27th or right at the heart of rush hour traffic. The test lasted 34.6 minutes and covered 2.76 miles during which time we downloaded 12.3 GB. Figure 15 provides a geo plot of the drive route.

Figure 15. 1815 Hours AM Drive Test Routes

Source: Signals Research Group

Figure 16 provides the distribution of downlink throughput during this drive test during which time the average throughput was a very impressive 48.3 Mbps (peak = 122.52 Mbps).

Figure 16. Downlink Throughput by Primary and Secondary Carrier (1815) – Probability Distribution and Pie Charts

Primary Throughput (Avg) = 26.02 Mbps	Primary Throughput (Max) = 69.41 Mbps
Secondary Throughput (Avg) = 23.82 Mbps	Secondary Throughput (Max) = 63.83 Mbps
Total Throughput (Avg) = 48.32 Mbps	Total Throughput (Max) = 122.52 Mbps

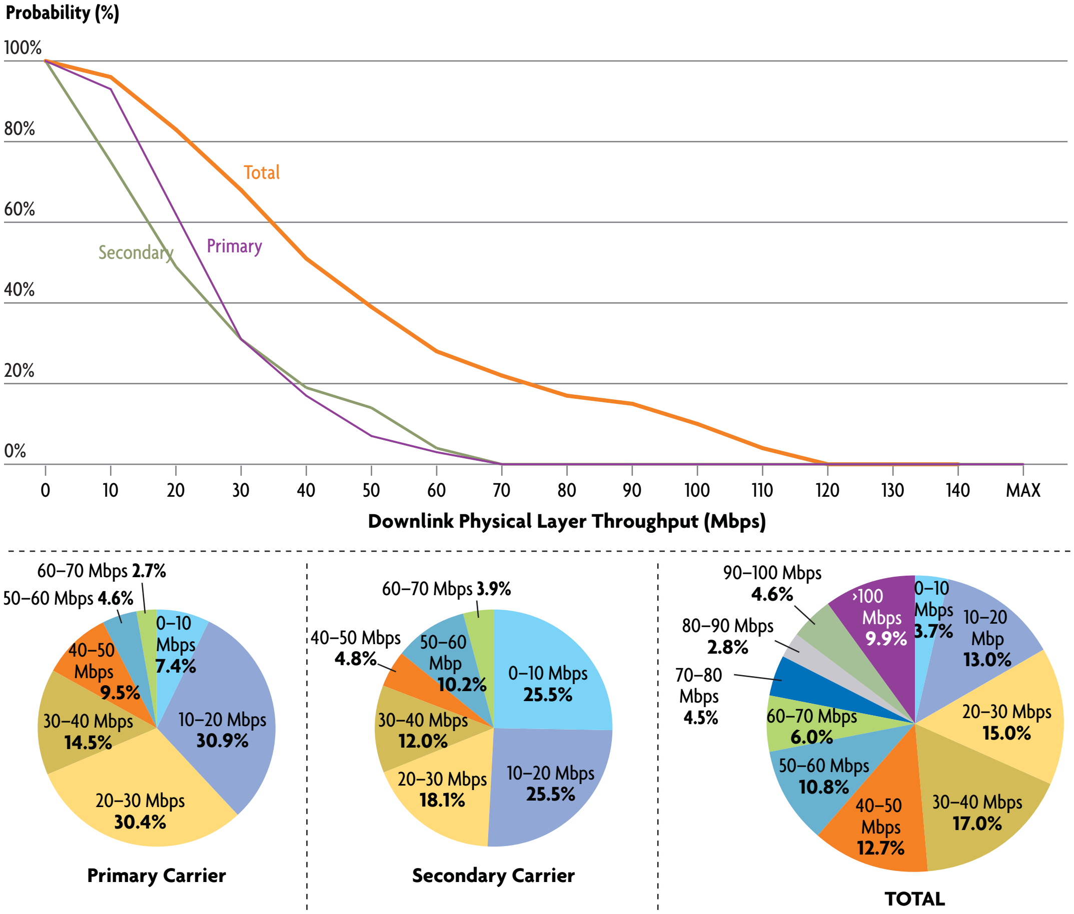

Source: Signals Research Group

Consistent with the Early AM drive test, during the 1815 hours drive test the primary carrier was at 2115 MHz and the secondary carrier was at 889 MHz. This information is implied in Figure 17 and Figure 18, in particular Figure 17 which shows that the secondary carrier had a more favorable RSRP distribution – consistent with the use of a lower frequency band.

Figure 17. RSRP by Primary and Secondary Carrier (1815 Hours) – Probability Distribution

RSRP - Primary (Avg) = -75.74 dBm RSRP - Secondary (Avg) = -72.06 dBm

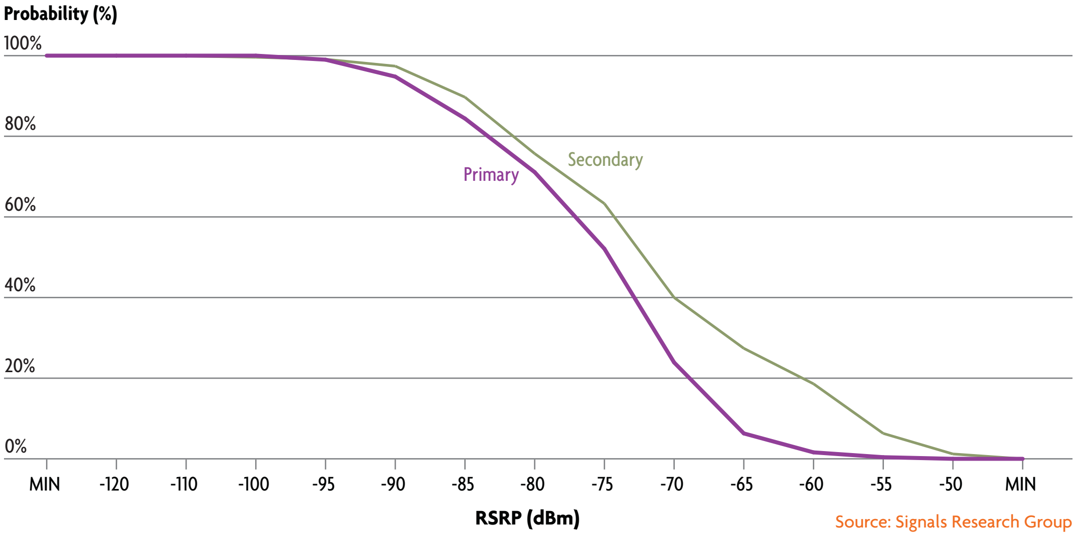

Source: Signals Research Group

Figure 18. SINR by Primary and Secondary Carrier (1815 Hours) – Probability Distribution

SINR - Primary (Avg) = 13.74 dB SINR - Secondary (Avg) = 10.47 dB

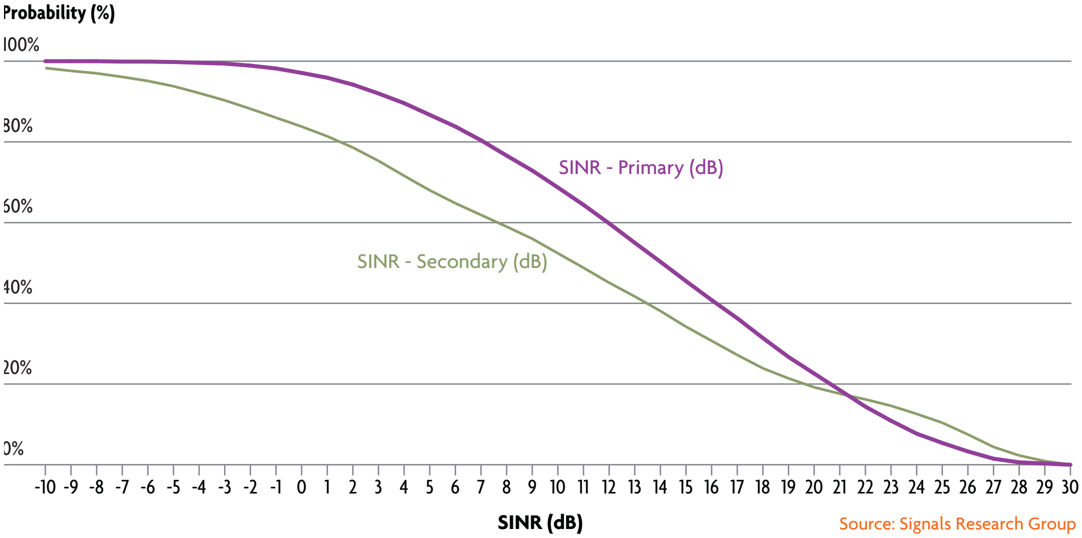

Source: Signals Research Group

Figure 19 provides probability distribution plots for the number of allocated RBs for the two radio carriers. Interestingly, the probability distribution was slightly better during this test than it was during the early AM drive tests.

Figure 19. Resource Block Allocation by Primary and Secondary Carrier (1815 Hours) – Probability Distribution

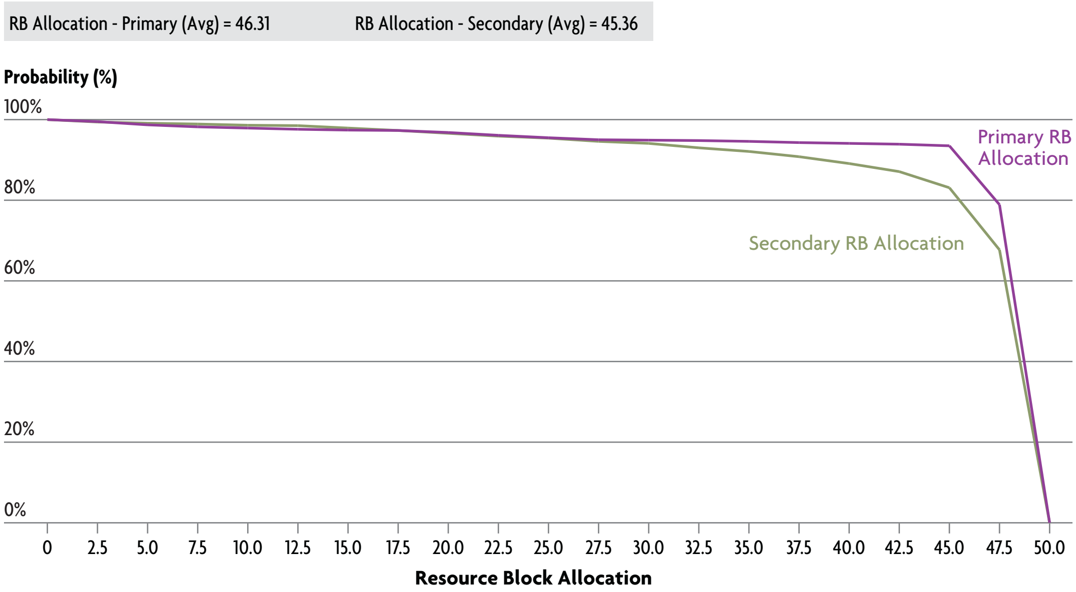

Source: Signals Research Group

The incremental benefits of a Category 4 device and MIMO were less apparent at 1815 hours than they were during the testing that occurred during the middle of the night.

Figure 20 provides the same set of information that we provided in Figure 7. Comparing the results between the two figures, it is evident that the incremental benefits of a Category 4 device were less apparent at 1815 hours than they were during the testing that occurred during the middle of the night. The Category 4 functionality was used less than 10% of the time and the incremental boost in throughput when the Category 4 device functionality was enabled was 13.0% (14.3% if we assume the Cat 4 device was needed to support a TBS value of 100,000 bits). Over the entire test, the benefits of the Category 4 device were inconsequential with less than a 2% boost in total throughput. Likewise, MIMO was used less frequently at 1815 hours than it was when the network was lightly loaded. Both results are consistent with what we would expect since more users in the network degrade the signal quality, which consequently limits the ability to use MIMO or send a larger amount of data in an individual subframe (e.g., a smaller TBS value).

During the 1815 hours test, the secondary carrier utilization was higher – 87.1% of the time the mobile device was using both radio carriers versus 79.5% during the 0313 hours drive test. It isn't clear why this was the case but since the low SINR in the secondary carrier was somewhat cell site specific, we suspect that the geographic locations of the two drive tests played a role.

Figure 20. Detailed Analysis of MIMO, Secondary Carrier and Category 4 Device Utilization (1815 Hours)

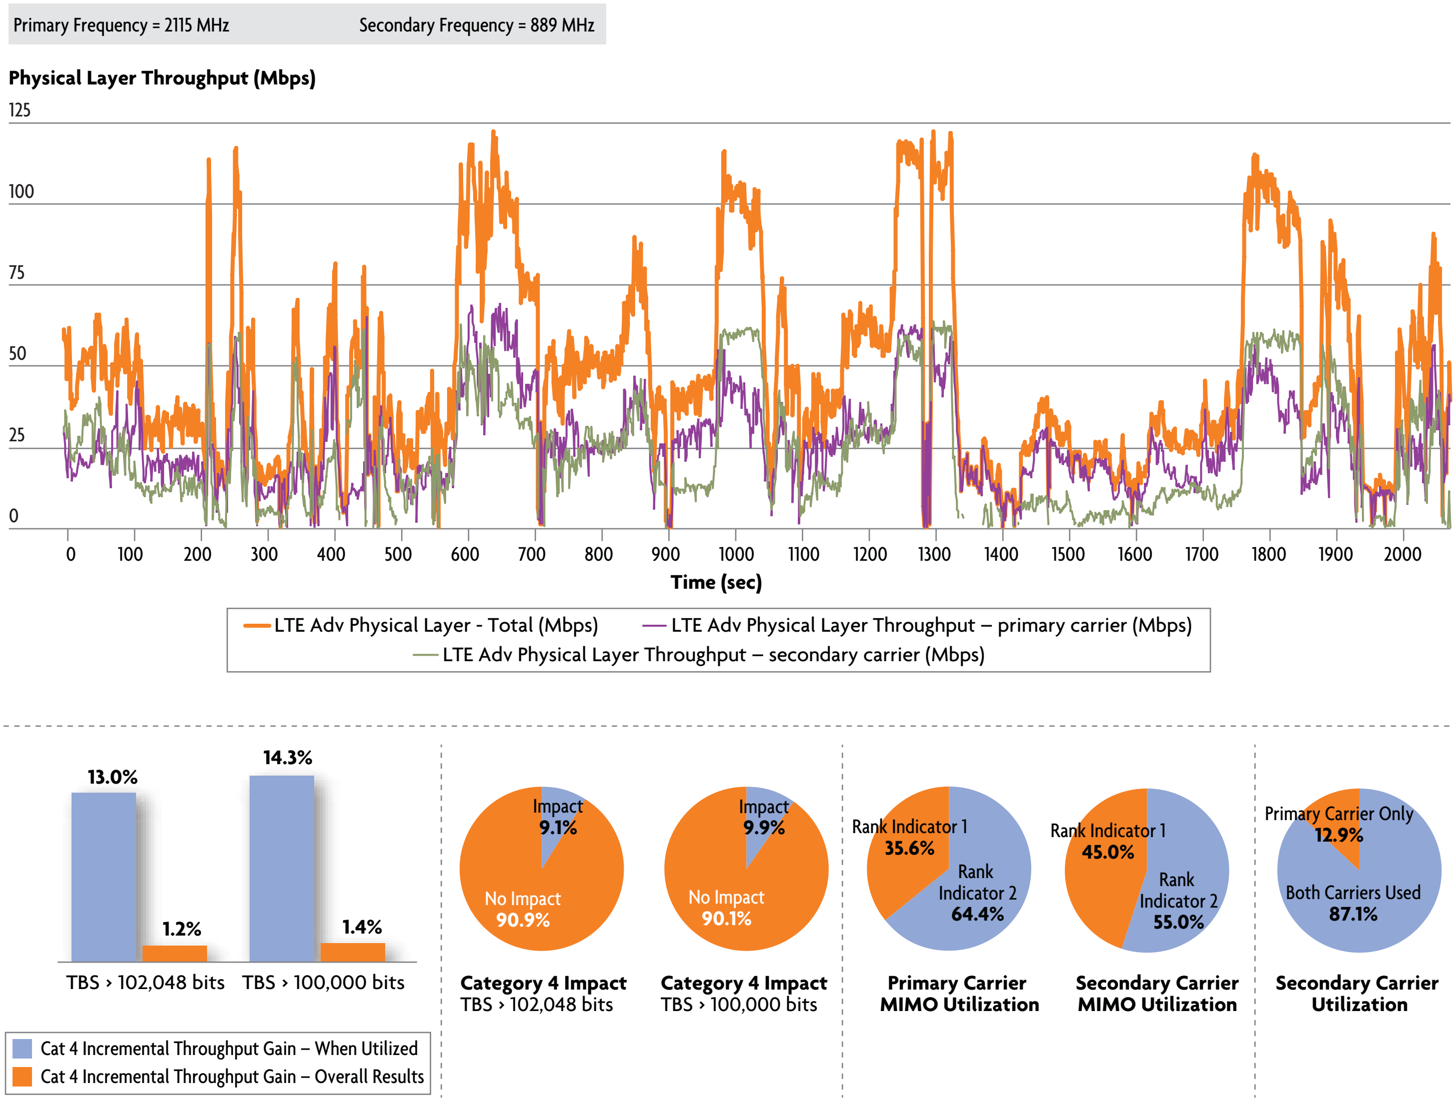

Source: Signals Research Group

The two carriers were using the same cells and that the handovers were occurring at the same time.

Figure 21 provides an enhanced view of the downlink throughput for the two radio carriers along with the corresponding serving cell PCI values. Once again, it is evident that the two carriers were using the same cells and that the handovers were occurring at the same time. We've highlighted an instance at roughly 1300 seconds into the drive test when the total throughput dropped appreciably and the contribution from the secondary radio carrier was nonexistent. In this case we know that the problem was not due to a cell handover and it wasn't due to low SINR (reference Figure 24). Conceivably, the phenomenon could have been due to other users in the network but the impact seems to be pretty dramatic. Further, the characteristics of the throughput are more consistent with another issue involving network + device behavior that we observed quite often during our testing.

