NIST Internal Report
NIST IR 8566

A QoE-Based Method for Quantifying and Comparing LTE Coverage

Wesley D. Garey
Yishen Sun
Richard A. Rouil

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8566

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8566

A QoE-Based Method for Quantifying and Comparing LTE Coverage

Wesley D. Garey
Yishen Sun
Richard A. Rouil
Wireless Networks Division
Communication Technology Laboratory

This publication is available free of charge from:
https://doi.org/10.6028/NIST.IR.8566

March 2025

U.S. Department of Commerce
Howard Lutnick, Secretary

National Institute of Standards and Technology
Craig Burkhardt, Acting Under Secretary of Commerce for Standards and Technology and Acting NIST Director

Publication History
Approved by the NIST Editorial Review Board on 2025-03-10

How to cite this NIST Technical Series Publication:
Garey WD, Sun Y, Rouil RA (2025) A QoE-Based Method for Quantifying and Comparing LTE Coverage. (National Institute of Standards and Technology, Gaithersburg, MD), NIST IR 8566. https://doi.org/10.6028/NIST.IR.8566

Author ORCID iDs
Wesley D. Garey: 0000-0003-1572-8208
Yishen Sun: 0000-0002-6334-253X
Richard A. Rouil: 0000-0003-0387-0880

Contact Information
wesley.garey@nist.gov

Abstract

To determine the coverage of a Long Term Evolution (LTE) system, measurements such as the Reference Signal Received Power (RSRP), Reference Signal Received Quality (RSRQ), and Signal to Interference and Noise Ratio (SINR) are often used. However, RSRP and RSRQ are only relevant to LTE, and all three measurements are representative of signal quality but not user experience. In this paper, we propose a method that relates signal quality to Quality of Experience (QoE) using simulation. With this method, the coverage of an LTE system is quantified using speech intelligibility which is agnostic to an underlying technology making it possible to compare different technologies while also allowing users to directly relate speech intelligibility scores to sampled audio. In addition to the proposed method, this paper also describes an implementation of the method, a case study of the method using drive test data, and validation of the proposed method.

Keywords

LMR, LTE, Modeling and Simulation, Quality of Experience, Speech Intelligibility

Table of Contents

List of Tables

List of Figures

Acknowledgments

We would like to thank the Public Safety and Communications Research (PSCR) division at the National Institute of Standards and Technology (NIST) in Boulder, Colorado. Our collaboration lead to the conceptualization of this work, and their efforts collecting measurements are what made the implementation and validation of the proposed method possible. This includes Sam Ray, Jordan O'Dell, Justin Sadinski, Lisa Soucy, and Gary Howarth.

1. Introduction

In 2012, the First Responder Network Authority (FirstNet) was created to establish and operate a broadband mobile network for public safety as a part of a larger initiative to advance communications for first responders. In essence, the goal of this initiative is to migrate public safety communications from the legacy Land Mobile Radio (LMR) technology to the more advanced Long Term Evolution (LTE) technology. This would provide public safety with access to broadband services that include digital voice, video, and data, to supersede the traditional, analog voice technology. However, as can be gleaned from the work in [1], LMR and LTE are two vastly different technologies that both come with their own strengths and weaknesses. In particular, LMR is the status quo for many public safety agencies, as it is widely deployed in that community and is best known for its resiliency and reliability, especially over large geographical areas. LTE, on the other hand, is the status quo for commercial deployments and is best known for its ability to provide low latency communication, high-speed data rates, and a large traffic capacity. With that said, several questions still remain for many public safety agencies that have the option to join FirstNet. These questions include, whether or not LTE should completely replace LMR, both systems should coexist independently, or if the two should be deployed and used together? Given these questions, we are proposing a method that can be used to estimate the coverage of an LTE deployment by mapping Radio Frequency (RF) measurements to speech intelligibility. This method is intended to provide public safety and other organizations in academia and industry with an approach that can be used in combination with the collection of RF measurement data to analyze the performance of an existing LTE deployment in a specific geographical location of interest. As will be described later in this document, using said approach comes with the benefit of easily relating speech intelligibility scores to audio, as well as the potential to provide a path to compare the coverage of different deployments in the same geographical area even when they use different underlying technologies (i.e., something other than LTE).

The remainder of this paper is organized as follows. Section 2 will describe the proposed method. Section 3 will cover an implementation of the proposed method using simulation and other tools. Section 4 will discuss how we verified the proposed method using Software Defined Radios (SDRs). Finally, Section 5 will provide a summary of this paper and discuss future work.

2. Coverage Evaluation Method

The goal of the evaluation method is to provide a way for organizations (e.g., first responders) to evaluate the coverage of a mobile network deployment in a given geographical area using Quality of Experience (QoE). With the method that we are proposing, speech intelligibility is the QoE metric that is used to measure network coverage. The purpose of using such a method to assess the network coverage of a mobile network is twofold. First, when using speech intelligibility, first responders will be able to relate a speech intelligibility score with what the audio sounds like since the audio that is used to derive an intelligibility score can be sampled. Second, since speech intelligibility is agnostic to an underlying technology, it would be possible for first responders to perform an apples-to-apples comparison of different deployments whether they use the same technology or not. For instance, Fig. 1 is an illustration of an area with both an LMR and LTE deployment, where the red circle outlines a particular location at which coverage is to be evaluated for each technology and then compared.

