ANALYSIS OF PRIVATE LTE AND 5G RADIO ACCESS NETWORK DATA SETS FOR PERFORMANCE OPTIMIZATION

December 2022

Introduction

EPRI AI Initiative and Data Sets

EPRI's AI initiative is a collaborative effort funded by Technology Innovation for developing artificial intelligence (AI) and machine learning (ML) solutions. As a global R&D leader, EPRI's effort is intended to bridge the gap between the electric power industry and AI community by bringing various cutting-edge methodologies and models to the power and energy sector and to utilities. EPRI's collaborative work with researchers from utilities and the AI community to identify the application of AI/ML in energy sector can produce game changing results and a pathway to the power grid of the future.

EPRI.AI focuses on three major objectives: 1) creating an Electric Power–AI community, 2) creating solutions by building EPRI10 data sets, and 3) delivering results through grand challenges. EPRI10 is a collection of the 10 most valuable data sets in the electric power industry that will accelerate the industry's use of AI to transform utility operation. Through this, EPRI is releasing the data sets to the AI community for developing and training models. One such example is a curated data set of long-term evolution (LTE) network operation that is described in this paper.

LTE Networks

Modern wireless networks, such as fifth-generation (5G) and long-term evolution (LTE), generate massive amounts of data for the operation and control of the network. Full optimization of network operation is beyond human capability but needs to be further investigated with tools such as artificial intelligence (AI). This type of optimization is especially valuable for private utility networks that operate in a limited wireless spectrum while being required to support critical monitoring and control for grid operation. Today, wireless network data are used internally by basic algorithms that manage network functions such as accessing the network, establishing and maintaining the connection, providing the best data rate given changing environmental conditions, prioritizing and limiting traffic, determining roaming, handover, network selection, and many other autonomous functions. These data are also used by the network operator for network management, monitoring, situational awareness, and cyber-security analysis. Increasingly, these data are being combined and analyzed in new ways to enable further optimization through AI and machine learning (ML).

Private LTE Networks

Private LTE (PLTE) networks, shown in Figure 1, are being deployed by utilities as a standards-based solution for field communications. Before private 5G and LTE deployment, many utilities operated and maintained dozens of separate, proprietary wireless networks. The use of 5G and LTE for multiple grid applications brings together these aspects of operation and control into a single network. Private 5G and LTE networks deployed by utilities are unique in the limited amount of radio spectrum in which they operate. Although commercial cellular operators may spend billions to own hundreds of MHz of spectrum, utilities are very constrained. Most utility networks operate in a small spectrum allocation: 3 x 3 MHz (that is, 3 MHz each of uplink and downlink) is typical. Southern Linc has a unique history that allowed it to acquire 3 x 3 MHz spectrum in the 800 MHz band. Other utilities are building networks on 900 MHz spectrum, which is also 3 x 3 MHz. Given the constrained spectrum resources and the critical operational functions of the network, it is even more important for utilities to fully optimize their networks to get the most value out of the available spectrum while maintaining reliability. Another driving factor is the exponentially increasing complexity of networks. Utilities are challenged to find skilled staff and train them in the basics of network deployment and operation, let alone have them become experts in network optimization.