Figure 21. Downlink Throughput by Primary and Secondary Carrier (1815 Hours) – Time Series

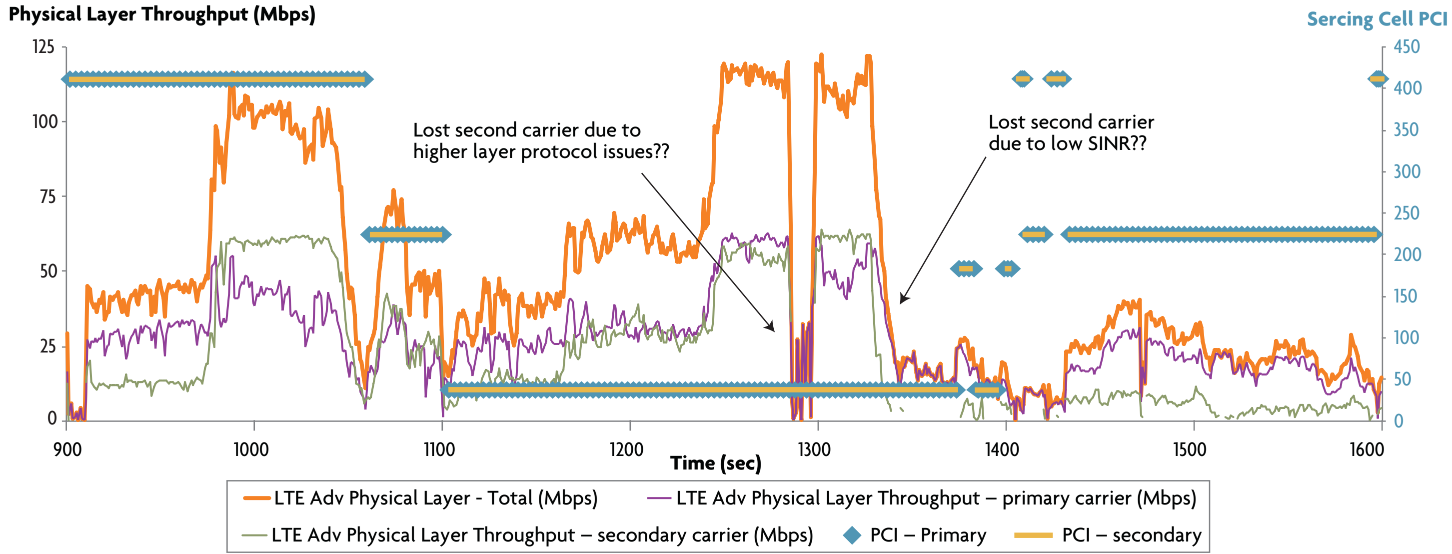

Source: Signals Research Group

Figure 22 and Figure 23 provide additional insight into the use of MIMO (code word 0 and code word 1) for both carriers.

Figure 22. Downlink Throughput by Primary Carrier with Individual Data Stream Contributions (1815 Hours) – Time Series

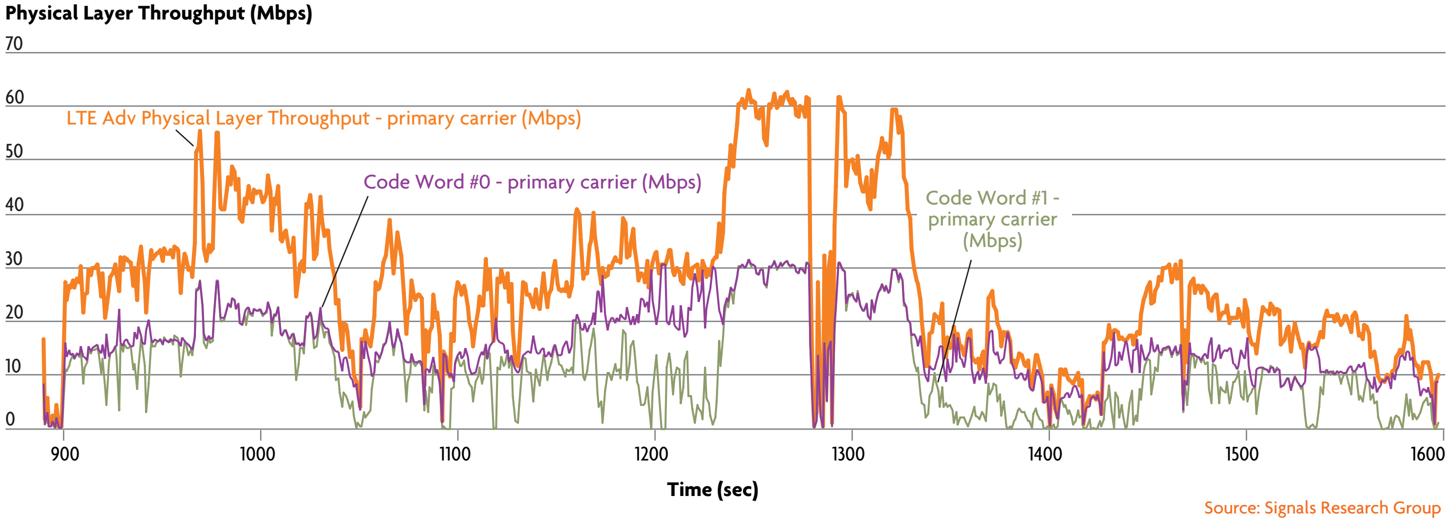

Source: Signals Research Group

Figure 23. Downlink Throughput by Secondary Carrier with Individual Data Stream Contributions (1815 Hours) – Time Series

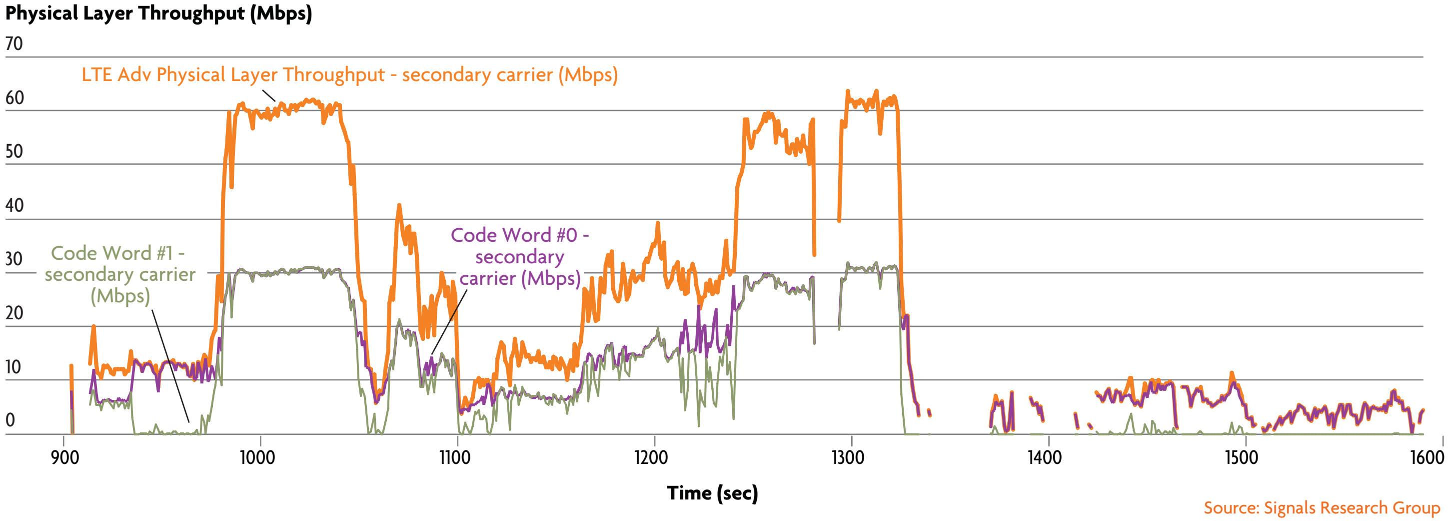

Source: Signals Research Group

We believe a higher layer protocol issue disrupted the secondary carrier throughput.

Initially, we used simultaneous FTP and UDP data streams to load the pipe. When we took this approach we observed the frequent loss of the secondary carrier and quite often the loss occurred in conjunction with the FTP or UDP application requesting another download. As part of our test methodology, we downloaded multiple files at a time and due to the relatively small data files on the host servers, combined with the bandwidth potential of the network, we had to periodically add new data files to the download queue. When the FTP or UDP application requested a new download the interruption in the throughput would occur even though we had other FTP/UDP data sessions running in parallel. In other words, there was, or at least should have been, plenty of data in the scheduler so the interruption in the data stream, which sometimes lasted for tens of seconds, had to have been due to something else. We now know that the low SINR in the secondary carrier played a role but we believe a higher layer protocol issue in the device or the network could exist. When we stopped using UDP the issue was far less prevalent but it still existed, as shown in the figure.

Figure 24. SINR and Resource Block Allocation by Primary and Secondary Carrier (1815 Hours) – Time Series

LTE Adv SINR - primary carrier = 13.74 dB	LTE Adv SINR - secondary carrier = 9.60 dB
LTE Adv RB Allocation - primary carrier = 46.3	LTE Adv RB Allocation - secondary carrier = 45.4
Primary Frequency = 2115 MHz	Secondary Frequency = 889 MHz

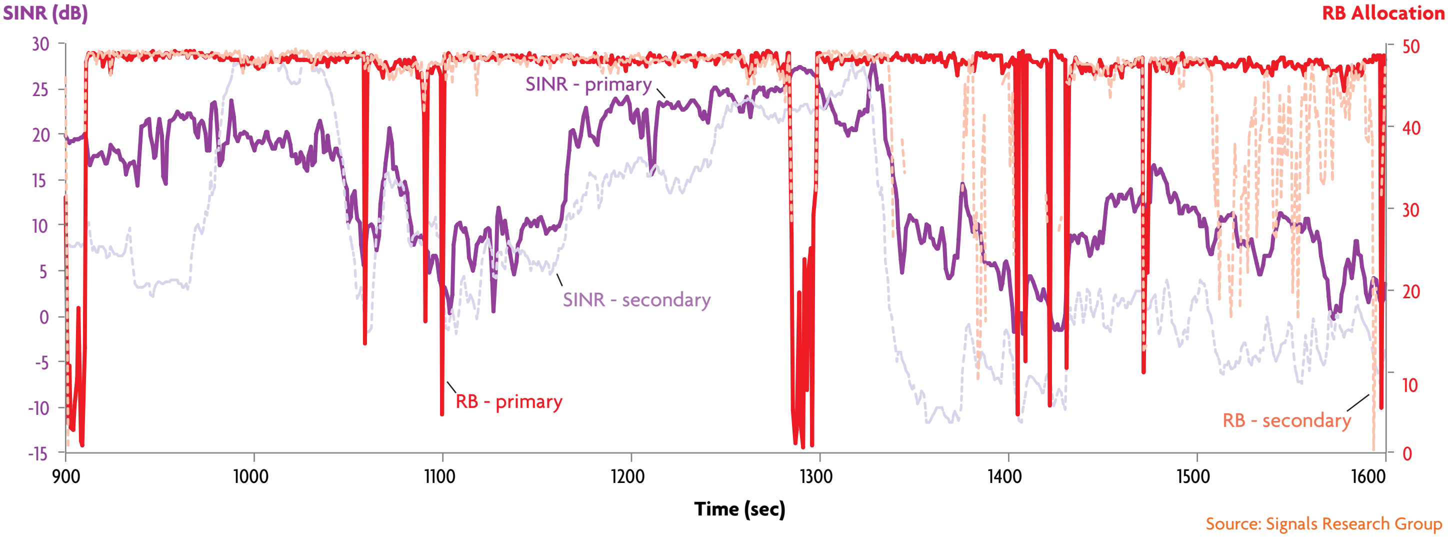

Source: Signals Research Group

Figure 25 provides scatter plots of the SINR and the corresponding downlink throughput for the two radio carriers.

Figure 25. SINR versus Downlink Throughput – primary and secondary carriers (1815 Hours) – Scatter Plots

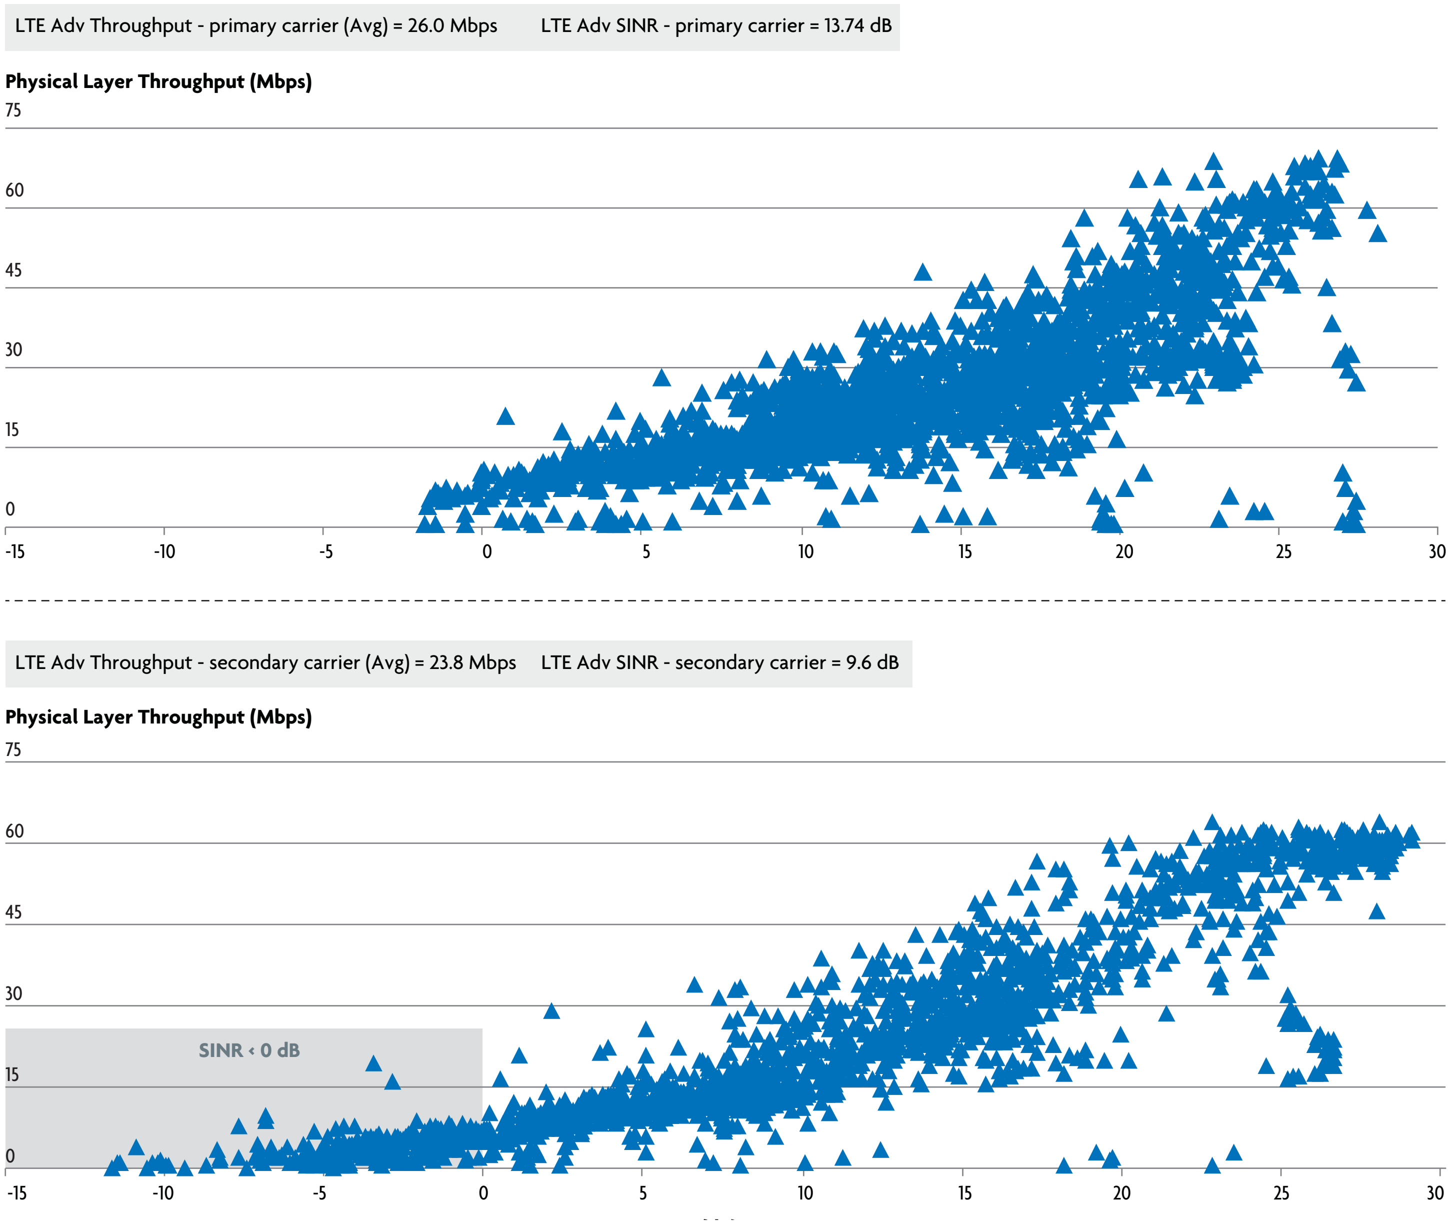

Source: Signals Research Group

3.3 0906 Hours Drive Test

For completeness sake we are including detailed analysis for one more drive test. This drive test occurred midmorning (0906 hours). The test lasted a relatively short 12.2 minutes, during which time we downloaded 5.2 GB. Figure 26 provides a geo plot of the 2.22 mile drive route.

Figure 26. 0906 Hours AM Drive Test Routes

Source: Signals Research Group

As shown in Figure 27, it was during this test that we observed the highest downlink throughput, or 141 Mbps. The throughput was recorded at 09:14:37 hours to be exact.