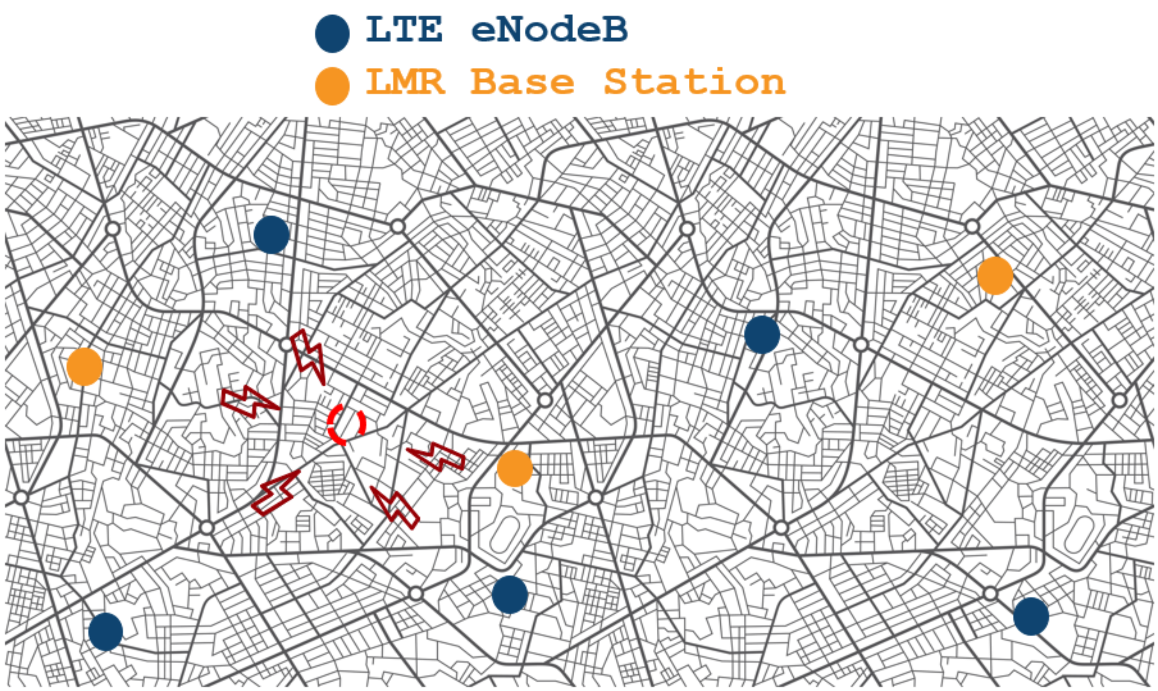

Fig. 1. Example of area with LMR and LTE deployments

Following the flow diagram, depicted in Fig. 2, the first step requires feeding input audio through both the LMR and LTE system. The output audio of both the LMR and LTE systems can then be analyzed to determine what the speech intelligibility at that particular location is. The resulting speech intelligibility (or QoE) score would then allow for a direct comparison of system performance between each deployment at a particular location. Furthermore, The coverage of the deployments can be compared over an entire area by performing this evaluation at each location of interest, and then plotting the speech intelligibility for each coordinate creating a coverage map for each technology. Once each coverage map is created, they can be overlaid and studied to understand how the coverage compares between the two deployments in a particular area.

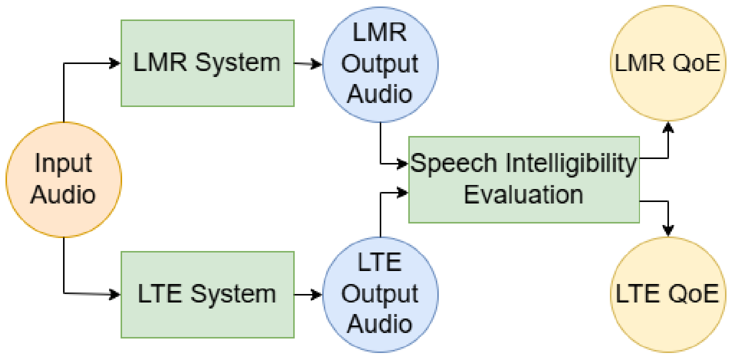

Fig. 2. Flow diagram showing the use of input audio to determine QoE metrics of each type of system

Now, let us discuss the method that we are proposing which is depicted in Fig. 3. This method consists of collecting measurements throughout a given area for the network deployment. For instance, RF drive testing across the area can be performed to collect the needed measurements, and then, using simulation tools and the collected data, create a simulation environment in which the simulator models the technology and accepts a subset of information from the measurements to realize the "channel conditions." Note that such an evaluation could be performed by using the deployed system directly, however, that would require authorization to use and/or control the network equipment which may or may not be possible. Thus, simulation can be used to circumvent this so that only measurement data is needed from the real deployment which can be collected at a location without using the network to transport the audio. Once the simulated system is up and running that reflects the channel condition, it is time to simulate sending the audio. This is accomplished by generating packetized voice using the encoder of the vocoder of choice, and then sending packets with similar characteristics in the simulator to represent the audio stream. The packet trace from the simulation can then be used to observe and control which packets are lost and the delay of those packets that are received. From there, the audio is reconstructed using the decoder of the selected vocoder, and then this output audio is stored in a file so that it can be evaluated later. At this point, the output audio and the impact that the system would have on that audio has been collected under the given channel conditions that are taken from the measurement data. From here, a speech intelligibility tool can be used to evaluate the resulting audio and provide a speech intelligibility score. Once the previously mentioned steps have been repeated for each location, the final step is to create a coverage map that associates each speech intelligibility score with its respective location. From here and throughout the remainder of this paper, we will focus on an LTE realization of this method since the implementation that we demonstrate in Section 3 is specific to LTE.