Table of Contents

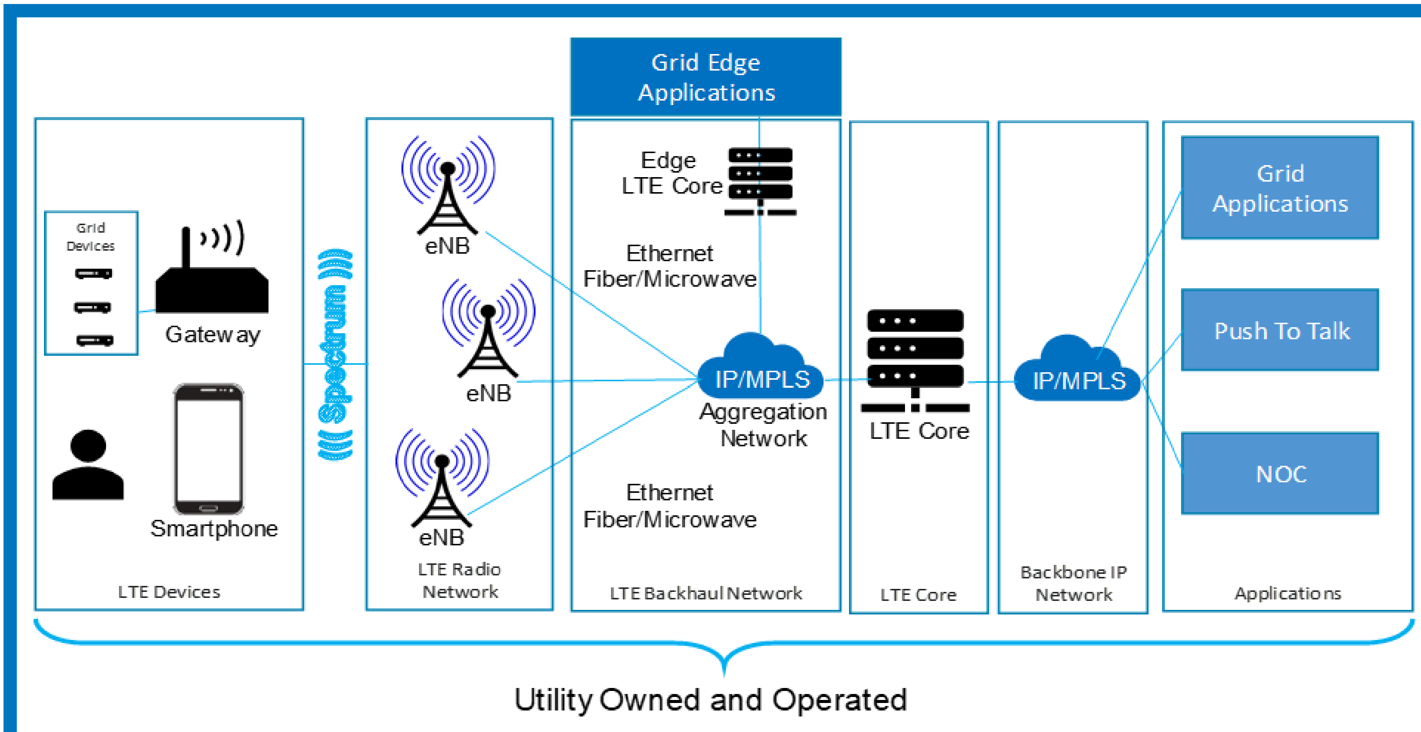

Figure 1. Private LTE Network Architecture

Table 1 provides private LTE band examples.

Table 1. Private LTE band examples

Common Name	3GPP Band Designation	Frequency Range, MHz (UL/DL)	Duplex Type	Channel Bandwidth
Southern Linc	Band 26	813.5–817.4/858.5–862.4	FDD	3.9 MHz
Anterix	Band 8	897.5–900.5/936.5–939.5	FDD	3 MHz
CBRS	Band 48	3550–3700	TDD	Fifteen 10 MHz channels

Network Architecture Domains

The Evolved Universal Terrestrial Radio Access Network (EU-TRAN; shown in Figure 2) is the domain that provides wireless connectivity (access to the network) based on 3GPP standards. The term ***EUTRAN*** is used to distinguish a 4G (LTE) RAN from prior generations (UMTS for 3G and GSM for 2G). It is commonly referred to simply as the ***RAN***. The complete network consists of various components that have been derived from 3GPP and earlier generation UMTS systems. Some of the components/terms are listed next.

User Equipment Domain

- **User equipment (UE):** The device that connects wirelessly to the network, which may be a phone, cellular modem, or equipment with embedded connectivity.

Radio Access Network Domain

- **Evolved Node B (eNodeB or eNB):** Traditionally known as the ***base station,*** this is the fixed infrastructure equipment to which the UE connects. It may be located on a tower or on a pole or wall in the case of "small cells."

- **S1 Interface:** Colloquially known as ***backhaul***, the S1 interface connects the set of eNBs to the EPC.
- **X2 Interface:** Typically implemented physically over the same transport network as S1, the X2 logically provides interconnection between eNBs.

Core Network Domain

- **Evolved Packet Core (EPC):** The centralized functionality that integrates the set of eNBs into a unified network. It is composed of Mobility Management Entity (MME), Serving Gateway (S-GW), and Packet Gateway (P-GW).
- **Mobility Management Entity (MME):** Supports subscriber authentication, roaming, and handovers to other networks.
- **Serving Gateway (S-GW):** Provides a common routing endpoint for UEs and manages mobility of UEs between eNBs within the network.
- **Packet Gateway (P-GW):** Routes packets between the RAN and an external IP network (Internet or utility private network).