Figure 27. Downlink Throughput by Primary and Secondary Carrier (0906) – Probability Distribution

Primary Throughput (Avg) = 32.42 Mbps
Secondary Throughput (Avg) = 26.28 Mbps
Total Throughput (Avg) = 58.38 Mbps
Primary Throughput (Max) = 71.39 Mbps
Secondary Throughput (Max) = 71.58 Mbps
Total Throughput (Max) = 140.96 Mbps

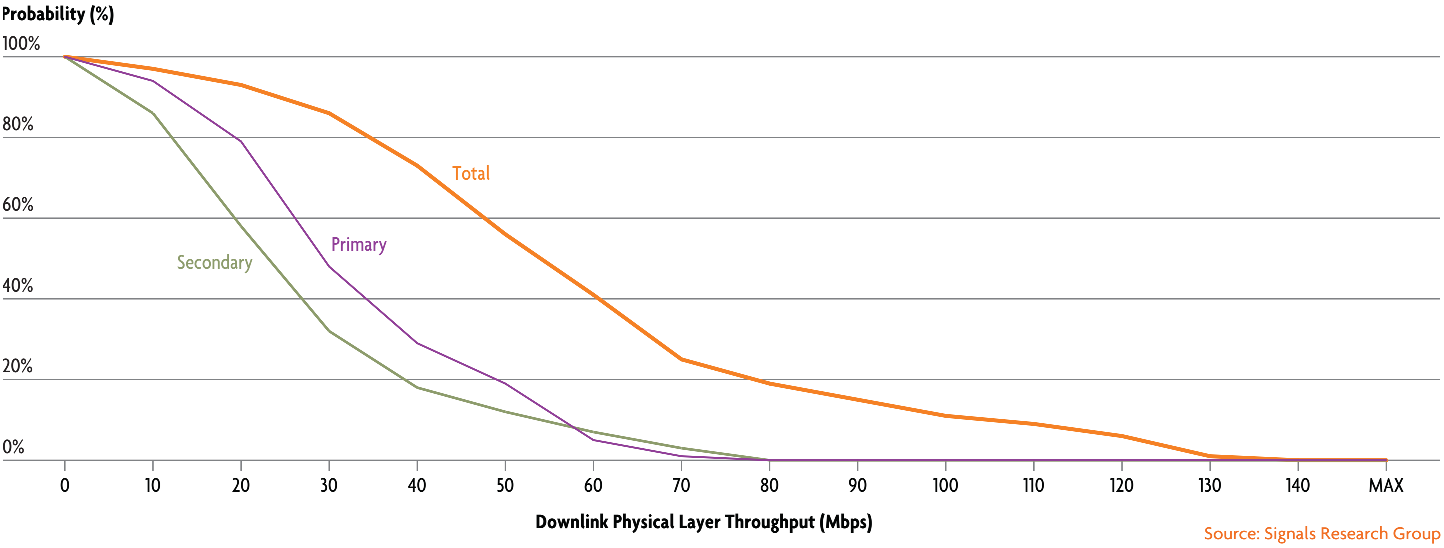

Source: Signals Research Group

Figure 28 provides the RSRP probability distribution plots for the primary and secondary carriers and Figure 29 provides comparable information for the SINR.

Figure 28. RSRP by Primary and Secondary Carrier (0906 Hours) – Probability Distribution

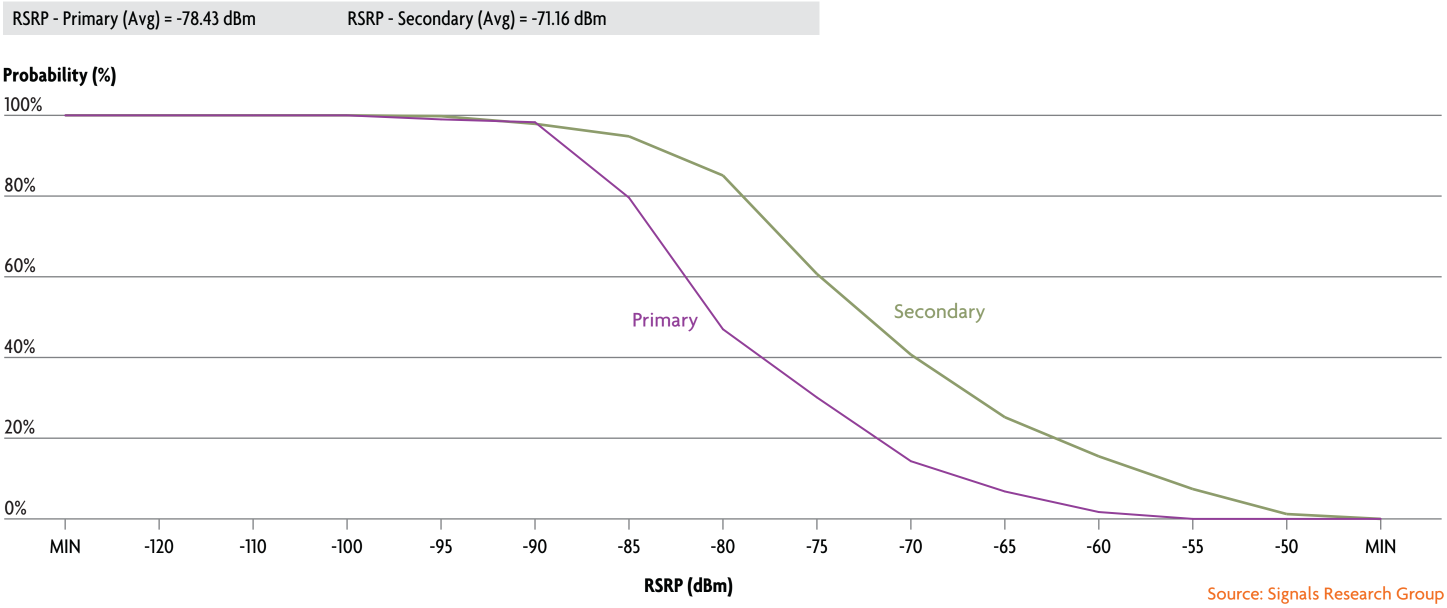

Source: Signals Research Group

Figure 29. SINR by Primary and Secondary Carrier (0906 Hours) – Probability Distribution

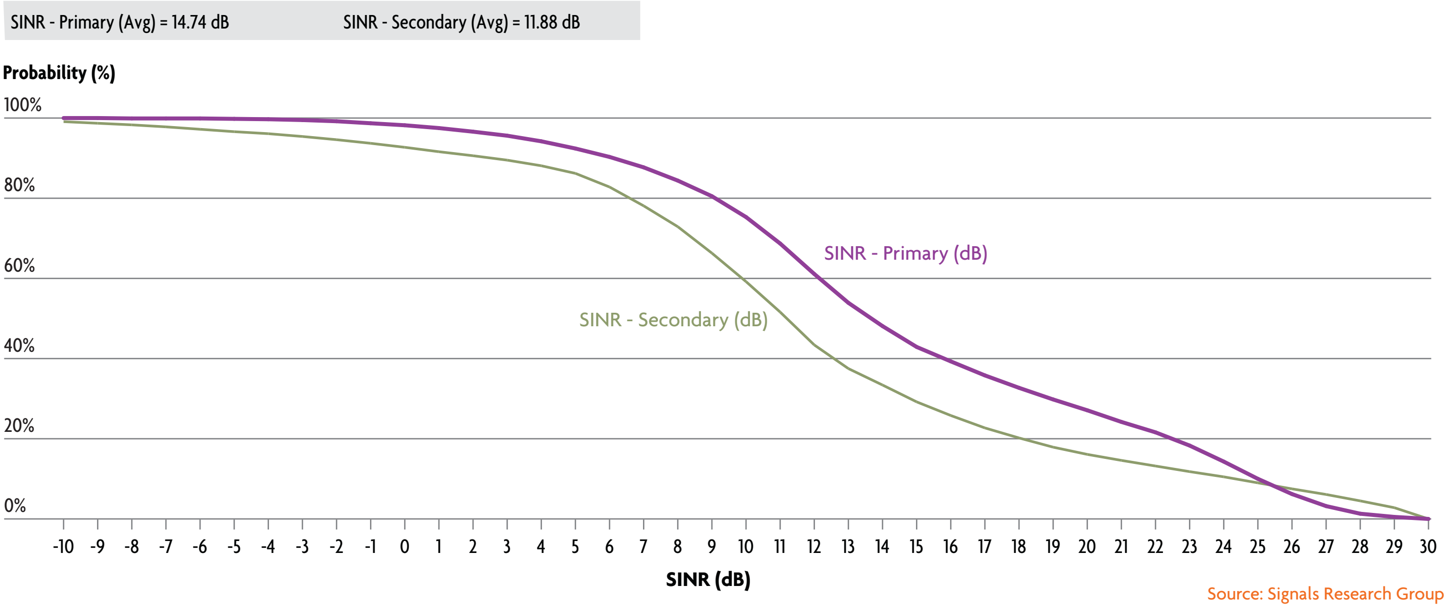

Source: Signals Research Group

Figure 30 highlights additional information about some of the underlying KPIs. In this relatively short drive test, both carriers were used 92.5% of the time. The benefits of the Category 4 device were also fairly meaningful, in particular with respect to the calculated boost in throughput when the Category 4 functionality was being used.

Figure 30. Detailed Analysis of MIMO, Secondary Carrier and Category 4 Device Utilization (0906 Hours)

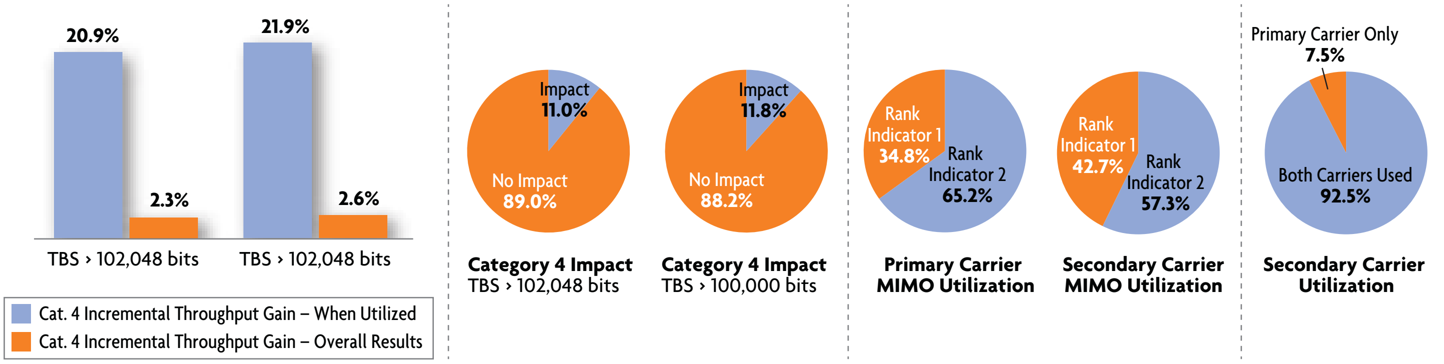

Source: Signals Research Group

3.4 Simultaneous LTE Advanced Carrier Aggregation and LTE Release 8 Drive Test

In this section we provide results from a drive test in which we tested two LTE modems running in parallel on two different notebook computers. One LTE modem supported LTE Advanced and the other LTE modem was limited to Release 8 functionality. Figure 31 provides a geo plot of the 3.6 mile drive route. During the 22.4 minute test we downloaded 10.44 GB of data – 6.95 GB with the Release 10 modem and 3.49 GB with the Release 8 modem.

Figure 31. Simultaneous LTE Advanced and LTE Release 8 Drive Test Routes

Source: Signals Research Group

Figure 32 and Figure 33 provide time series plots of the throughput. For the LTE Release 10 modem we include the total throughput and the individual contributions from the primary and secondary carriers. In Figure 33, which is a time enhanced version of Figure 32, we also include the serving cell PCI values for the two modems. Throughout the entire test, the Release 8 modem used the radio channel at 889 MHz while the Release 10 modem started with the primary carrier using 889 MHz before switching the primary carrier to 2115 MHz approximately 150 seconds into the test.

Figure 32. Downlink Throughput by LTE Advanced (Total, Primary and Secondary) and Release 8 Carriers – Time Series

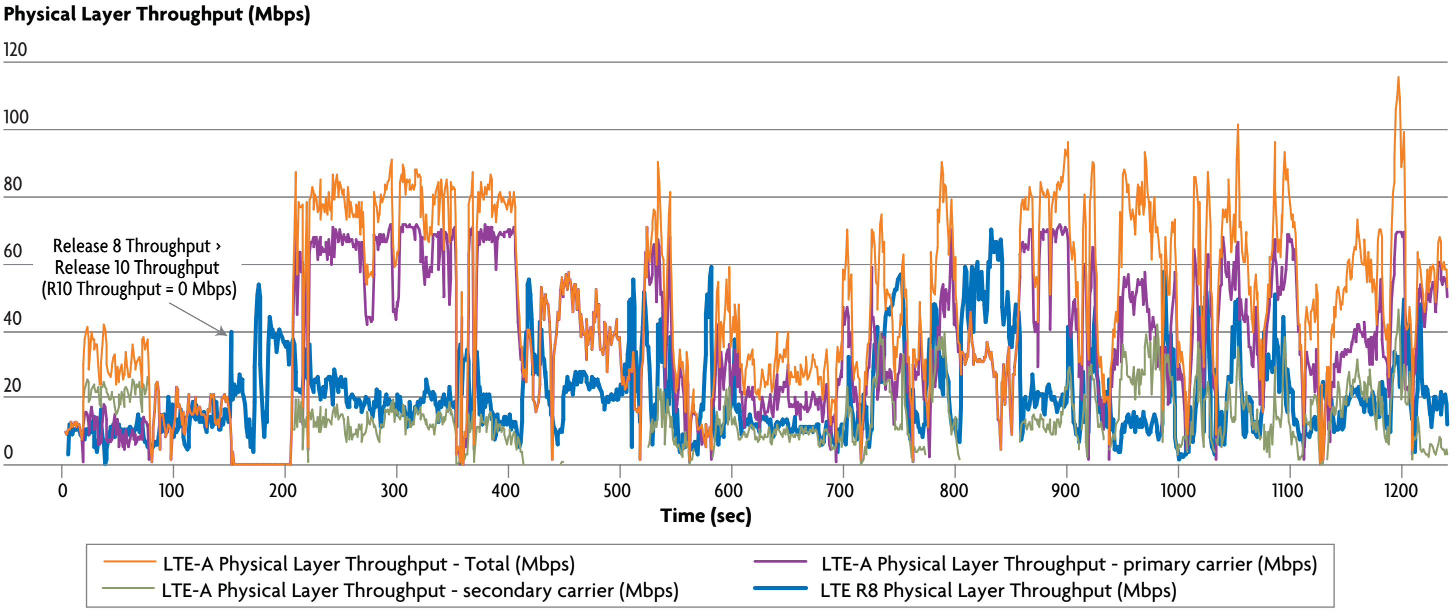

Source: Signals Research Group

Figure 33. Downlink Throughput by LTE Advanced (Total, Primary and Secondary) and Release 8 Carriers with Serving Cell PCI Values and Carrier Frequency Assignments – Time Series

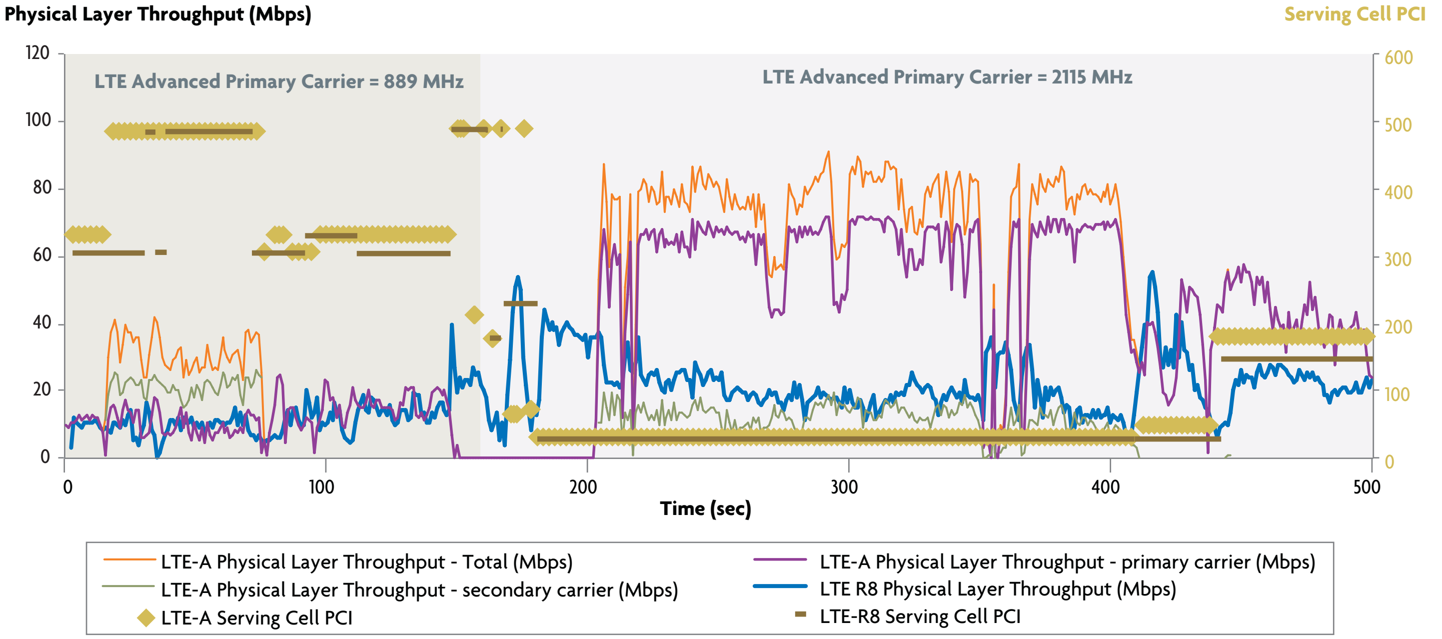

Source: Signals Research Group

Carrier aggregation is all about getting the most that the network can deliver under all circumstances.