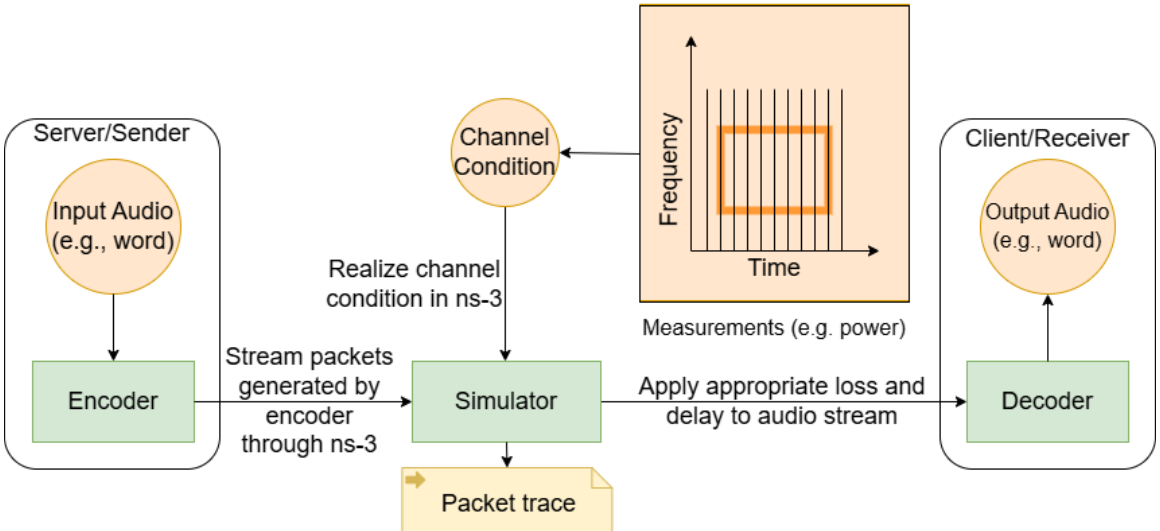

Fig. 3. Workflow of LTE system side of the method

3. An Implementation of the LTE Evaluation Method

The LTE implementation of the proposed method makes use of several different tools to both simulate the LTE network and characterize speech intelligibility, as shown in Fig. 4. These tools are what make it possible to relate measurement data with the corresponding QoE metrics. This includes the use of ns-3 [2] and the Articulation Band Correlation Modified Rhyme Test 16 (ABC-MRT16) algorithm [3]. ns-3 is used to simulate the LTE network for a particular configuration and Carrier-to-Interference plus Noise Ratio (CINR) value. Through the use of Test Access Point (TAP) devices and network bridges, two Linux Containers (LXCs) are deployed to act as a server on the internet and a User Equipment (UE) in the Radio Access Network (RAN) via ns-3 as it simulates the LTE network. The Fast Forward Moving Picture Experts Group (FFmpeg) [4] multimedia framework is deployed on each LXC and is the application used to encode, transmit, and decode the Real-time Transport Protocol (RTP) audio that is transmitted from the server to the UE. The audio that is received at the UE is saved for later so that the ABC-MRT16 algorithm can process the set of audio files associated with an LTE configuration and CINR value. This process is what allows us to quantify the speech intelligibility for a given LTE configuration and CINR pair, which is later used to classify measurements that are collected in the field.

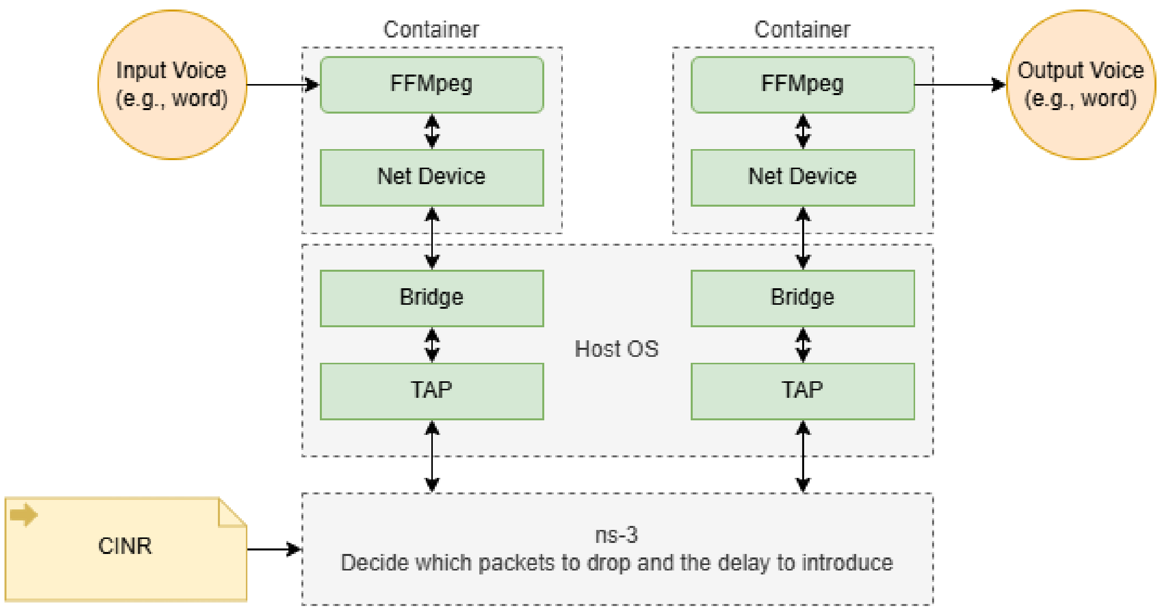

Fig. 4. Diagram of LTE simulation components

3.1. Audio Codec

In Fig. 5, an illustration depicts how audio is transported through an LTE system. Note that the traffic demand of an audio stream can typically be represented by a Constant Bitrate (CBR), however, each codec can have a different CBR which will likely impact the maximum quality of audio that can be achieved and the demand that this would put on the underlying communication technology. When it comes to mission critical communication, the 3rd Generation Partnership Project (3GPP) has specified that the Adaptive Multi-Rate Wideband (AMR-WB) codec must be supported [5]. Therefore, it is the codec that we use in our implementation, specifically the VisualOn (encoder) and Open Core (decoder) implementations that are supported by FFmpeg [4]. AMR-WB supports several different bit rates, with the lowest being 6.60 kbits/s and the highest being 23.85 kbits/s. In our implementation we use the highest bit rate since it is expected that it will be the most demanding on the system, and from the user's perspective, can provide the highest quality audio. Furthermore, since we assume that the AMR-WB codec will be used at the 23.85 kbit/s bit rate, this means that when audio is packetized, 60 byte packets will be generated every 20 ms.

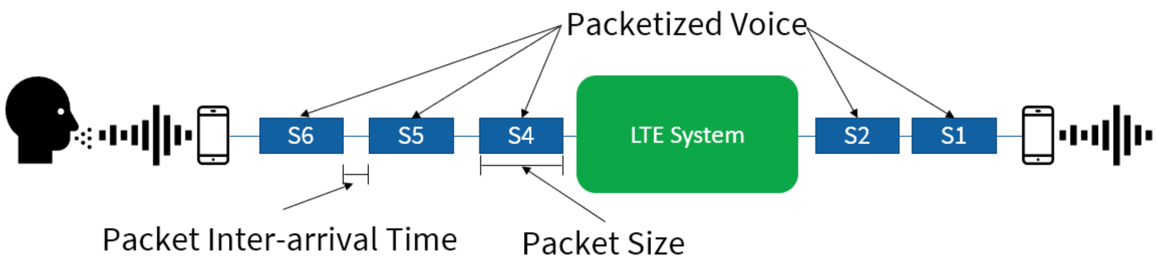

Fig. 5. Diagram of packetized voice.

3.2. Simulating LTE

ns-3 provides a comprehensive simulation model of LTE [2]. As is depicted in Fig. 6, this includes simulation models for UEs, Evolved Node Bs (eNBs), the air interface between UEs and eNBs, an LTE core that is comprised of a Service Gateway (SGW) and Packet Data Network Gateway (PGW), and servers on the internet. Since ns-3 is capable of simulating all of the previously mentioned components to some degree, this means that the simulated system, which is a combination of the channel, network layers, control traffic, and data traffic, should behave similar to a true LTE system. Hence, ns-3 is able to model the delay, loss, and overhead that an audio packet would incur as it is transported from the server to a UE as it traverses the LTE core and RAN. It is the over-the-air component where we inject our channel condition to reflect the effect that a measurement would have on speech intelligibility. In ns-3, there are simple and complex propagation models that are used to simulate the channel based on the configuration assumptions, distance between the UE and eNB, as well as any other simulated obstacles in the environment. However, since we want to control the channel condition, we modified the simulation so that no propagation related characteristics are considered, and we simply set the Signal-to-Interference plus Noise Ratio (SINR) value of all transmissions that take place in the Downlink (DL) between a UE and eNB in a simulation to a value that corresponds to a measured CINR. Therefore, each simulation that we run for a particular SINR value we can relate to a CINR from the collected measurement data. In our case, the CINR from our measurement data is the SINR of the reference signal that is received from the eNB at a given location.