Over the wireless interface between the UE and the eNB, the standards for LTE describe two methods of structuring the simultaneous communication in the uplink (from UE toward eNB) and downlink (from eNB toward UE):

- **Frequency-division duplex (FDD):** A system in which the uplink and downlink occupy different radio channels separated in frequency.
- **Time-division duplex (TDD):** A system in which a single frequency radio channel alternates between uplink and downlink operation, operating in each direction for a short time interval.

FDD originated with 1G analog cellular, where simultaneous analog voice circuits simulated the landline telephone experience. A separate frequency was used for the voice in each direction. These

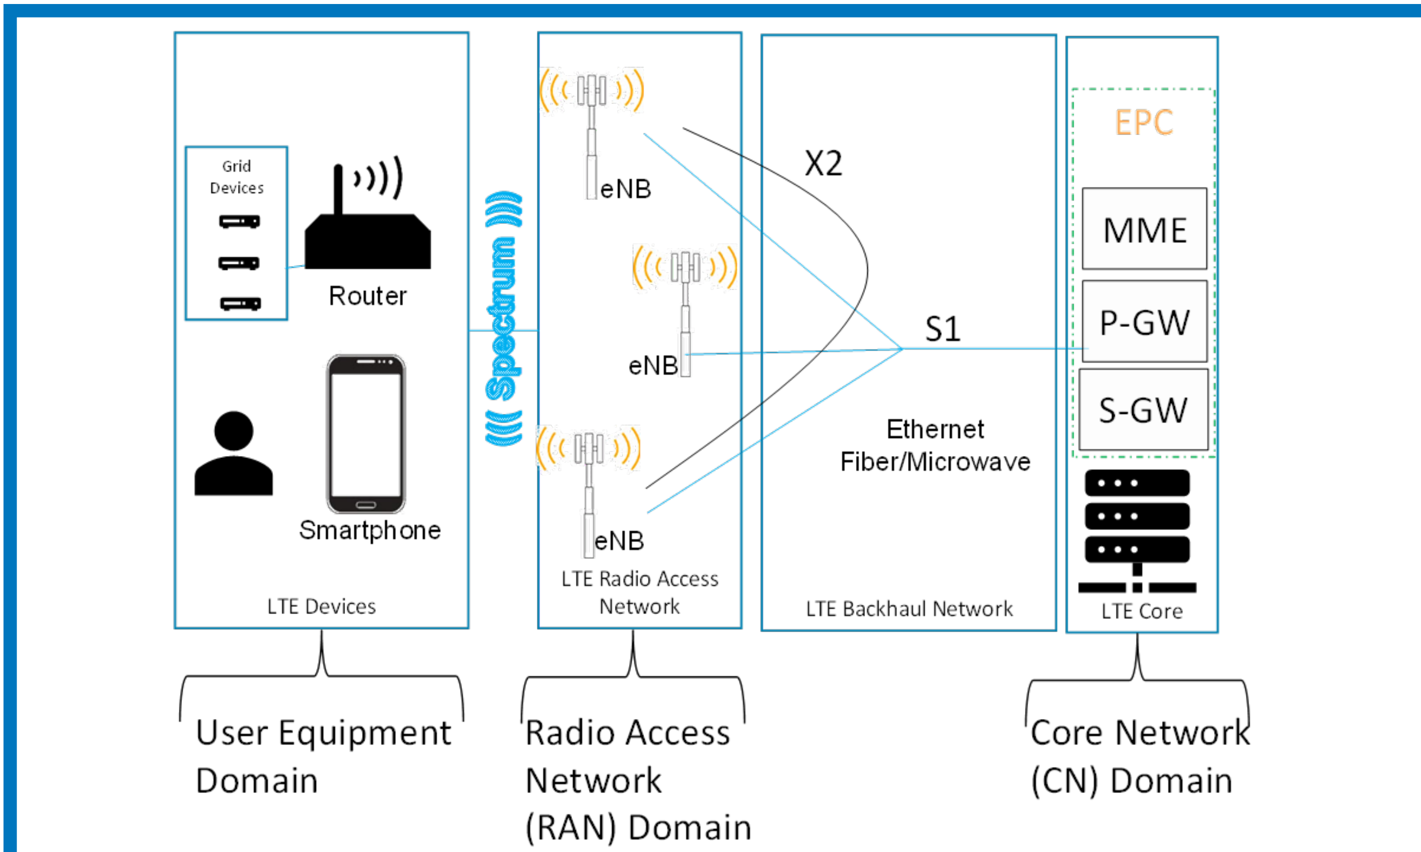

Figure 2. EUTRAN

frequencies had to be widely separated to allow the mobile phone to transmit and receive simultaneously without interference. This uplink/downlink frequency separation, also called ***duplex spacing*** or ***gap***, must be wide enough for RF filter components in the UEs and eNB to isolate the two channels from one another. In early cellular systems, a 45 MHz duplex gap was used, but technology improvements have allowed this to be reduced to 10 MHz in newer systems. This style of paired spectrum allocation has continued through all the cellular generations. FDD was the only supported option through 3G, and the option of TDD was introduced with 4G LTE. Although FDD has a fixed allocation for uplink and downlink, TDD offers the potential for more flexibility by varying the ratio of uplink to downlink time as dictated by the traffic flow in each direction.

The data set examined for this project is based on Band 26 (800 MHz) FDD. Southern Linc is also starting to deploy Citizens Broadband Radio Service (CBRS) sites using Band 48 (TDD) at 3.6 GHz, but data for those sites were not available during this project. In addition to the RAN data analyzed here, a CBRS network also can provide data from the Spectrum Access System (SAS).

The CBRS band is the first spectrum allocation to implement a spectrum sharing approach. The requirements of federal incumbents that could not easily be relocated resulted in the creation of spectrum sharing rules based on spectrum sensing and databases for controlling device access to the spectrum. The SAS is a commercial service designated by the FCC to provide authorization to use CBRS spectrum.

The eNodeB (or base station) installed as infrastructure for private LTE networks in the CBRS band is referred to as ***CBRS Service Devices (CBSDs)***. A CBSD must obtain authorization from the SAS before it can transmit in the CBRS band. A group of SAS administrators has been designated and approved by the FCC.

Future analysis of the SAS data could reveal patterns and probabilities in the way CBRS network operating frequencies and/or power limits may be changed by the SAS based on incumbent activities and the resulting impact on availability and network capacity.

Research Objective