Among other things, Figure 33 shows that the two modems didn't always share the same serving cell during the drive test. This isn't too surprising given the density of the network and we've seen this situation numerous times in the past when testing multiple devices other networks. More importantly, the figure shows perhaps the key benefit of carrier aggregation. The Release 8 modem remained at 889 MHz throughout the test even though the carrier at 2115 MHz had the better signal quality (SINR) and potentially the ability to deliver higher throughput. Conversely, the Release 10 modem used both carriers simultaneously so it didn't matter which carrier delivered the higher throughput. Carrier aggregation is all about getting the most that the network can deliver under all circumstances.

Finally, Figure 34 provides probability distribution plots for the combined LTE Advanced + LTE Release 8 throughput. The average throughput for the two devices operating in a combined channel bandwidth of 20 MHz (10 MHz + 10 10 MHz) was an impressive 71.4 Mbps. However, it is worth reminding readers that this value is somewhat overstated since there were moments during the drive test when the two dongles were using different serving cells. The average throughput for the LTE Release 10 modem was also 126% higher than the average throughput for the Release 8 modem.

Figure 34. Combined Downlink Throughput for LTE Advanced and Release 8 – Probability Distribution

LTE Advanced Throughput (Avg) = 49.5 Mbps LTE Release 8 Throughput (Avg) = 21.9 Mbps
Combined LTE Adv and LTE Rel 8 Throughput (Avg) = 71.4 Mbps

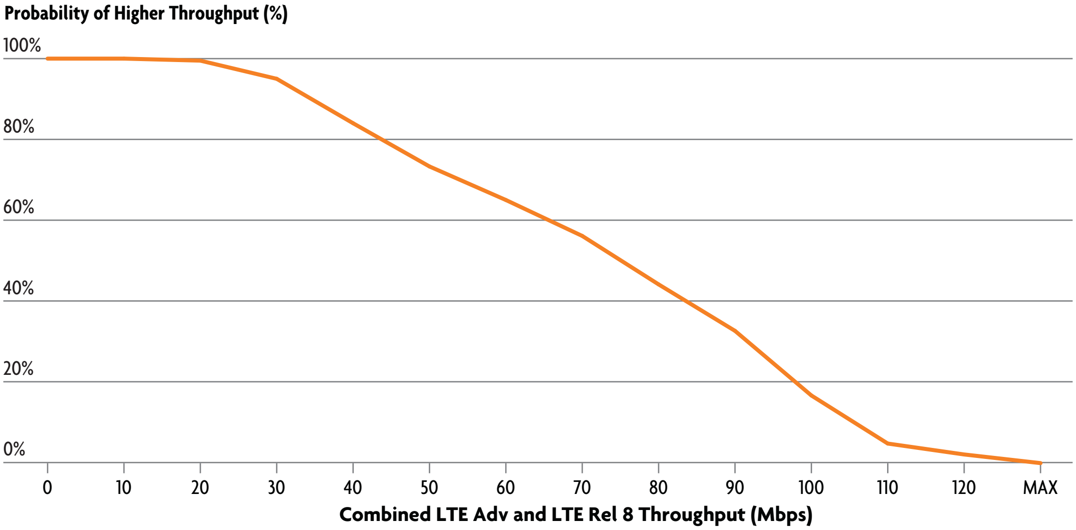

Source: Signals Research Group

4.0 Uplink Drive Tests – Detailed Analysis and Commentary

LTE Advanced Carrier Aggregation is a downlink feature so there isn't any impact on the uplink performance. However, given the overall stellar performance of the network we wanted to take the opportunity to see how the uplink performed. Needless to say we were not disappointed. In fact, to some extent we were more impressed, or at least equally impressed, with the uplink throughput as we were with the downlink throughput.

Figure 35 provides a geo plot of the 3.9 mile drive route that we used for the uplink drive test. This particular test started at 0747 on August 29th and it lasted for 20.9 minutes.

Figure 35. Uplink Drive Test Route

Source: Signals Research Group

90.6% of the time the uplink throughput was higher than 10 Mbps.

Figure 36 provides the probability distribution plot for the Physical Layer uplink throughput and the same information in a pie chart format. The average uplink throughput was a very impressive 20.21 Mbps. Nearly 96% of the time (95.7% to be exact) the uplink throughput was higher than 5 Mbps and for 90.6% of the time the uplink throughput was higher than 10 Mbps. For comparison purposes, when we tested AT&T's LTE network in Houston (pre-commercial launch) the average uplink throughput was 15.2 Mbps.

Figure 36. Uplink Throughput – Probability Distribution and Pie Charts

Uplink Physical Layer Throughput (Avg) = 20.21 Mbps
Uplink Physical Layer Throughput (Max) = 22.77 Mbps

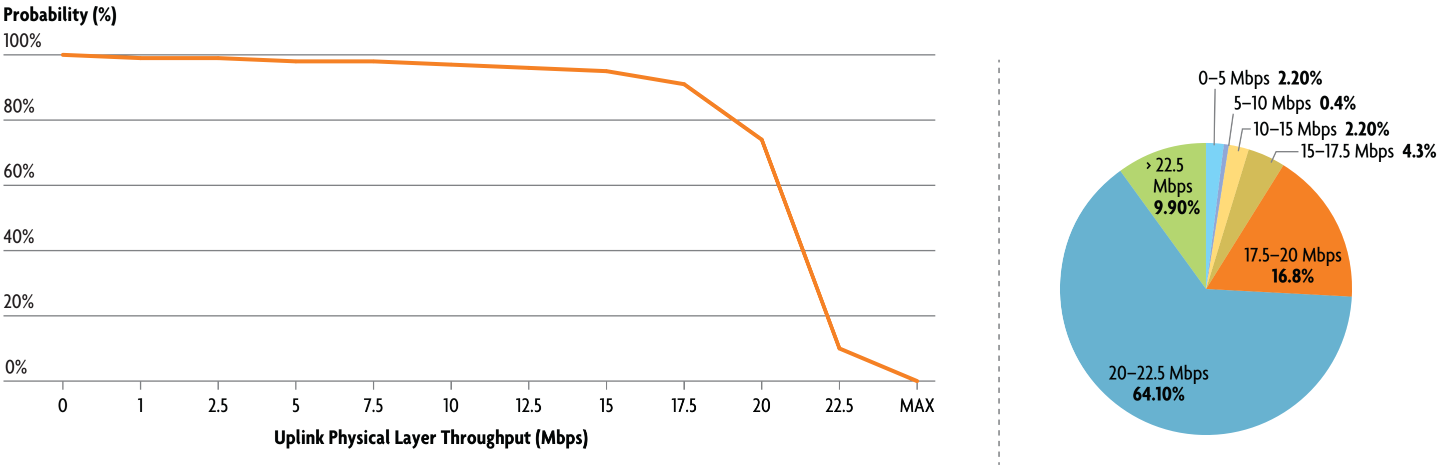

Source: Signals Research Group

With a maximum of only 45 PUSCH RBs available in the uplink, the maximum throughput was a bit lower than in other networks that we have tested.

Figure 37 provides a probability distribution plot for the transmit power and Figure 38 provides a probability distribution plot for the number of assigned uplink resource blocks. Since the uplink resource block allocation was limited to a maximum of 45 RBs, the peak uplink data rate was somewhat lower than what we have observed in other networks. We assume that the operator was using four PRBs for the uplink control channels in order to support a higher number of concurrent mobile

Figure 37. PUSCH Transmit Power (Uplink Drive Test) – Probability Distribution

PUSCH Transmit Power (Avg) = 10.89 dBm
Frequency = 1925 MHz

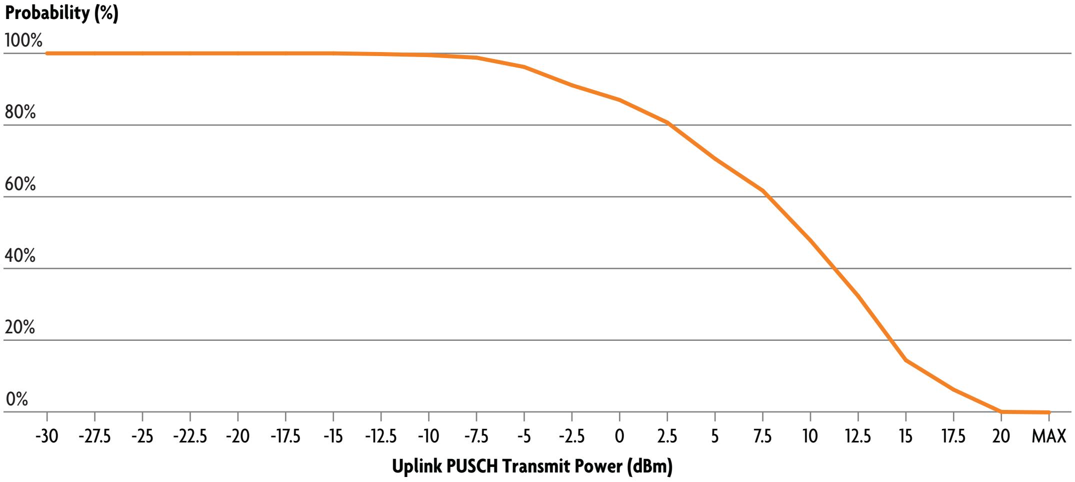

Source: Signals Research Group

data users. Trading off uplink control channels for additional PUSCH RBs would have increased the maximum throughput but it would have negatively impacted the number of concurrent mobile data users in the network. Given the high average uplink data rates and the large number of mobile data subscribers in the network, we can't complain.

Figure 38. Uplink Resource Block Allocation (Uplink Drive Test) – Probability Distribution

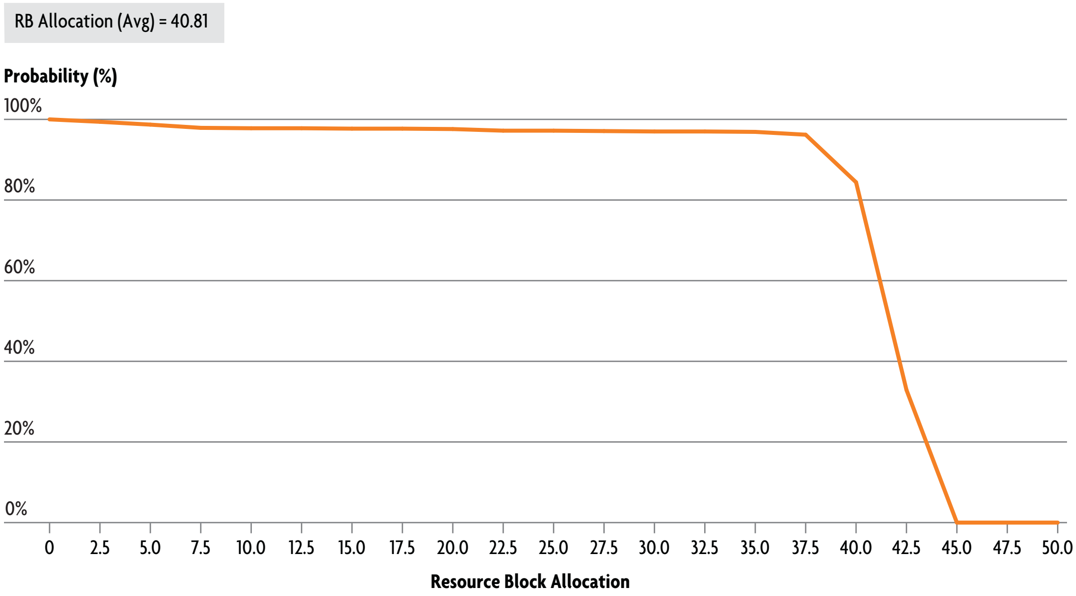

Source: Signals Research Group

The following set of figures provide a wealth of information pertaining to the relationships between the uplink throughput, the transmit power, and the number of allocated RBs. The first figure (Figure 39) is from the testing that we did in Seoul. The last three figures are from the testing that we did in Tokyo. Figure 40 shows the LTE TDD results, Figure 41 shows the LTE FDD 2 x 10 MHz results and Figure 42 shows the LTE FDD 2 x 5 MHz results.

In all of the figures, the three dimensional scatter plot shows the uplink Physical Layer throughput (color coded) as a function of the number of allocated RBs (X axis) and the average transmit power (Y axis). Note that we used a slightly different color scheme and X axis values for the LTE TDD and LTE FDD 2 x 5 MHz results. Further, we did not adjust the LTE TDD transmit power for the duty cycle – with TDD the transmitter is not turned on when the mobile device is receiving data. The information in the lower left-hand corner is perhaps the most interesting. It provides an efficiency metric in the form of the average transmit power per RB – grouped by uplink throughput values. We also include the average number of allocated RBs for each grouping of throughput values. The pie chart in the lower right-hand corner provides probability information for each grouping of throughput values.

A lower transmit power per RB is preferable, assuming that the average throughput is similar.

For example, one can see in Figure 39 that the average uplink throughput was higher than 20 Mbps for 74.5% of the time in the Seoul network. This grouping of throughput values required an average allocation of 42.3 RBs with an average transmit power of 0.255 dBm per RB. Conversely, in the LTE FDD 2 x 10 MHz network in Japan (reference Figure 41), the average uplink throughput was in the range of 15 – 20 Mbps for 80.6% of the time. This grouping of throughput values required an average of 35.4 RBs with an average transmit power of 0.120 dBm per RB. The lower transmit power per RB in the Tokyo network is preferable to the higher value in Seoul, but since the average throughput was also lower in the Tokyo network, the mobile device would have to transmit for a longer period of time to upload the same amount of data.

Figure 39. Detailed Analysis of Uplink Physical Layer Throughput versus Transmit Power versus Resource Block Allocation (South Korea)

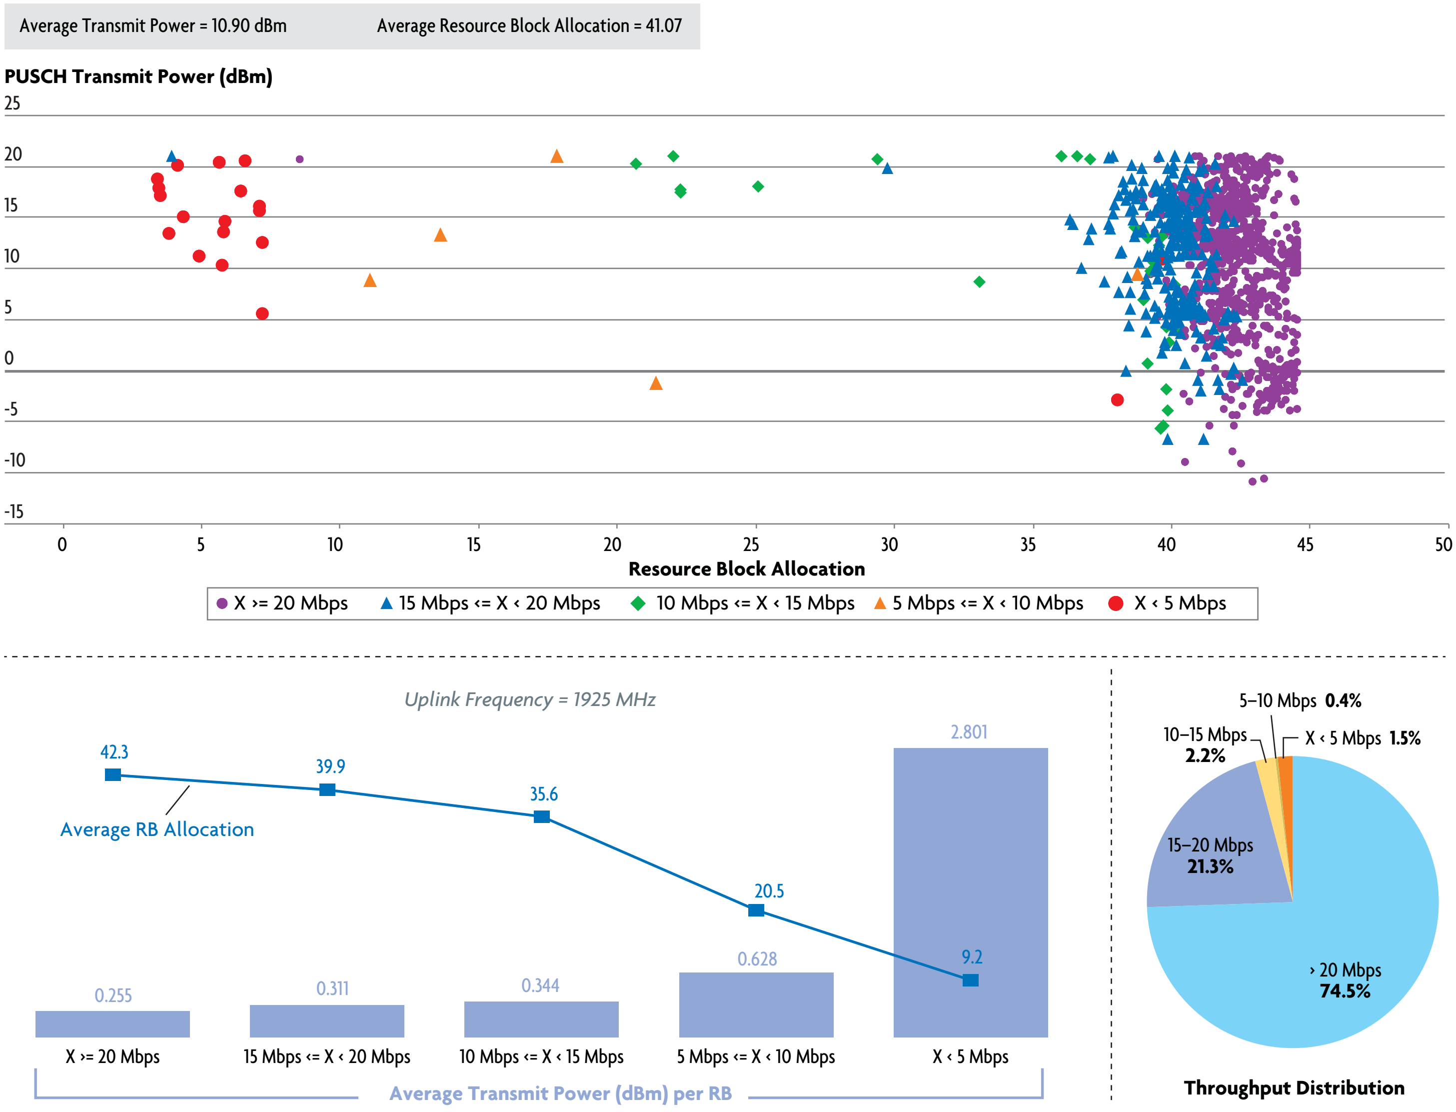

Source: Signals Research Group

From a power efficiency perspective, the LTE FDD 2 x 5 MHz network in Tokyo was "lights out fantastic" while the LTE TDD network was impacted by the use of a higher frequency.