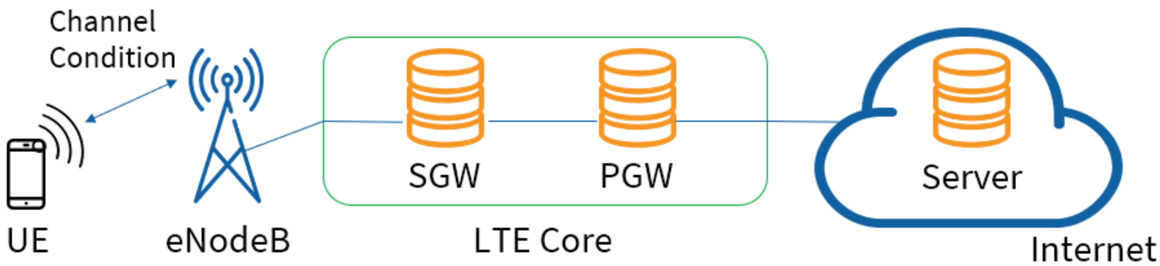

Fig. 6. Diagram of simulated LTE architecture.

ns-3 also makes it possible to configure many different aspects of the LTE network, all the way from the Non Access Stratum Layer (NAS) to the Physical Layer (PHY). However, the most important aspects that this method requires for consideration are the frequency and bandwidth. The reason why those two characteristics are the most important is because those two, in combination with the SINR, are what drive the Block Error Rate (BLER). The BLER, which is a part of the LTE error model in ns-3, is ultimately what determines the quality of the resulting audio since it directly impacts which audio packets will be lost or received. It is also important to note that while the SINR, frequency, and bandwidth are the inputs to our simulation, there are also several internal components that impact the BLER. This includes the Modulation and Coding Scheme (MCS) and number of Physical Resource Blocks (PRBs) that are selected by an eNB to perform a DL transmission. The SINR affects the highest MCS that can be used and the demand of the audio stream will impact how many PRBs the scheduler allocates to fulfill each transmission. Thus, the sample size and rate of a codec determines how large of a Transport Block (TB) is required for each transmission. Therefore, once the highest MCS is determined, depending on how much data needs to be sent, it can be determined how much of the available bandwidth (i.e., number of PRBs) is required to meet the demand. A TB is the lowest level at which data can be transmitted in ns-3, and the MCS and PRB associated with a TB determine which BLER curve is used. It is then the SINR value that determines the chance of successfully decoding or dropping a TB depending on where that SINR value is located on the curve.

Content	Type	Size (bytes)
20 ms Audio Sample	Data	60
Adaptive Multi-Rate Wideband (AMR-WB)	Header	2
Real Time Transport Protocol (RTP)	Header	12
User Datagram Protocol (UDP)	Header	8
Internet Protocol (IP)	Header	20
Packet Data Convergence Protocol (PDCP)	Header	2
Radio Link Control (RLC)	Header	2

Table 1. Packet contents.

It is also worth mentioning that while the size of a TB will be correlated with the size of the packets generated by the codec, the TB will likely need to be larger than the sample size.

This is due to the fact that the audio sample that is encoded by the codec will eventually be the payload of many different layers in the network. For instance, as can be seen in Table 1, the 60 byte sample generated by the AMR-WB codec with a 23.85 kbit/s bit rate results in a 106 byte datagram at the Medium Access Layer (MAC) layer due to the additional headers added by all of the other network layers. Hence, each TB will be sized to accommodate a 106 byte payload.

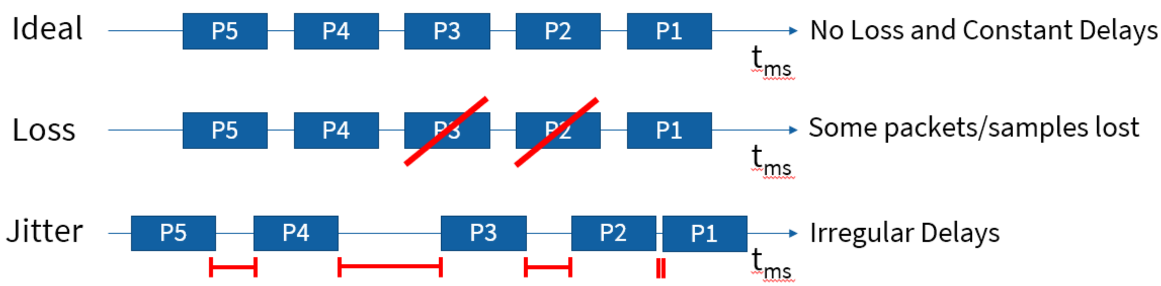

Fig. 7. Example of an ideal packet flow, a packet flow with loss, and a packet flow with jitter.

It is well known that irregular packet delays, known as "jitter," and packet loss play a significant role in the quality of audio that is transported over a network. Examples of jitter and loss are depicted in Fig. 7. Therefore, it is worth mentioning that in our simulations packet delay is almost constant and jitter is negligible since there is no traffic in addition to the audio that is sent from the server to a UE that could cause congestion in the network. In fact, the largest delay an audio packet will incur is when it is transported through the core from the server since our simulation includes point-to-point links with a constant 10 ms delay from the server to the SGW, SGW to PGW, and PGW to eNB. This results in a constant 30 ms delay that is observed for all packets up to that point. Any additional delay will be variable due to the eNB having to schedule resources to transport a TB over the air. However, those delays are typically under 1 ms. Hence, in our simulation it is packet loss, which is based on the input SINR value, that is affecting speech intelligibility not jitter.