The purpose of this research was to gather a requisite network data set to build a data-centric architecture capable of supporting data-driven efforts such as the development of AI applications, software, and tools. Data-driven applications can help optimize 5G and LTE networks by reducing labor costs while improving the performance and reliability of the network. This will facilitate widespread use of 5G in the industry and enable the realization of more of the associated benefits from wireless networks. In this research, EPRI obtained an LTE network data set from a domestic industry utility. The data set is 12 months of operational data from the utility's private network. Having such a data set builds the foundation for a data-centric architecture and is invaluable for any future data-driven application. Another objective of this research is to evaluate and describe the data set by assessing its content and taxonomy and providing recommendations on how to work with this data set as well as how it can be further used.

This data set will enable the enhancement and optimization of utility-owned wireless networks supporting critical operational communications by gathering data from the network itself. It can counter a potential misperception that 5G is simply a neutral medium used only to transport data from other types of devices to enable consumption by AI operating in other realms. This project sets the stage for further work in this area, engaging with additional utilities and vendors to expand optimization for utility operational networks.

Finally, EPRI is in a unique position to work with utilities to identify, gather, and use data sets from private 5G and LTE networks that operate the electrical grid. These data are different from what would be gathered from commercial cellular operators, which have different business goals and customers. Private utility networks have unique use cases and data flows and will therefore have unique optimizations.

Cellular Background

An introduction to certain cellular concepts is needed to understand the meaning of some aspects of the data set and how it relates to performance optimization.

Sector Size, Downtilt, and User Density

Typically, cellular infrastructure for utility private LTE or 5G networks is optimized for the broadest coverage. This is accomplished by mounting antennas on top of towers (see Figures 3 and 4). Typically, three antennas are installed on a tower, each serving a 120-degree sector of the azimuth around the tower.

In the sections that follow, data from a tower site will be categorized as Sector 1, Sector 2, and Sector 3.

Cellular antennas also support downtilt, which allows the "beam" of radio-frequency (RF) energy to be directed toward the ground instead of straight toward the horizon (Figures 5, 6, and 7). Downtilt can be used to limit the cell size to prevent interference with adjacent sites operating on the same frequency. Downtilt is a system configuration parameter and may be varied remotely. In addition, it may be mechanical (where the antenna unit is physically moved by a motor in the mounting) or electronic (most prevalent in array antennas or multi-input-multi-output [MIMO] antennas in which phasing and beamforming are used to electronically steer the energy up or down). The same multi-element antenna structures that support electronic beamforming for downtilt (vertical axis) often support beamforming in the horizontal (azimuth) direction, which can help with range and reduce interference.