Another way to look at the data is to compare the LTE TDD results (reference Figure 40) with the Seoul results (reference Figure 39). If we focus just on the resource requirements associated with achieving an uplink throughput of 5 – 10 Mbps, one sees that the average transmit power per RB was nearly twice as high in the LTE TDD network as it was in the network in Seoul. Although the network in Seoul also used 41% more RBs to achieve a comparable throughput, it still appears to us that the LTE FDD network in Seoul was more efficient with respect to power utilization. Clearly the use of a lower frequency band helped in that regard. For the same grouping of throughput values, the LTE FDD 2 x 5 MHz network in Japan was "lights out fantastic" with an average transmit power of 0.071 dBm per RB and a comparable number of allocated RBs.

Figure 40. Detailed Analysis of Uplink Physical Layer Throughput versus Transmit Power versus Resource Block Allocation (LTE TDD - Japan)

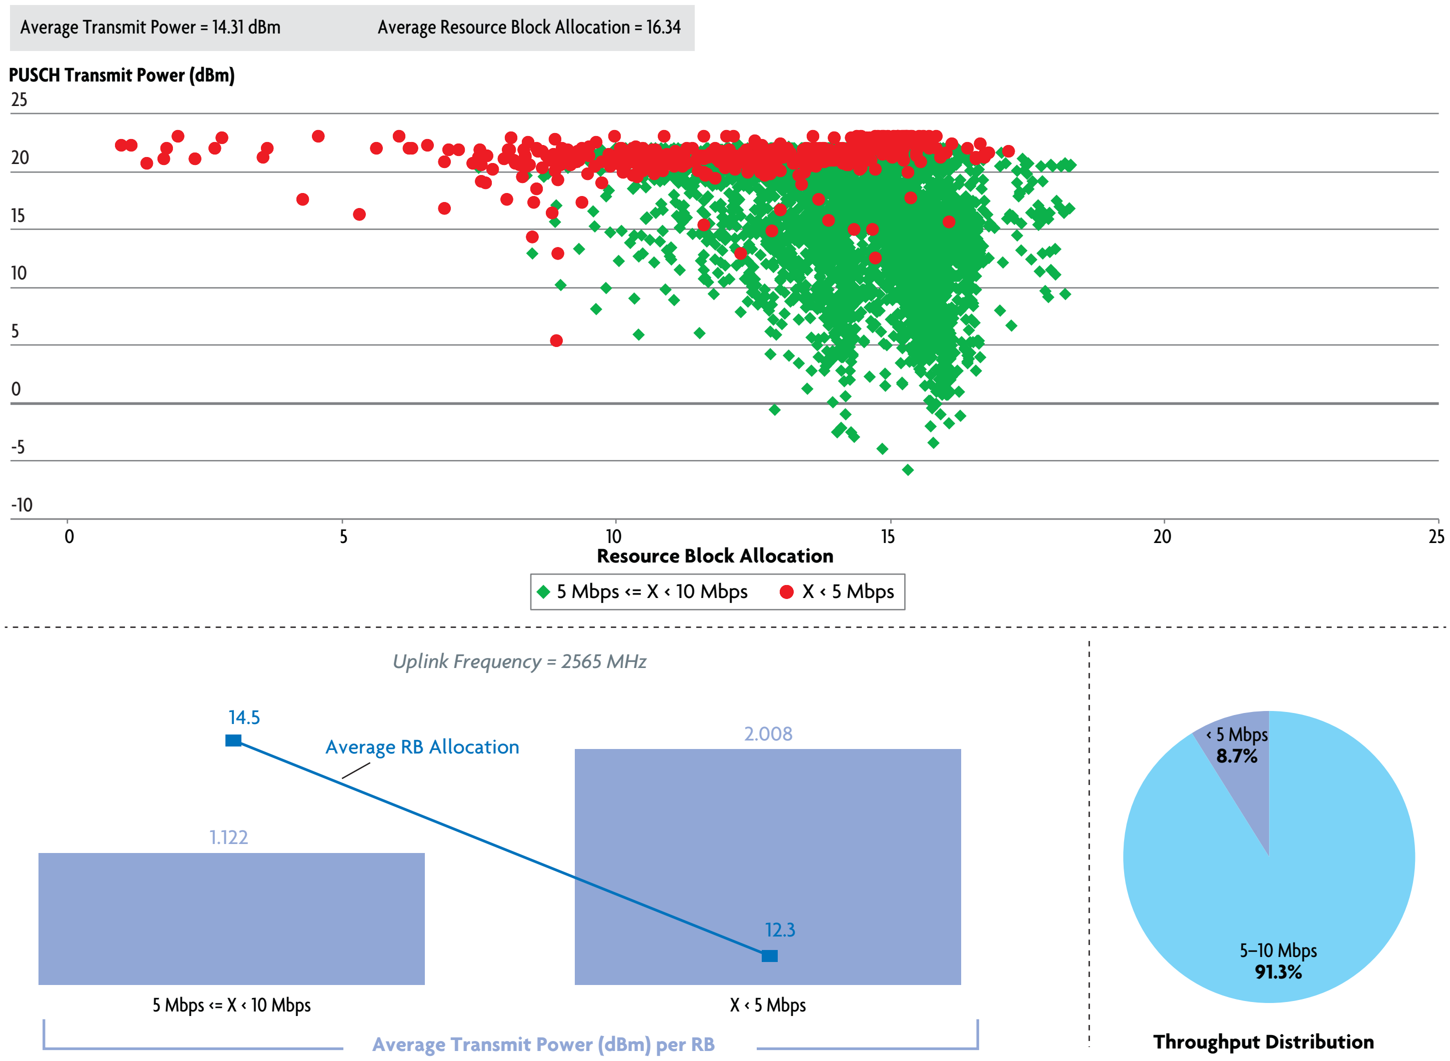

Source: Signals Research Group

Figure 41. Detailed Analysis of Uplink Physical Layer Throughput versus Transmit Power versus Resource Block Allocation (LTE FDD 10 MHz - Japan)

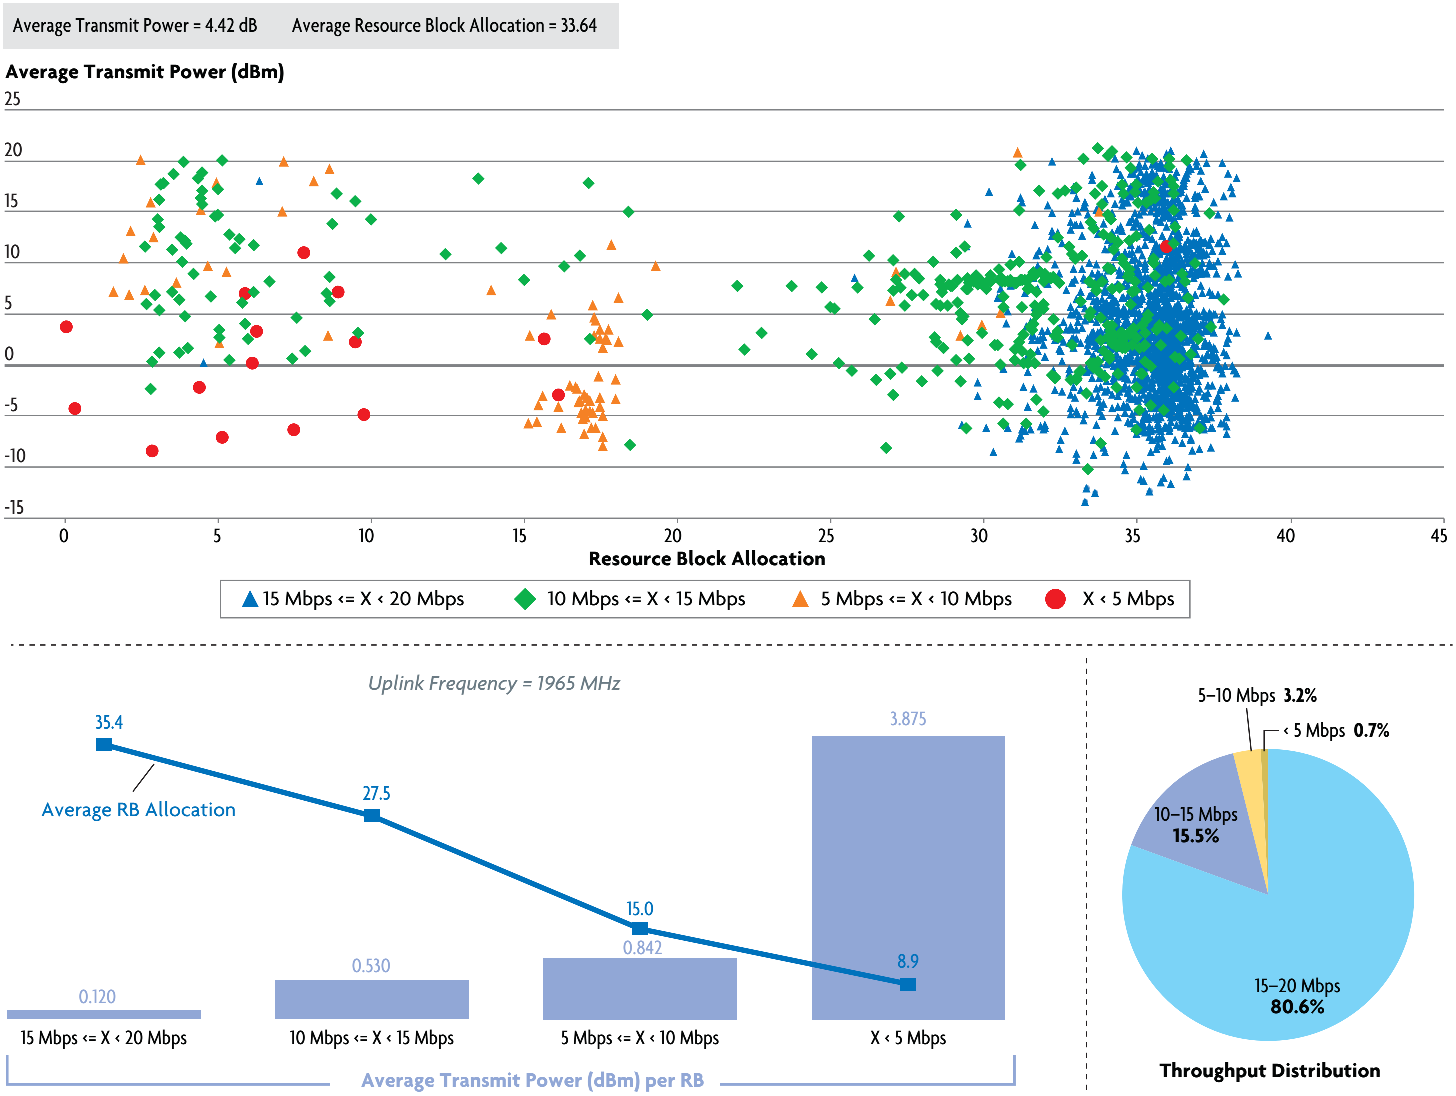

Source: Signals Research Group

Figure 42. Detailed Analysis of Uplink Physical Layer Throughput versus Transmit Power versus Resource Block Allocation (LTE FDD 5 MHz - Japan)

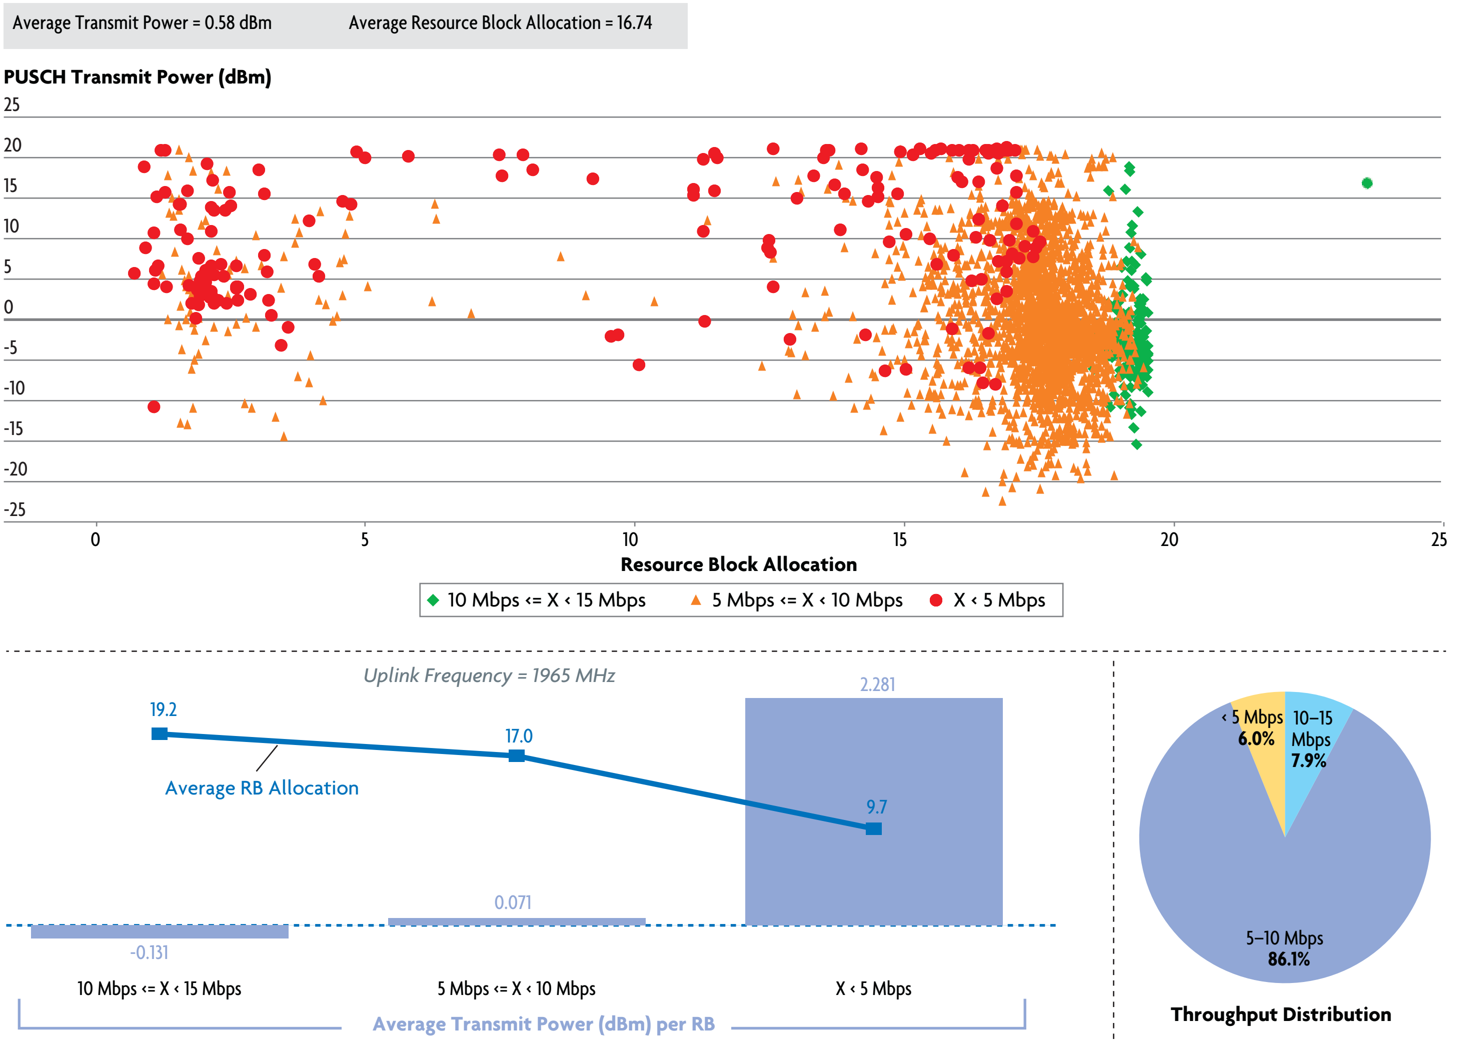

Source: Signals Research Group

PART II – USER EXPERIENCE TESTS

TEST RESULTS AND TEST METHODOLOGY

5.0 User Experience Tests – Detailed Analysis and Commentary

In addition to the basic downlink and uplink throughput tests, we conducted a series of user experience tests to determine how carrier aggregation impacted the user experience. For the user experience testing, we primarily focused on web browsing, using the capabilities of Spirent's Datum application, a solution that it brought in-house when it acquired Metrico Wireless.

5.1 Web Browsing

For the web browsing tests we selected several web sites that we felt were popular in South Korea. We made the selection of web sites based on the results of several Google searches. All of the web browsing tests occurred from stationary positions, including our hotel (fixed wireline) or while parked along the city streets.

We ultimately zeroed in on four websites: NATE, English.Gmarket.kr, Naver, and Chosun. For each test we loaded each web site 25 times. When we analyzed the results we observed more than a few instances when the web page load time was meaningfully higher (e.g., 30 seconds versus 3 seconds) than the preceding and subsequent values. This situation was also prevalent in the wireline test results from the hotel room test so we believe that the problem wasn't necessarily network related. We elected to remove these values before averaging the results that we present in this section. With sufficient time and energy we probably could have identified the root cause but given that it didn't happen all that frequently and it impacted all results, we elected to not pursue the matter at this time.

When conducting the web browsing tests, we used the Datum application to automatically load each web page.

Figure 43 presents the results from the tests that we conducted. When conducting the web browsing tests, we used the Datum application to automatically load each web page. The application would wait for the web page to fully download before moving on to the next web page. After loading all four web pages, the test suite automatically repeated the test for 25 iterations. In addition to "managing the browser" the Datum application recorded how long it took to load each web page and transmitted this information to a server that we could easily access to get the information that we needed.