3.3. Analyzing Speech Intelligibility

Now that we have covered how to simulate the channel condition and the effects that the simulation will have on audio, let us talk about how speech intelligibility is measured. To characterize the performance of the LTE system, we use FFmpeg to stream several audio files between two LXC instances connected by ns-3 for a particular evaluation of SINR value, frequency, and bandwidth. The source of the audio stream is the LXC representing the server while the receiving LXC that decodes the stream and stores the resulting audio represents the UE. The ABC-MRT16 algorithm is then used to process all of the saved audio files and determine an intelligibility score. The 1200 audio files that we stream for each evaluation comes prepackaged with the ABC-MRT16 algorithm. These audio files are made up of 4 groups of 50 blocks, with each block consisting of 6 phrases. The 4 groups correspond to 4 different speakers (i.e., voices), the 50 blocks correspond to 50 sets of 6 phrases, where each phrase in a group starts with, "Please select the word" followed by the keyword [3]. Each keyword in a group must rhyme with all other keywords in that group, and only differ by a leading or trailing consonant. For example, the set "bear," "chair," "dare," "tear," "mare," and "pair." Each audio file is approximately 2 seconds long resulting in a little over 36 minutes of audio. Before the audio is streamed, all 1200 audio files are converted down from a mono, Pulse-Code Modulation (PCM), format that uses 16 bits per sample and 48 000 samples each second, to a mono, AMR-WB format that uses 14 bits per sample and 16 000 samples each second before they are sent through the simulated LTE system. However, once received, the audio files must be converted back to the original PCM format so that they can be processed by the algorithm.

When using the ABC-MRT16 algorithm the following takes place. First, the algorithm uses the file name of the streamed file to determine which speaker and block the keyword belongs to. Using this information the algorithm then uses articulation band indices to compare the received file with the 6 other files in that particular block. The algorithm then selects a file from that block to "guess" which keyword is contained in the received audio. This selection is performed by comparing the articulation band indices of the audio from the received file with the audio from the other files in the block, and then choosing the file from the block that has the highest correlation. Since the algorithm knows which keyword is contained in the audio based on the file name, it can determine whether the "guess" is correct or not. After a guess is made for all of the received files, the ratio of correctly selected words becomes the ABC-MRT16 score for that iteration. Thus, a score of 0.9 would mean that out of the 1200 audio files, 1080 were "heard" correctly.

It should be noted that in our analysis, it is possible that the SINR is so low, that either no audio is received for a streamed file, or there are too few samples for the algorithm to make a guess. We refer to the this as a "complete loss of audio." The original ABC-MRT16 algorithm did not account for this, and so we added an additional condition such that if there is a complete loss of audio, then a "complete guess" is made. Meaning that when there is no audio to compare for a received file, the algorithm uses a uniformly distributed random variable to select one of the 6 words from a block when guessing.

3.4. Boulder Case Study

Fig. 8. Image of drive test van.

Now let us take a look at a case study using a deployment from Boulder, Colorado. For this case study we analyzed the coverage for a Band 14, LTE deployment. Band 14 is the band that was licensed to FirstNet by the United States (US) Congress in North America specifically for first responders. Since FirstNet has partnered with American Telephone and Telegraph (AT&T), they have been able to roll out and deploy Band 14 in several areas, including Boulder.

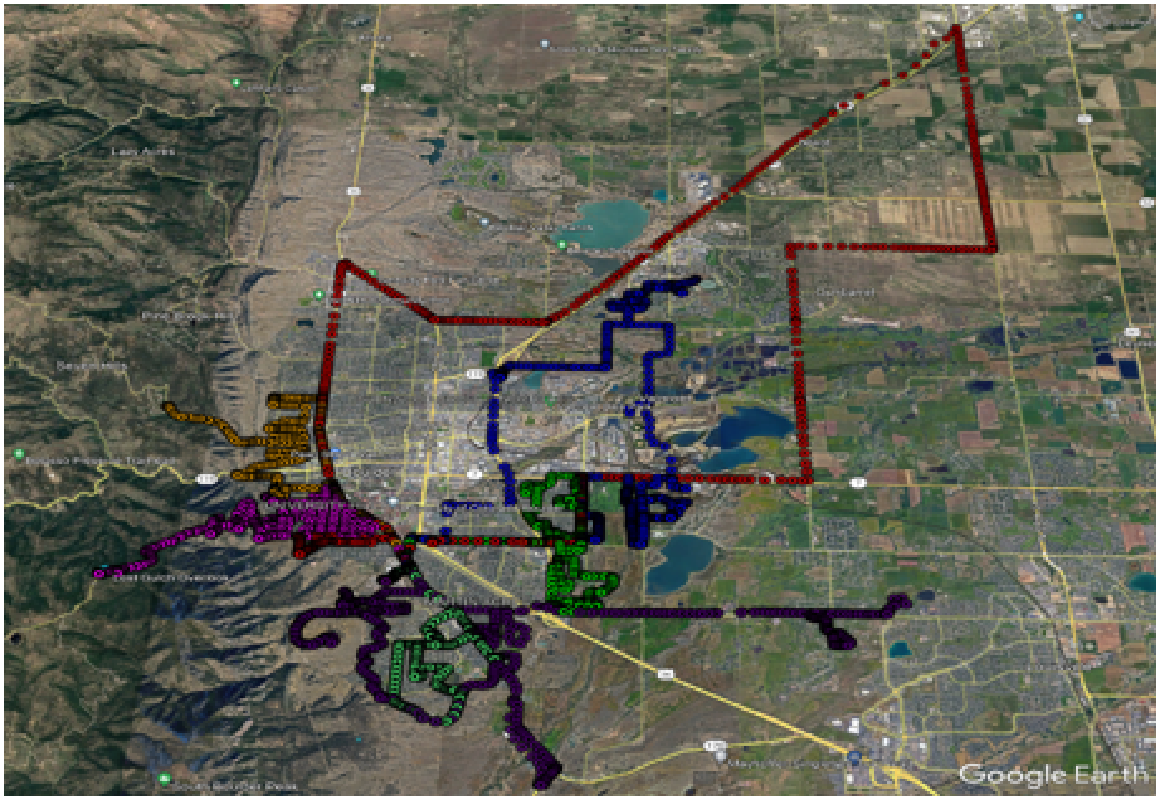

Fig. 9. Image of drive test van routes where measurements were collected.

The first step in the process was to collect measurements in the Boulder, Colorado area. Our colleagues used a specialized drive test van to collect LTE measurements for Band 14. This van, shown in Fig. 8, is equipped with a drive test scanner and Global Positioning System (GPS). With this equipment our colleagues drove several routes, shown in Fig. 9, and collected measurements at around half-second intervals. Each drive route resulted in a log file with several thousand rows of data that includes the date, time, latitude, longitude, cell identifier, and many reference signal attributes among several other metrics. The reference signal attributes that were collected include the time offset, delay spread, received power, received quality, and CINR. Therefore, with the measurement data that was collected we are able to associate a specific location (i.e., latitude and longitude) with a CINR value for the DL reference signal that was detected by the van.