Figure 3. Cellular sector antennas on a tower

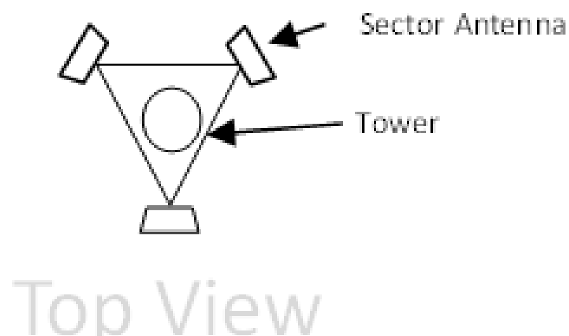

Figure 4. Top view of three-sector tower and antennae

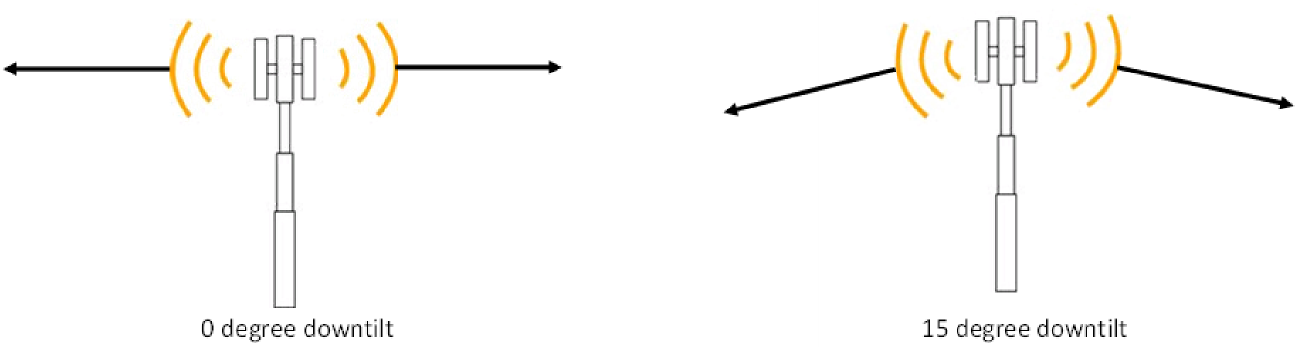

Figure 5. Depiction of downtilt from side view of tower

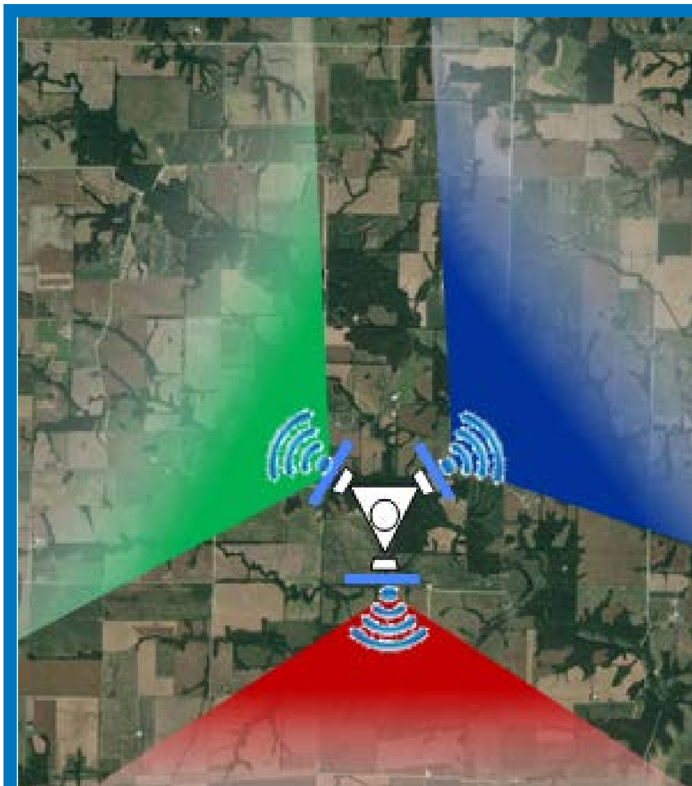

Figure 6. Coverage of tower site: no downtilt

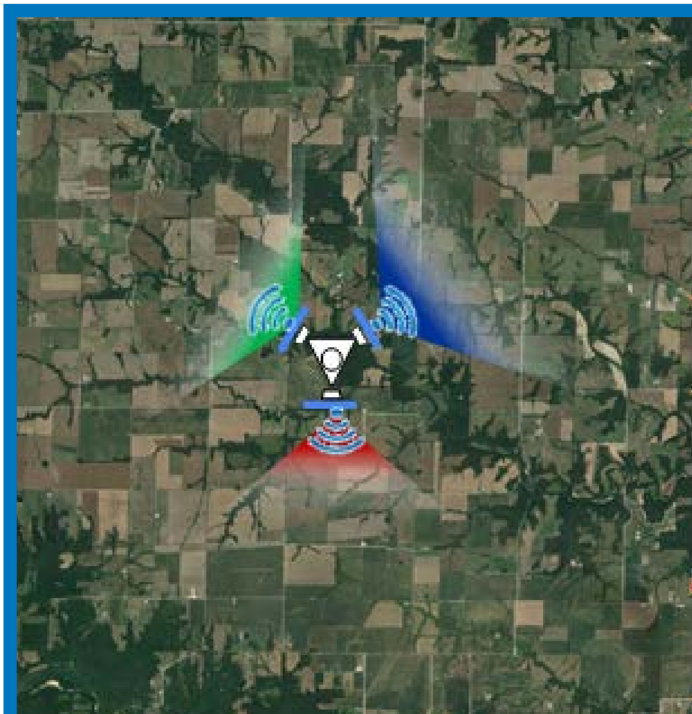

Figure 7. Coverage of tower site: with downtilt

Density and Capacity Management

In some utility networks, when the number of connected devices is low, it is desirable to have the largest cell size so that the territory can be covered with the smallest number of cell sites (which are expensive to install). As the network grows and more devices are added, the network will eventually become capacity constrained. This happens when the total amount of data transmitted and received by all the devices in the coverage area of a sector exceeds the bandwidth capacity of that sector.