The Notebook PC – LTE Advanced (Avg) and Notebook PC – wired connection (Avg) results are based on the averages of two different test periods (50 loads per web page) and four different test periods, respectively. Unfortunately, the other test results are not as interesting as we would have liked. In one test period involving the head-to-head competition of a Release 8 smartphone and a Release 10 smartphone, the Release 10 smartphone did not use carrier aggregation for reasons that we do not fully understand, meaning that the Release 10 smartphone behaved as if it were a Release 8 smartphone. In another test, we thought we were testing a Release 8 smartphone and a Release 10 smartphone but it turned out that the Release 8 smartphone was actually a Release 10 smartphone. Both smartphones used carrier aggregation to load the web pages during the test so we didn't get any good comparative data.

The next time we conduct user experience testing involving web pages we will make it the primary focus of the study instead of an afterthought. Still, it is possible to extract some interesting information from the results that we collected.

Figure 43. Web Page Load Times – by access network technology

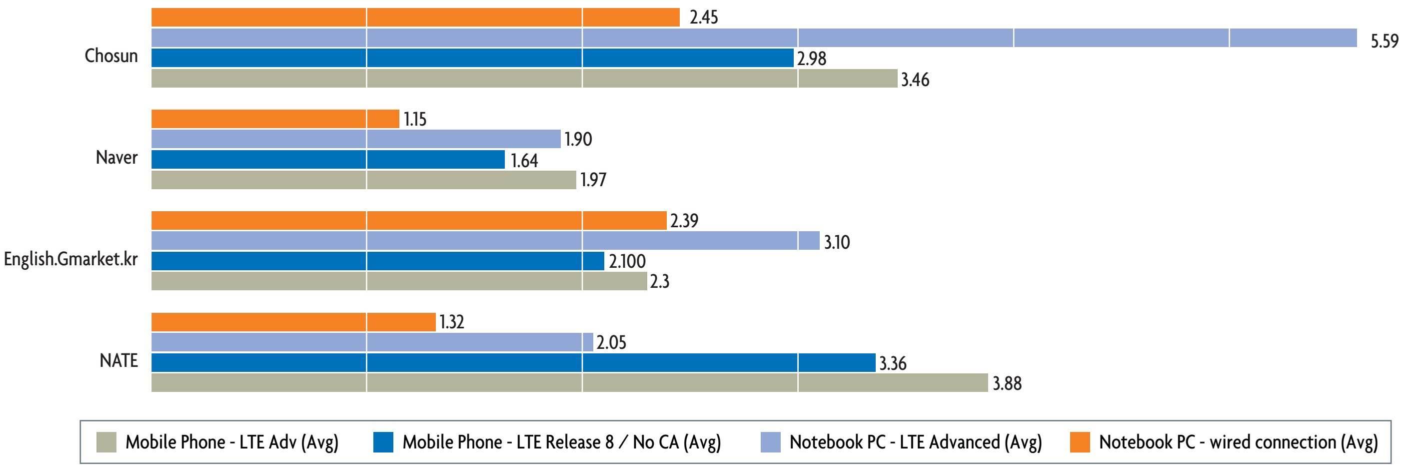

Source: Signals Research Group

The wireline network (presumably fiber) delivered the best web browsing user experience.

First, it is apparent that despite the stellar performance of the LTE Advanced network, the wireline network (presumably fiber) delivered the best web browsing user experience. With the exception of the Chosun website, the LTE Advanced network performed reasonably well with the notebook computer. In the Chosun results there were several results that were somewhat higher than the norm but we elected to leave the results in the log file when calculating the average. Excluding these quasi-outliers would have lowered the average load time to slightly below 5 seconds versus the 5.59 second value that we show in the figure.

Figure 44 shows the first 80 seconds of a web browsing test involving a notebook computer that was tethered to a Release 10 smartphone. The figure shows the TBS allocations for the primary and secondary carriers. For each carrier, the figure also shows the TBS allocations for each code word. With 2 x 2 MIMO there are two possible code words with a single code word (code word #0) used if the network reverts back to transmit diversity. Among other things, the figure shows that carrier aggregation was used when appropriate and when necessary when transmitting data to the mobile device. Although it isn't readily visible in the figure, we noticed in the raw data that the TBS values

Figure 44. Web Page Browsing with a Notebook Computer and its Impact on Resource Allocation – per TTI subframe Time Series

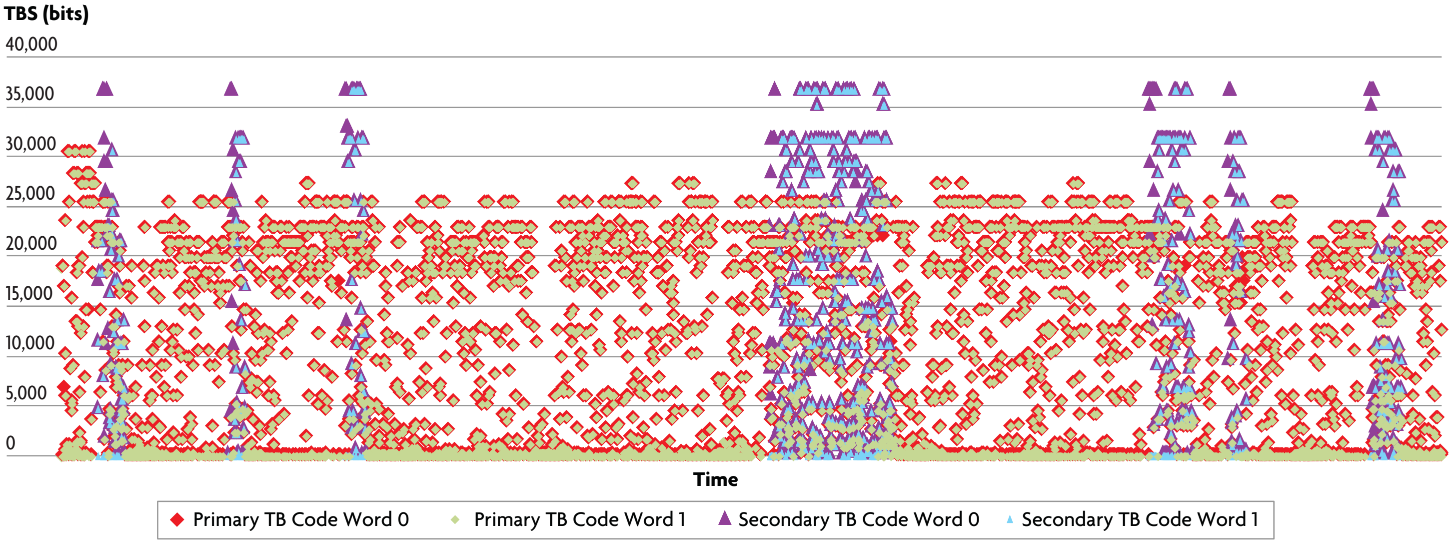

Source: Signals Research Group

for both code words were always identical. The only exception was those instances when transmit diversity was used, in which case the TBS value for code word 1 was 0 bits.

The ability to schedule data on either carrier in a near-seamless fashion is more compelling than the ability to schedule data on both carriers currently.

Figure 45 and Figure 46 provide information on the use of carrier aggregation and MIMO during one of the Notebook PC – LTE Advanced web browsing tests. As shown in Figure 45, both carriers were used concurrently for 20.14% of the time. We assume both radio carriers were used when there was enough data in the scheduler's buffer to warrant its use. The interesting observation in this figure is that the secondary carrier was used by itself (i.e., no primary carrier) for 3.7% of the time. In many respects, the ability to schedule data on either carrier in a near-seamless fashion is more compelling than the ability to schedule data on both carriers currently. A lot of data transactions involve snacking and relatively small amounts of data, in which case the use of both radio carriers may not be triggered. However, as networks become more loaded it will become more important to schedule data transmissions on the most appropriate carrier, based on which carrier can deliver the traffic in the most efficient manner. This capability is good for the operator in the form of increased network capacity and it is good for the consumer because it delivers a better user experience. Most importantly, these benefits are realized even if both carriers are not used at the same time.

The bottom four pie charts provide information about the use of MIMO. The top two pie charts show what the mobile device requested and the bottom two pie charts indicate what the network scheduled. The secondary carrier results only include those instances when the secondary carrier was active. It isn't surprising that the actual use of MIMO (i.e., Rank Indicator 2) was lower than the device-requested use of MIMO since the network will not use MIMO if there isn't sufficient data in the scheduler's buffer to justify its use. Depending on the vendor's implementation, the network scheduler may also take a less aggressive stance toward the use of MIMO relative to what the mobile device requests.

Figure 45. Carrier and MIMO Rank Indicator 2 Utilization while Web Browsing with a Notebook Computer

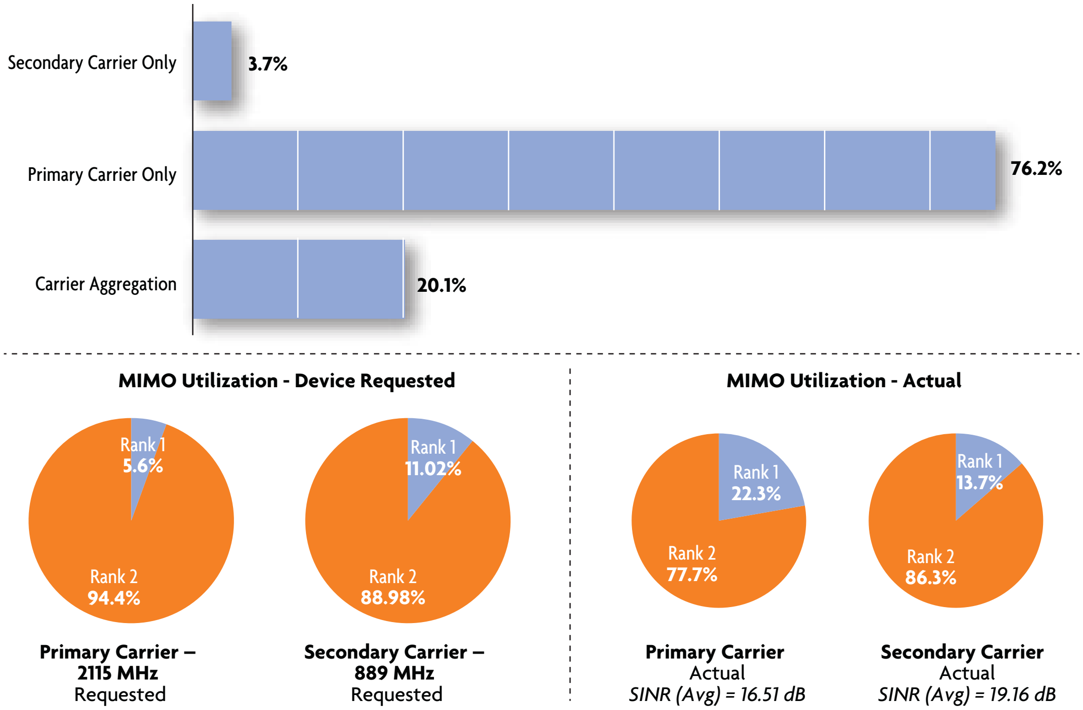

Source: Signals Research Group

Although we didn't come back with good data that would allow us to directly compare Release 8 and LTE Advanced web browser performance, we can still look at how carrier aggregation behaved while browsing with a LTE Advanced smartphone. Figure 47 and Figure 48 are identical in nature to the two previous figures that we just discussed. Comparing the two sets of figures, it is evident that both radio carriers were used more frequently with the smartphone browser (31.31%) than with the notebook PC browser (20.14%). This finding is seemingly counter-intuitive since one would assume that the distribution of TBS values for the notebook computer would contain a greater mix of larger TBS values. In fact, we prove this hypothesis in Figure 49.

Figure 47 shows the use of both radio carriers as a function of the TBS value. Obviously, both the primary and secondary carriers had to be used when the TBS value was greater than 75,000 bits since it isn't possible to send 75,000 bits in one sub-frame of a 10 MHz LTE channel. The figure also shows that as the TBS value decreases there was a lower probability that both carriers were used to carry the payload. This observation is intuitive and hardly surprising but the figure does back up the point with hard numbers.

The behavior of carrier aggregation will change based on the mobile data application that is being used.

For comparison purposes, recall that during the 0906 drive test the TBS value was greater than 100,000 bits for 21.9% of the time (reference Figure 30) versus only 2.7% during the web browsing test. In the same 0906 drive test the TBS was less than 25,000 bits for 10% of the time versus 42.3% of the time in this web browsing test. We raise these points to illustrate that applications have varying bandwidth requirements and different distributions of data packet sizes. It is, therefore, only natural that the behavior of carrier aggregation will change based on the mobile data application that is being used.

Figure 46. Carrier Aggregation Utilization while Web Browsing with a Notebook Computer – as a function of TBS

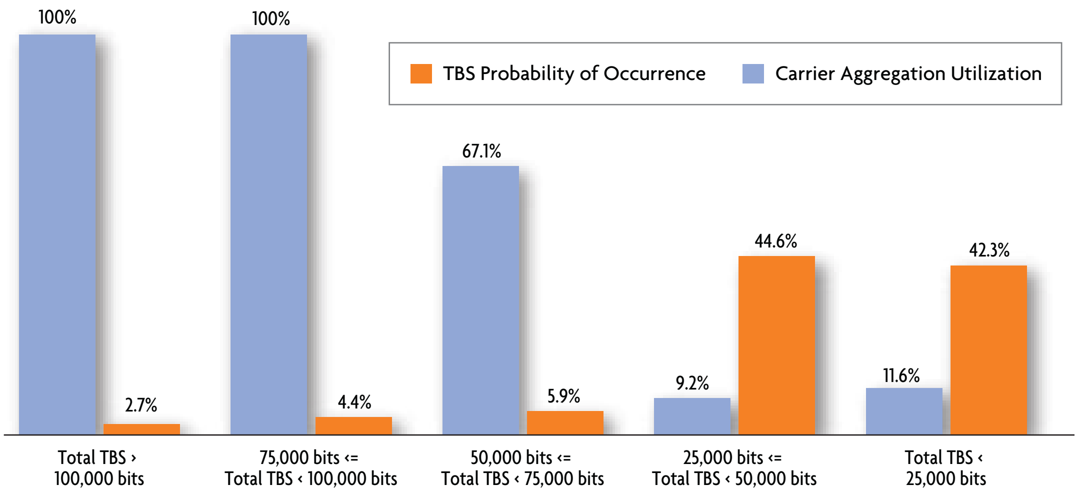

Source: Signals Research Group

Although we didn't come back with good data that would allow us to directly compare Release 8 and LTE Advanced web browser performance, we can still look at how carrier aggregation behaved while browsing with a LTE Advanced smartphone. Figure 47 and Figure 48 are identical in nature to the two previous figures that we just discussed. Comparing the two sets of figures, it is evident that both radio carriers were used more frequently with the smartphone browser (31.31%) than with the notebook PC browser (20.14%). This finding is seemingly counter-intuitive since one would assume that the distribution of TBS values for the notebook computer would contain a greater mix of larger TBS values. In fact, we prove this hypothesis in Figure 49.

Figure 47. Carrier and MIMO Rank Indicator 2 Utilization while Web Browsing with a Smartphone Carrier Aggregation Utilization while Web Browsing with a Smartphone

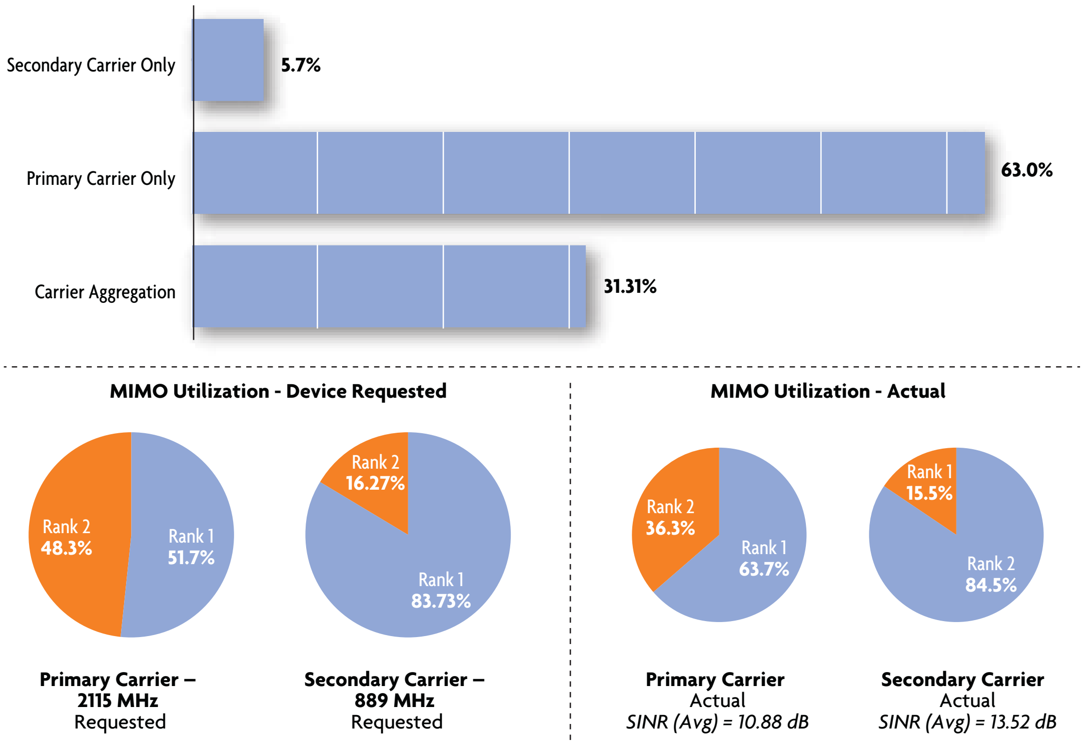

Source: Signals Research Group

We believe that the differences in the MIMO utilization between the two tests partially explain what is happening. During the notebook computer web browsing tests the network assigned MIMO for 77.7% of the time with the primary carrier and 86.3% of the time with the secondary carrier. During the smartphone web browsing tests the MIMO utilization rate was only 48.3% and 36.3%, respectively. Our conclusion is that the greater dependence on MIMO during the notebook computer web browsing tests resulted in the less frequent use of both radio carriers. During the smartphone web browsing tests, MIMO wasn't used as frequently – largely due to the poorer network conditions – so there was a greater opportunity to send data using both radio carriers. With web browsing and the relatively small data packets there is only so much data that can be transmitted at any point in time.