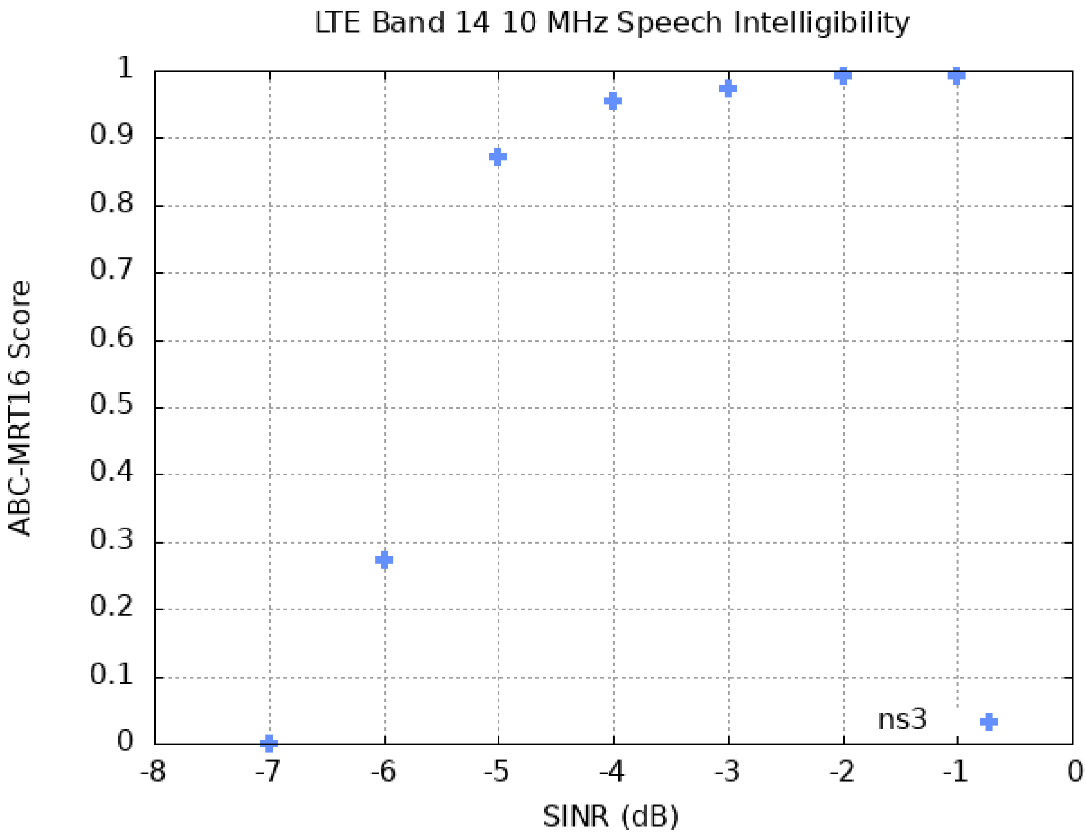

Fig. 10. Graph of speech intelligibility obtained from the ns-3 simulation for Band 14 using a 10 MHz bandwidth.

While the team in Boulder was collecting measurements, we started running simulations and evaluating the speech intelligibility to associate SINR values with an ABC-MRT16 score for Band 14. This means that we configured the simulation, taking into account Band 14 characteristics, to use a 763 MHz center frequency and 10 MHz bandwidth (i.e., 50 PRB). From there, all 1200 audio files were streamed through the simulated LTE system 8 different times to evaluate the speech intelligibility for integer SINR values ranging from -7 dB to -1 dB. This resulted in the graph shown in Fig. 10. In this graph, the ABC-MRT16 score is on the y-axis, while the SINR, in dB, is on the x-axis. As can be seen from the graph, with an SINR of -2 dB or higher, the speech intelligibility is almost perfect as the ABC-MRT16 scores are close to 1. However, for values at -7 dB and below the speech intelligibility is 0 indicating that the audio is unintelligible. For values greater than -7 dB and less than -2 dB, however there is some variation. It is with this information that we are able to map the collected measurement data with QoE.

$$SCORE(CINR) = \begin{cases} 0.00, & \text{for } CINR < -6 \\ 0.27, & \text{for } -6 \leq CINR < -5 \\ 0.87, & \text{for } -5 \leq CINR < -4 \\ 0.95, & \text{for } -4 \leq CINR < -3 \\ 0.97, & \text{for } -3 \leq CINR < -2 \\ 0.99, & \text{for } -2 \leq CINR < 0 \\ 1.00, & \text{for } 0 < CINR \end{cases} \tag{1}$$

To create the coverage map, we performed the following steps. First we parsed the measurement log file and created a Comma Separated Value (CSV) file with the latitude, longitude, reference signal CINR, and speech intelligibility score for each location. To determine the speech intelligibility score, the reference signal CINR is binned based on the function in Eq. 1 which is derived from the results in Fig. 10. From there, a Keyhole Markup Language (KML) file is created that plots place markers with a gradient from red, denoting a low ABC-MRT16 score, to green, denoting a high ABC-MRT16 score, a long with the score and measured reference signal CINR at each location, as shown in Fig. 11.

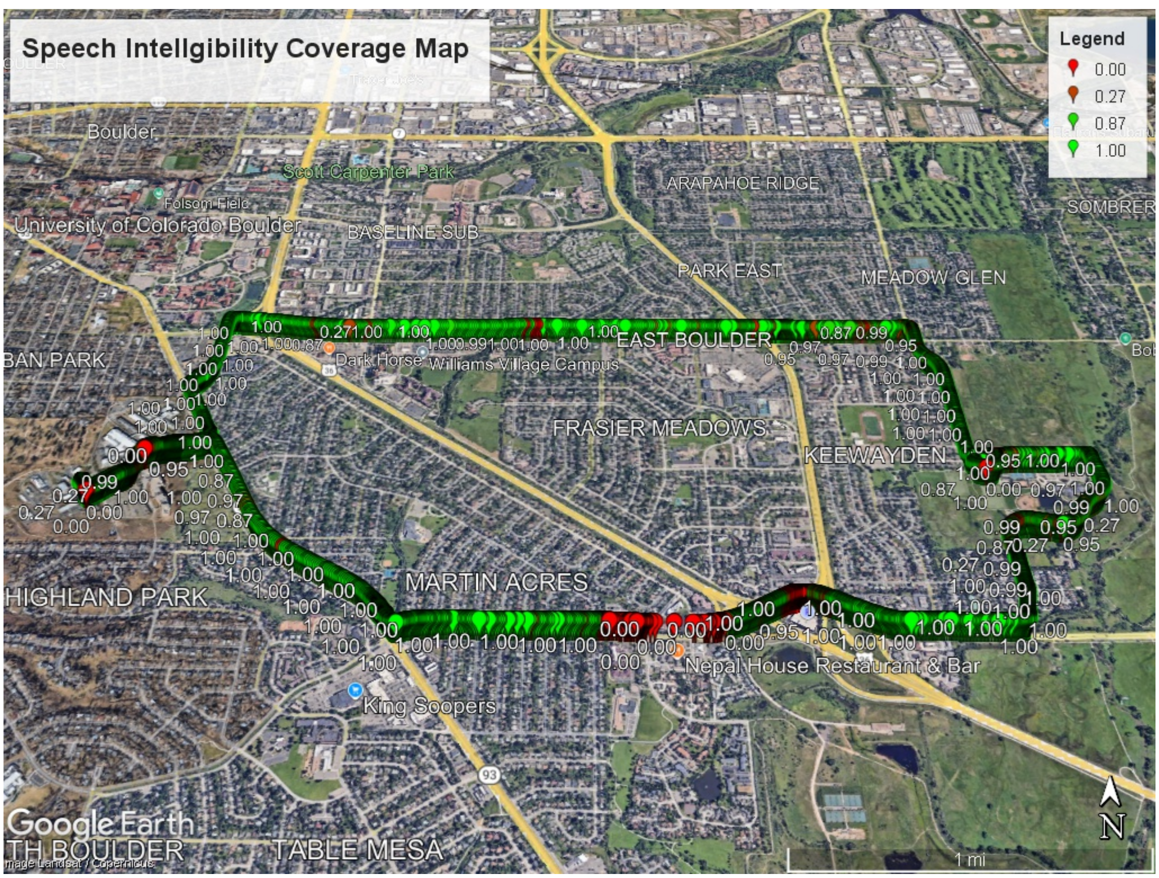

Fig. 11. A map of the Boulder, Colorado region that relates measurement data collected for Band 14 with a 10 MHz bandwidth to the speech intelligibility scores derived the from ns-3 simulations.

4. Validation of LTE Evaluation Method

In order to validate our proposed method for the LTE implementation, a testbed with an Open Air Interface (OAI) LTE RAN software stack [6] was used. As shown in Fig. 12, the testbed consists of two servers: a source and a receiver. Each server is deployed with an instance of FFmpeg which is used for encoding and decoding AMR-WB audio, as well as a respective instance of the OAI eNB and OAI UE. The Universal Software Radio Peripherals (USRPs) are connected to each server and are controlled by their respective OAI instances to perform the RF signaling. Both USRPs connect to a channel emulator which emulates the air interface and allows us to control the quality of each channel independently. This allows us to specify what the SINR is for the DL channel between the eNB and UE. With this setup, we repeat the same process that is described in Section 3.4 but in this case we are using an emulated LTE deployment as opposed to a simulated one.

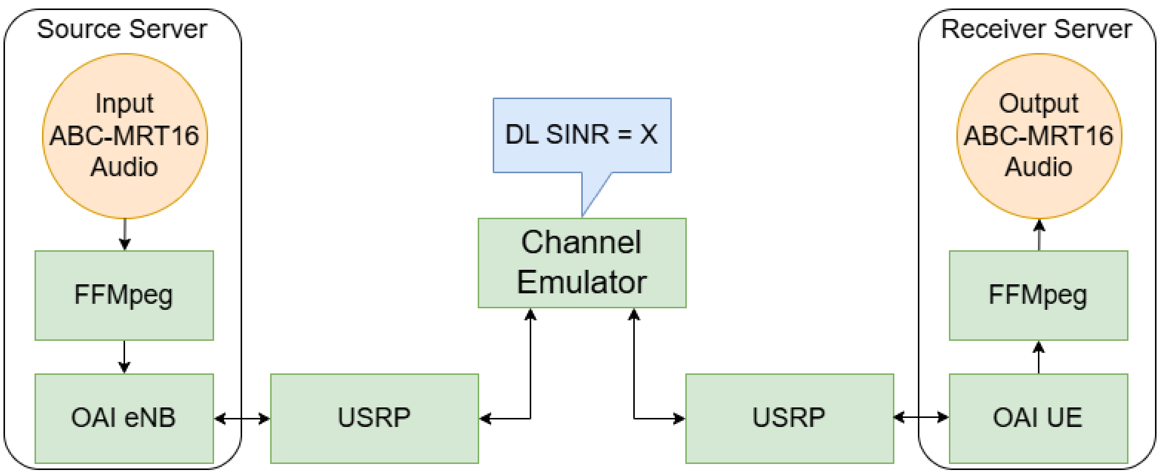

Fig. 12. Illustration of LTE testbed that depicts components and connections.

Unfortunately, when attempting to validate Band 14 which has 10 MHz of bandwidth available, it was discovered that with our OAI deployment this was not supported. To circumvent this, we opted to use Band 14 with a 5 MHz (i.e., 25 PRBs) which is supported with our OAI deployment. Because of this, we also had to generate simulation data as described in Section 3.4, however we simply used a 5 MHz bandwidth instead of 10 MHz bandwidth. This resulted in the graph shown in Fig. 13. On the x-axis is the SINR that is set in ns-3 and on the channel emulator for the DL channel. On the y-axis is the ABC-MRT16 score that is achieved after processing all 1200 audio files. As can be seen in the graph, the results above -4 dB and below -6 dB match very well, however there are visible differences on the graph for all other values between -6 dB and -4 dB.