Figure 48. Carrier Aggregation Utilization while Web Browsing with a Smartphone – as a function of TBS

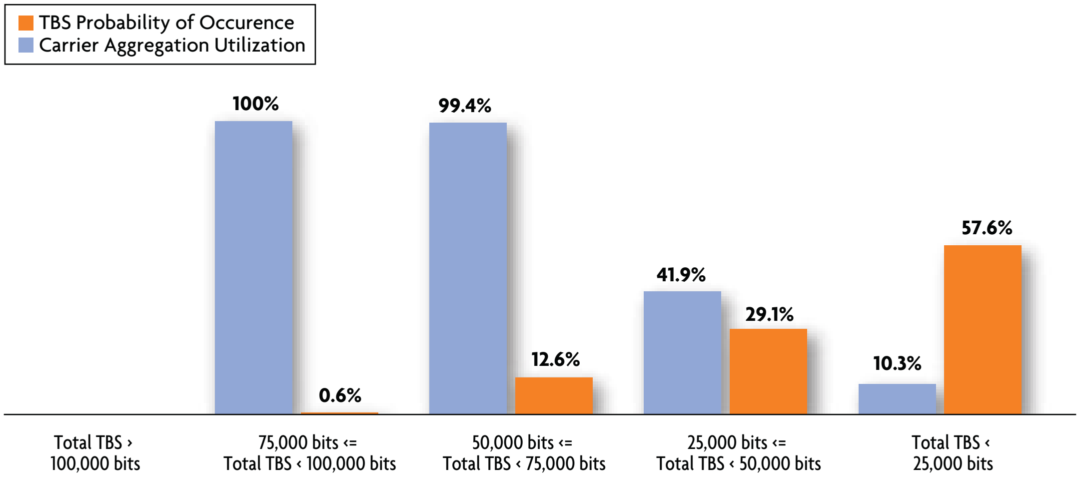

Source: Signals Research Group

Figure 49. A Comparison of TBS Allocations – notebook computer versus smartphone

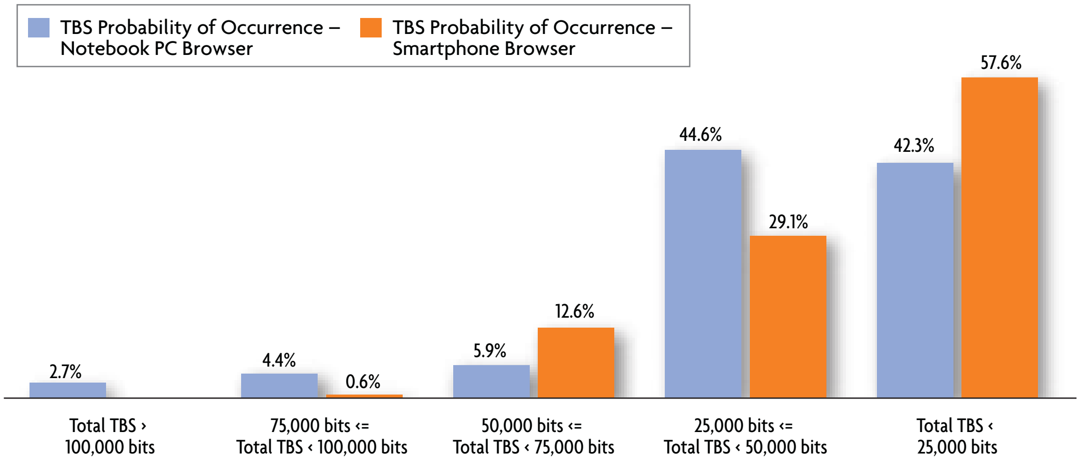

Source: Signals Research Group

Figure 50 shows results from another smartphone web browsing test. In this case the LTE Advanced smartphone didn't use carrier aggregation. Further, the smartphone also didn't request, nor was it assigned, MIMO, even though the channel conditions would have supported its use. We can't explain why this situation took place but it did prevent us from collecting "good data" that would have allowed us to compare the web browsing performance between LTE Advanced and Release 8 smartphones. Interestingly, the Release 8 smartphone also didn't request MIMO during this test and the channel conditions that it reported were even better than the conditions reported by the LTE Advanced smartphone.

Figure 50. Carrier Aggregation and MIMO Rank Indicator 2 Utilization while Web Browsing with a Smartphone – as a function of TBS

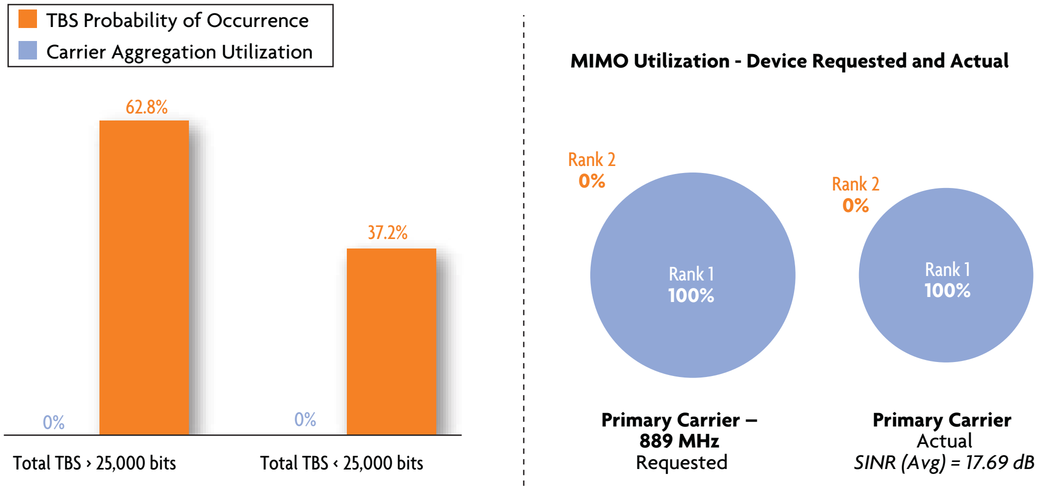

Source: Signals Research Group

5.2 FTP Server versus Speedtest.net versus Google Play

We conducted a series of tests from a stationary location where we had fairly ideal network conditions – in particular, given that we were testing mid-morning. The test consisted of downloading files from the high bandwidth FTP server, using Speedtest.net to measure the network throughput, switching back to the high bandwidth FTP server to download some more data, and then concluded by downloading a couple of Angry Bird games from the Google Play web site. Figure 51 shows the maximum throughput that we measured with each application/test procedure and Figure 52 shows a time series plot of the throughput.

Figure 51. Maximum Achievable Throughput – by Test Methodology and Application

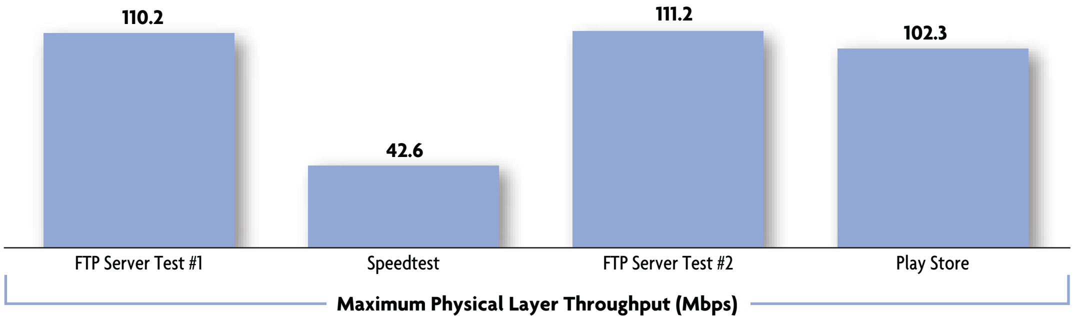

Source: Signals Research Group

Figure 51. Maximum Achievable Throughput – by Test Methodology and Application

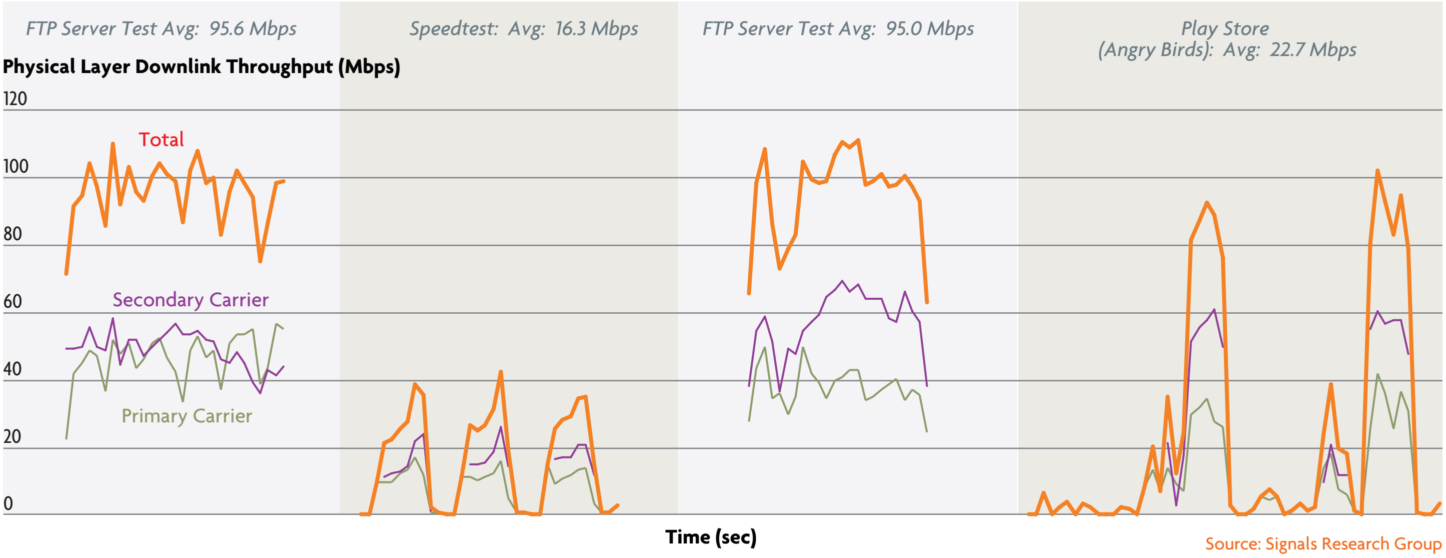

Source: Signals Research Group

We use third-party websites on occasion to measure network throughput, but we also know enough to take the answer with a bit of skepticism.

The results are interesting for a couple of reasons. First, it is apparent that the peak throughput that we measured via Speedtest.net was substantially lower than the actual capabilities of the network. Speedtest.net may not be a popular application in South Korea, but it is available and we know that we used a local server to do the tests. We observed a similar situation involving Speedtest.net when we tested the Rogers Wireless network in Vancouver. During the Vancouver tests, we tried several different Speedtest.net recommended servers that were in the vicinity of Vancouver. A couple of the servers provided what we felt were reasonably accurate results while a couple of the servers provided results that we felt were well below the capabilities of the network at the time and place we were doing the tests. We aren't sure where the fault lies and we use third-party websites on occasion to measure network throughput, but we also know enough to take the answer with a bit of skepticism. As wireless networks become faster the issue could become more prevalent.

The other interesting observation is that the download from Google Play took full advantage of everything the network had to offer. In the past we have seen a few instances when we felt the host server (Google, iTunes, etc.) was the bottleneck but in this particular instance this wasn't the case. The average data rate during the Google Play sequence seems low but that is because we based the average over the entire time period instead of trying to filter the values based on when the download was actually taking place.

5.3 Voice and Video Telephony Applications

On our last day in Seoul we did some tests with VoLTE, video telephony, Skype Voice and Skype Video. We didn't have the opportunity to evaluate voice quality but we were able to analyze the data based on the required amount of network resources. Table 1 provides values for several KPIs and Figure 53 shows the allocation of downlink and uplink resource blocks when the calls were active.

It is our understanding that a call placed between an LTE smartphone and a 3G smartphone will use NB-AMR (12.2 Kbps) and that a VoLTE call between two compatible devices will use WB-AMR. In the case of South Korean operators, they use a fixed codec of 23.85 Kbps since they care more about the quality of the voice call than they do about maximizing the use of available network resources. Of course, the networks also seem to have ample resources available due to the cell density.

We consider the information presented in this section to be preliminary since we would like to do more extensive testing and include some sort of voice/video quality analysis. We also greyed the

VoLTE – NB result for the average amount of downlink bytes used per second. To us the number doesn't make sense compared with the VoLTE – WB results. There is a chance that what we thought was a call using NB-AMR actually wasn't so we need to revisit this matter in the future.

What we can conclude at this time is that Skype Voice seems to require more network resources than VoLTE. This observation isn't surprising since the Skype codec uses between 6 Kbps and 40 kbps – most likely toward the higher end in the LTE network that we tested. The operator's video telephony service was the big bandwidth hog compared with the plain vanilla voice service, but to be fair it was very compelling to use and the networks seemed to have plenty of available capacity at the moment. Skype Video still takes the prize for consuming the most network resources.

Table 1. Key Network Utilization Parameters – by Application

Application	Downlink RB Allocation (Avg)	Uplink RB Allocation (Avg)	Rank Indicator 1 (requested)	Rank Indicator 2 (requested)	Rank Indicator 1 (actual)	Rank Indicator 2 (actual)	Downlink Bytes (Avg per sec)
VoLTE - NB	0.28	0.16	14.4%	85.6%	19.0%	81.0%	5,905.4
VoLTE - WB *	0.45	0.24	30.2%	69.8%	48.8%	51.2%	4,611.0
Skype Voice	0.71	0.22	21.9%	78.1%	55.7%	44.3%	7,374.1
Video Telephony *	2.36	1.73	26.6%	73.4%	50.6%	49.4%	73,365.4
Skype Video	2.50	1.59	27.0%	73.0%	58.5%	41.5%	96,283.0

**Carrier Aggregation used very sporadically*

Source: Signals Research Group

We expected to find the VoLTE /video telephony calls assigned to a particular frequency band, but this wasn't the case.

We expected to find the VoLTE /video telephony calls assigned to a particular frequency band, but this wasn't the case. In the "VoLTE-NB" and "video telephony" tests the two smartphones used different frequency bands during the call, despite the phones being adjacent to each other in the moving vehicle. In the "VoLTE – WB" test both smartphones used 2115 MHz. During the "Skype Voice" test one of the phones actually switched frequency bands during the call – it started at 2115 MHz and then switched to 889 MHz during a cell handover when the SINR dropped to ~5 dB. The other smartphone remained at 889 MHz throughout the entire test.

Figure 53. Key Network Application Parameters – by Application

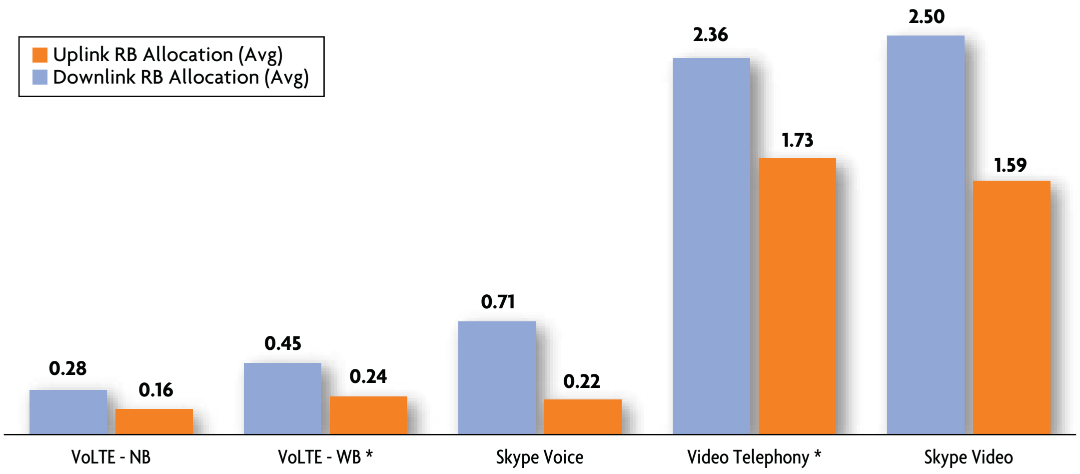

** Carrier Aggregation used very sporadically*

Source: Signals Research Group

Before leaving Seoul we wanted to view a 1080p video from YouTube. We selected a Planet Earth video and nearly wrecked the car, thanks to the incredible quality of the video. Our Samsung notebook computer probably helped in that regard. The video playback was flawless and the start time seemed nearly instantaneous. The last three figures in this section showcase how carrier aggregation was used to deliver the video content.

Figure 54 shows the Physical Layer throughput as a function of time. We have identified where we believe the video was initially buffering before beginning to play. The time preceding the buffering was when we were navigating the YouTube site to select a good video. As the figure shows, carrier aggregation was used to download the video but it was used very sparingly, especially after the initial download that filled the buffer. We've included the last two figures since they show the primary and secondary carriers flipping frequency bands when the primary carrier's RSRP and SINR dropped to appreciably low values – roughly -105 dBm for RSRP and -1 dB for SINR.