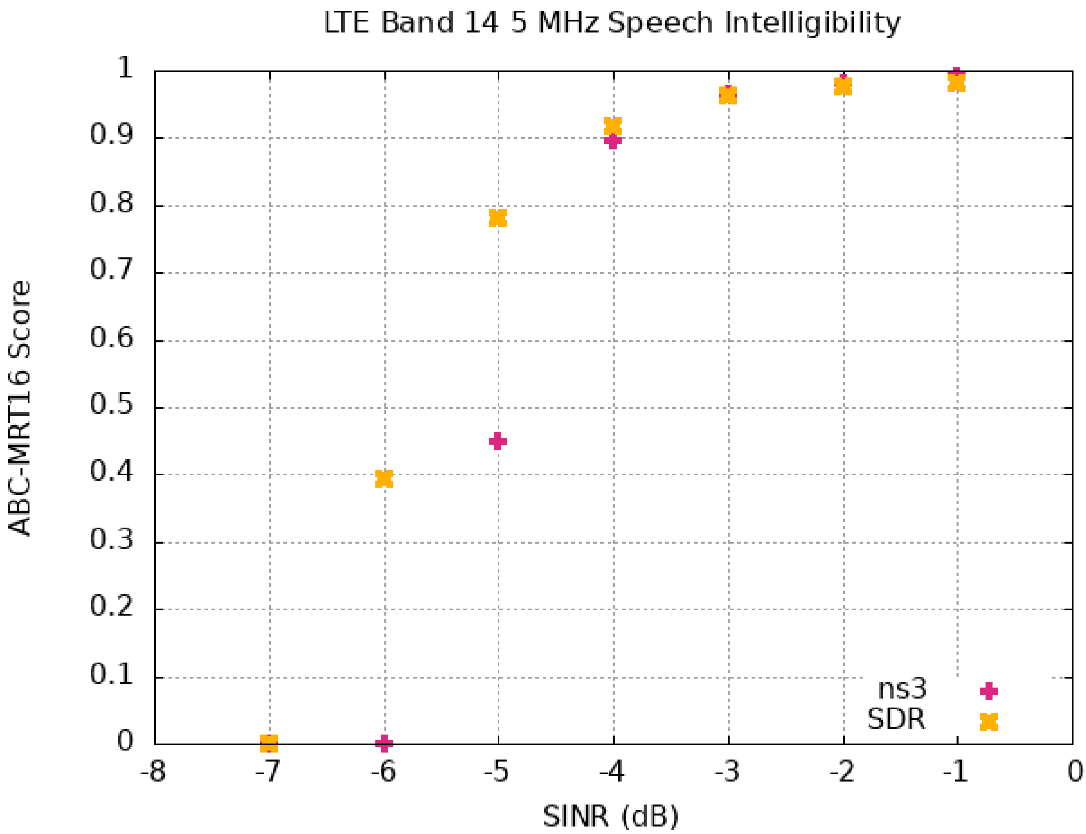

Fig. 13. A graph that plots the speech intelligibility achieved for the simulated LTE system (ns-3) and the emulated LTE system (SDR).

Due to this difference we also studied the variability of the CINR that is measured at a single location for just over an hour time period. These results are shown in Fig. 14, and are measurements collected for Band 14 with a 10 MHz bandwidth. In this graph, the measured CINR is on the x-axis while the Probability Density Function (PDF) of measurements is on the y-axis. As can be seen in the graph, of the 7900 samples that were collected, 80 % of the measurements fall between 3 dB and 5 dB (25 % at 3 dB, 29 % at 4 dB, and 26 % at 5 dB). From this graph we can see that the 1 dB shift that exists in Fig. 13 between -6 dB and -5 dB is negligible due to the variability that can be observed at a single location is greater than 1 dB. Therefore, the method that is proposed appears to accurately represent the speech intelligibility of an LTE system. This graph also indicates that when performing drive tests, it is important to take multiple measurements to ensure that an accurate representation of the location is captured.

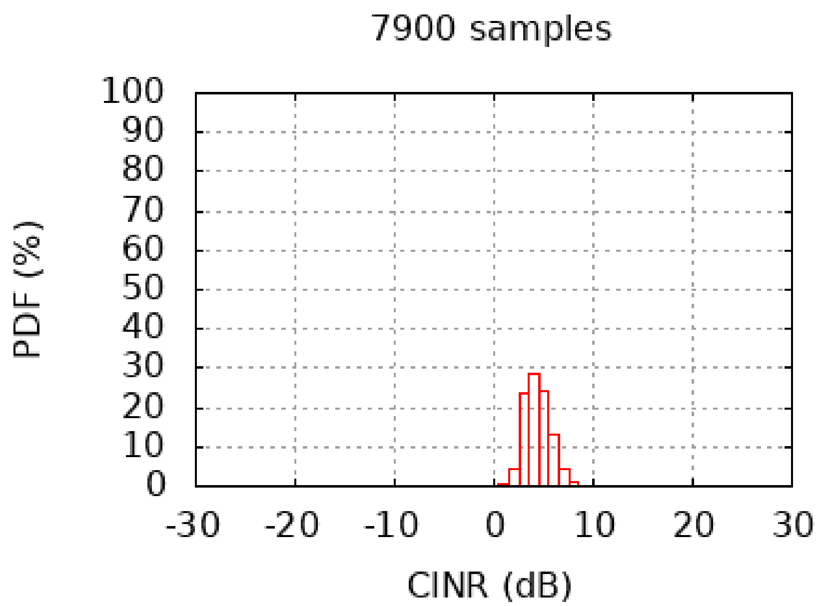

Fig. 14. A PDF that shows the distribution of 7900 reference signal CINR measurements at a single location for over an hour of time.

5. Conclusion

In this paper, a method is proposed to measure the performance of a deployed mobile network. This included a description of that method, how it could be used to compare the performance of two different networks, an LTE implementation of that method, and how that implementation was validated. For the future, it is desirable to come up with an LMR implementation of this method so that coverage maps can be overlaid to perform a direct comparison between LMR and LTE. We believe that this would be very valuable to public safety and that it would also make this effort whole. We would also like to continue relating measurement data with speech intelligibility for more LTE configurations, and release a software package that includes all of the tools that we used for the LTE implementation so that the public safety, commercial, and academic community can make use of them.

References

[1] Walton C, O'Dell C (2020) Bridging Analog LMR to LTE MCPTT Communications (National Institute of Standards and Technology (NIST)), https://doi.org/https://doi.org/10.6028/NIST.IR.8338

[2] Riley GF, Henderson TR (2010) *The ns-3 Network Simulator* (Springer Berlin Heidelberg, Berlin, Heidelberg), , pp 15–34. https://doi.org/10.1007/978-3-642-12331-3_2. Available at https://doi.org/10.1007/978-3-642-12331-3_2

[3] Voran SD (2017) A multiple bandwidth objective speech intelligibility estimator based on articulation index band correlations and attention. *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp 5100–5104. https://doi.org/10.1109/ICASSP.2017.7953128

[4] Fast Forward Moving Picture Experts Group (FFmpeg), https://ffmpeg.org/. Accessed: 2025-03-03.

[5] 3GPP (2024) Technical Specification Group Services and System Aspects; Mission Critical Push To Talk (MCPTT); Codecs and media handling (Release 18) (3rd Generation Partnership Project (3GPP)), Technical Specification (TS) 26.179. Available at https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1420.

[6] Open Air Interface (OAI), https://openairinterface.org/. Accessed: 2025-03-03.