Figure 54. You Tube 1080p Video Playback – Time Series

Primary Throughput (Avg) = 6.16 Mbps
Secondary Throughput (Avg) = 8.83 Mbps
Total Throughput (Avg) = 9.22 Mbps
PEAK TBS SIZE = 146,784 bits

Primary Throughput (Max) = 37.95 Mbps
Secondary Throughput (Max) = 42.40 Mbps
Total Throughput (Max) = 62.12 Mbps

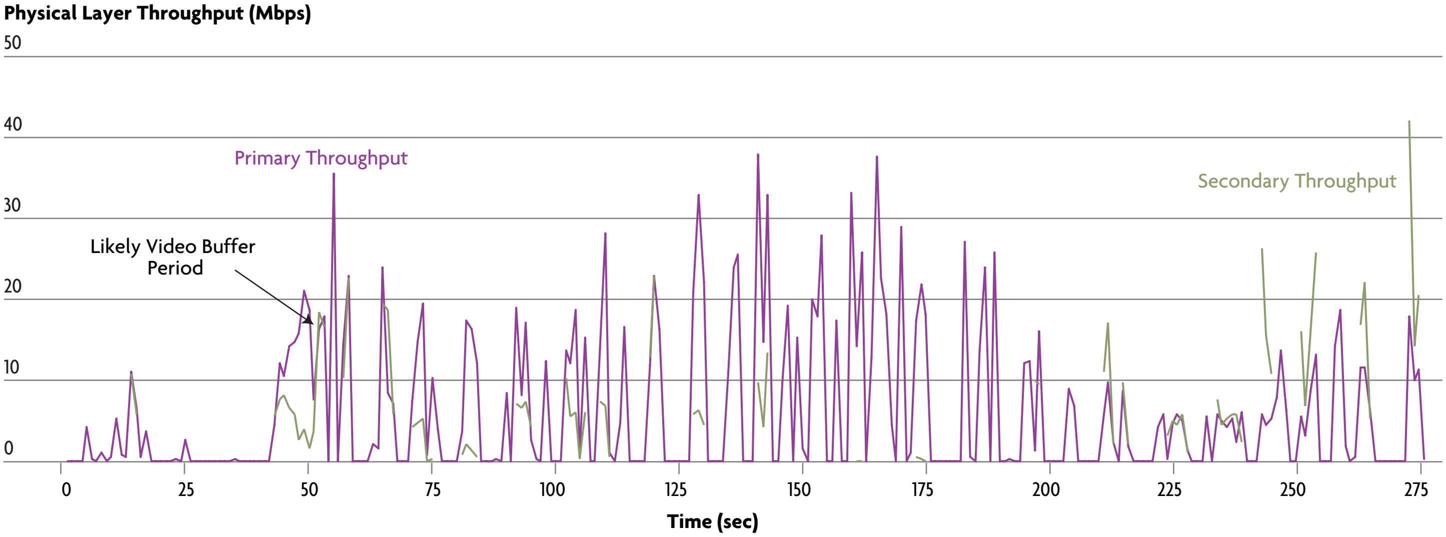

Source: Signals Research Group

Figure 55. SINR by Primary and Secondary Carrier (YouTube 1080p playback) – Time Series

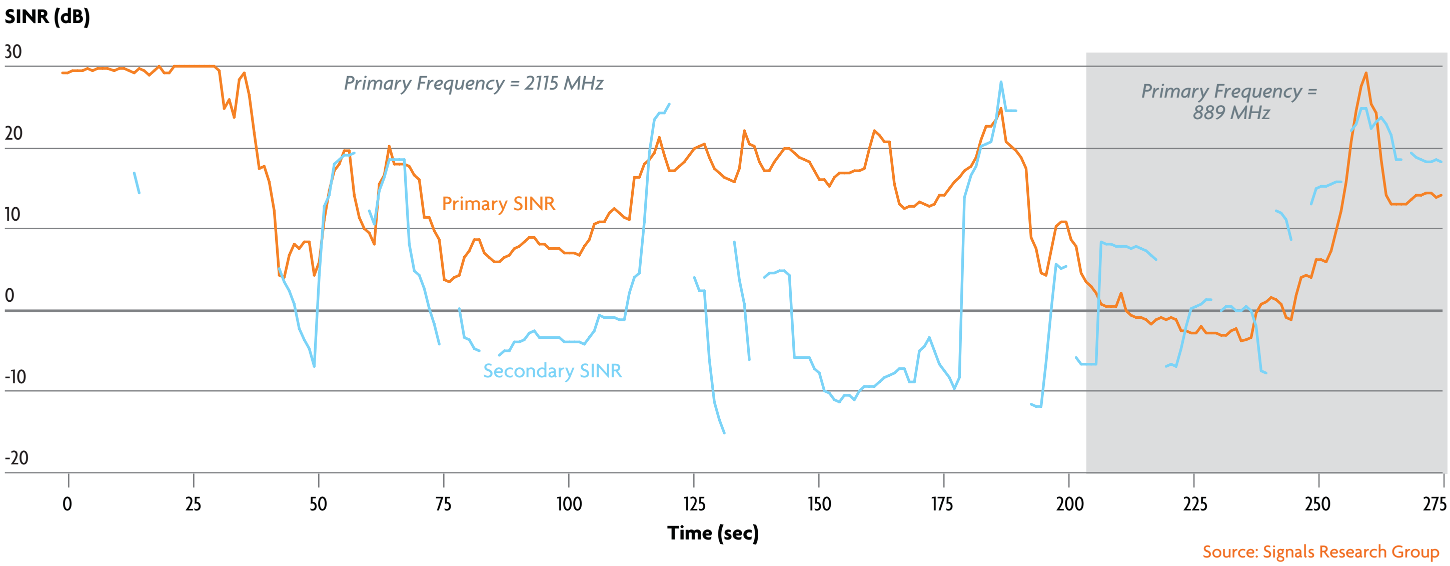

Source: Signals Research Group

Figure 56. RSRP by Primary and Secondary Carrier (YouTube 1080p playback) – Time Series

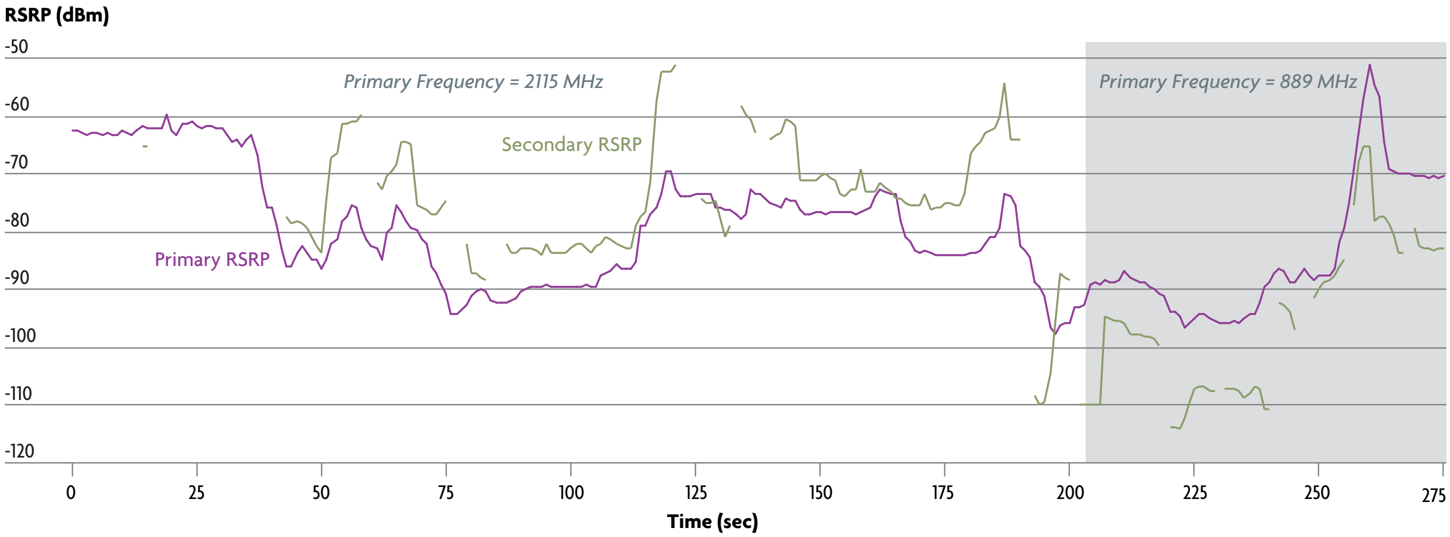

Source: Signals Research Group

6.0 Test Methodology

For the LTE Carrier Aggregation testing we once again used the Accuver XCAL drive test tool to collect the data and the Accuver XCAP post-processing tool to analyze the data and to help us create the figures that appear in this report. Figure 57 shows a screen shot of the XCAL tool in action. The top row shows the total Physical Layer throughput, or the sum of the primary and secondary carriers. The middle row shows the Physical Layer throughput from the primary carrier and the bottom row shows the Physical Layer throughput from the secondary carrier. The screen shot stems from the 1815 hours – note the total throughput shown in the figure is over 100 Mbps.

Figure 57. XCAL in Action

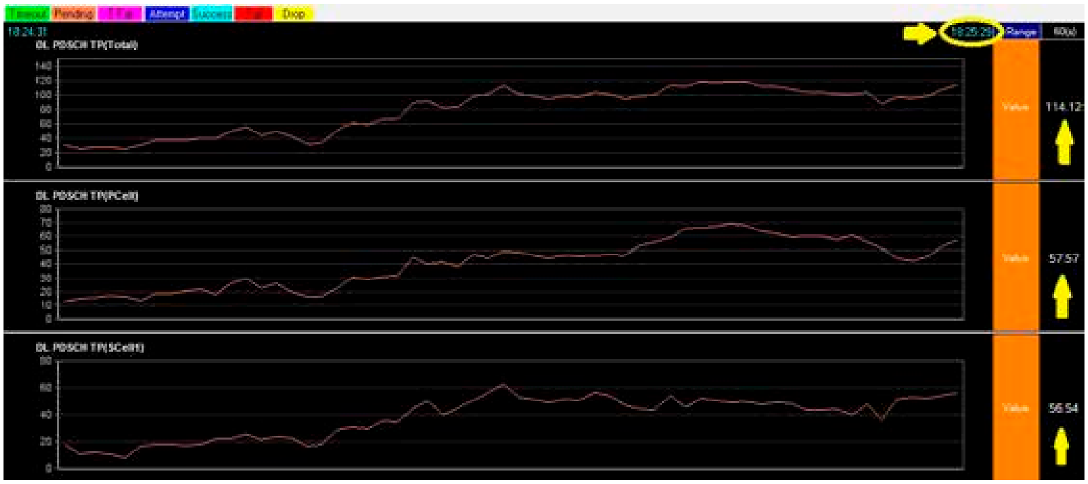

Source: Signals Research Group

For the web browsing tests we used Spirent's Datum application.

For the web browsing tests we used Spirent's Datum application to load the web pages and to automatically track and record the amount of time that it took to load each web page. We've used the Datum tool in the past – most recently the MIMO drive test study that we published in August – as well as in a major benchmark study that we conducted for an operator. It was the operator's preferred solution for the study.

As is the case with all *Signals Ahead* reports, this entire endeavor was self-funded. We rented two Release 8 mobile phones at the airport and used their SIMs in the Release 10 smartphones. We separately purchased a LG G2 Release 10 smartphone and a LG Optimus Release 8 smartphone. We were also able to leverage two Qualcomm test phones that we used occasionally during our testing.

We also could not have done this report without the support of Accuver.

We also could not have done this report without the support of Accuver who provided us with its suite of drive test tools and post-processing software. SRG takes full responsibility for the analysis and conclusions that are documented in this report.

All of our testing took place from a moving vehicle, unless we were stuck in traffic or at a traffic light.

All of our testing took place from a moving vehicle, unless we were stuck in traffic or at a traffic light. We tested at all hours of the day, from 0300 in the morning until 1900 in the evening. The drive routes that we selected were entirely random, but we elected to remain in the Gangnam area of Seoul. Therefore, the results that we present in this report may not reflect the network performance across the greater Seoul area. Testing indoors would have produced different results and we know that building penetration loss would have degraded the signal over what we observed. However, we also know that merely testing in a few in-building locations would result in statistically meaningless results. We believe it is better to obtain statistically meaningful results and then let readers apply their own adjustments to compensate for in-building performance.

For the throughput tests, the Release 10/Release 8 modems were tethered to notebook computers. For many of the user experience tests the data connection terminated at the smartphone although we still connected the smartphones to our notebook computer so that we could log the data with XCAL.

Some of the results that we present in this report stem from combining multiple log files into a single log file. This approach ensures that our analysis and conclusions are based on the representative performance of the network. We did, however, focus our analysis on isolated results in order to illustrate interesting findings that we observed in the data.

The following information identifies how frequently the KPIs were reported in the log files.

- Vehicular Speed – once per second/data collected and averaged over the entire interval
- Serving PCI – ~once every 40 ms/data collected and averaged over the entire interval
- SINR – ~once every 40 ms/data collected and averaged over the entire interval
- RSRP – ~once every 40 ms/data collected and averaged over the entire interval
- Rank Indicator 1/Rank Indicator 2 – once per second/data collected and averaged over the entire interval
- CQI – ~once every 10 ms/data collected and averaged over the entire interval
- Number of Assigned Resource Blocks – once per second/data collected and averaged over the entire interval
- MCS Code Word 0/Code Word 1/data collected and averaged over the entire interval
- Modulation Rate (QPSK, 16 QAM, 64 QAM) – ~once every 50 ms/data collected and averaged over the entire interval
- BLER – once per second/data collected and averaged over the entire interval
- PDSCH Throughput – once per second/data collected and averaged over the entire interval

For the scatterplots, we linked the two applicable KPIs together and then did the necessary averaging. For example, for the SINR versus Physical Layer throughput plots, we averaged all reported SINR values plus or minus one second from the reported throughput value in order to obtain the corresponding SINR values. In some cases we also sorted the throughput values into discrete buckets in order present color coded information. For example, we used this methodology to create the three dimensional plots that show uplink throughput versus resource block allocation versus transmit power.

For the Rank Indicator and Category 4 analysis, we post-processed the log files so that information from every 1 ms subframe was present in the log file.

For the Rank Indicator and Category 4 analysis, we post-processed the log files so that information from every 1 ms subframe was present in the log file. We could then determine the Rank Indicator value that the network assigned the mobile device and compare it to the value that the mobile device requested – this information is found elsewhere in the log file. Similarly, by having information for every single subframe we could see the exact TBS value that the mobile device used versus an average over a one second increment. Our approach was more precise since a one second average could mask those brief instances when the TBS value was greater than what a Category 3 device could support.

7.0 Final Thoughts

If you made it this far without skipping a page then thanks! Until next time, be on the lookout for the next *Signals Ahead*….

Michael Thelander

Michael Thelander is the CEO and Founder of Signals Research Group. In his current endeavor he leads a team of industry experts providing technical and operator economics analysis for clients on a global basis. Mr. Thelander is also responsible for the consultancy's *Signals Ahead* research product, including its widely acclaimed "Chips and Salsa" series of reports that focus on the wireless IC industry.

Previously, Mr. Thelander was an analyst with Deutsche Bank Equity Research. Prior to joining Deutsche Bank, Mr. Thelander was a consultant with KPMG (now known as BearingPoint) and a communications officer with the United States Army. Mr. Thelander has also published numerous articles for leading trade publications and engineering journals throughout his career.

He has been an invited speaker at industry conferences around the world and he is frequently quoted by major news sources and industry newsletters, including T*he Economist, The Wall Street Journal, Investors Business Daily, Reuters, Bloomberg News,* and *The China Daily*. Mr. Thelander earned a Masters of Science in Solid State Physics from North Carolina State University and a Masters of Business Administration from the University of Chicago, Graduate School of Business.

SIGNALS AHEAD SUBSCRIPTION

THE SIGNALS AHEAD NEWSLETTER is available on a subscription basis. We offer four distinct packages that have been tailored to address the needs of our corporate users. The **GROUP LICENSE** includes up to five users from the same company. The **GLOBAL LICENSE** is the most attractive package for companies that have several readers since it is offered to an unlimited number of employees from the same organization. Finally, the **PLATINUM** package includes the Global License, plus up to five hours of analyst time. Other packages are available.

CORPORATE RATES (12 issues of *Signals Ahead* plus 4-6 issues of *Signals Flash!*)
☐ Group License ($3,995) ☐ Global License ($7,995) ☐ Platinum ($9,495)

PAYMENT TERMS
☐ American Express ☐ Visa ☐ MasterCard Credit Card # ______________________ Exp Date __ / __ / __
☐ Check Check Number ____________________
☐ Purchase Order PO Number ____________________
Name: ____ Title: ______________________
Affiliation: ____ Phone: (__) ____________________
Mailing Address: ____

MAILING ADDRESS
Signals Research Group
10 Ormindale Court
Oakland, CA 94611

Our **FAX** number is **(510) 338-1284.**

Alternatively, you may contact us at **(510) 273-2439** or at information@signalsresearch.com and we will contact you for your billing information. We will not process your payment until after the trial subscription period is completed.

TERMS AND CONDITIONS: Any copying, redistributing, or republishing of this material, including unauthorized sharing of user accounts, is strictly prohibited without the written consent of SRG.

PLEASE NOTE DISCLAIMER: The views expressed in this newsletter reflect those of Signals Research Group and are based on our understanding of past and current events shaping the wireless industry. This report is provided for informational purposes only and on the condition that it will not form a basis for any investment decision. The information has been obtained from sources believed to be reliable, but Signals Research Group makes no representation as to the accuracy or completeness of such information. Opinions, estimates, projections or forecasts in this report constitute the current judgment of the author(s) as of the date of this report. Signals Research Group has no obligation to update, modify or amend this report or to otherwise notify a reader thereof in the event that any matter stated herein, or any opinion, projection, forecast or estimate set forth herein, changes or subsequently becomes inaccurate.

If you feel our opinions, analysis or interpretations of events are inaccurate, please fell free to contact Signals Research Group. We are always seeking a more accurate understanding of the topics that influence the wireless industry. Reference in the newsletter to a company that is publicly traded is not a recommendation to buy or sell the shares of such company. Signals Research Group and/or its affiliates/investors may hold securities positions in the companies discussed in this report and may frequently trade in such positions. Such investment activity may be inconsistent with the analysis provided in this report. Signals Research Group seeks to do business and may currently be doing business with companies discussed in this report. Readers should be aware that Signals Research Group might have a conflict of interest that could affect the objectivity of this report. Additional information and disclosures can be found at our website at www.signalsresearch.com. This report may not be reproduced, copied, distributed or published without the prior written authorization of Signals Research Group (copyright ©2013, all rights reserved by Signals Research Group).