The solution to capacity limitations is cell splitting, where additional cell sites and towers are installed to reduce the cell size and reduce the number of devices per sector. Figure 8 shows a single three-sector site, with about 16 user devices per sector. In Figure 9, the cells are subdivided, and additional towers and sectors are used, resulting in about four devices per sector.

The solution to capacity limitations is cell splitting, where additional cell sites and towers are installed to reduce the cell size and reduce the number of devices per sector. Figure 8 shows a single three-sector site, with about 16 user devices per sector. In Figure 9, the cells are subdivided, and additional towers and sectors are used, resulting in about four devices per sector.

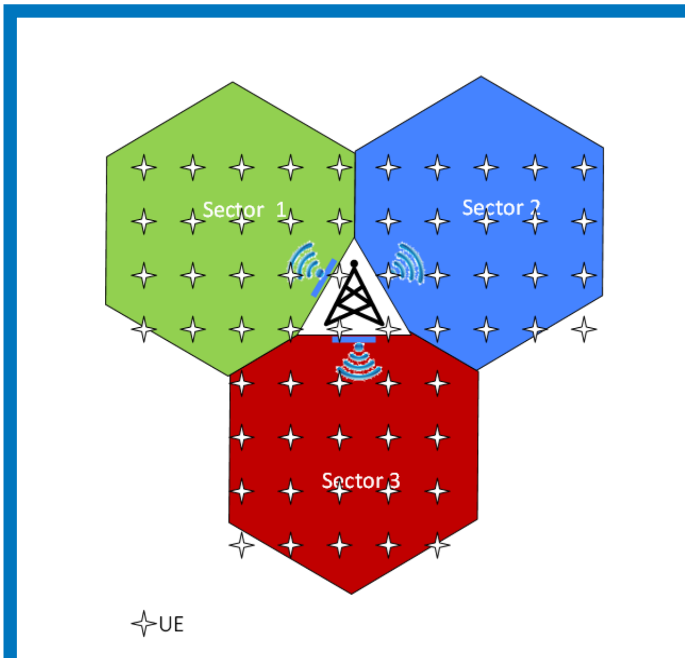

Figure 8. Three large sectors

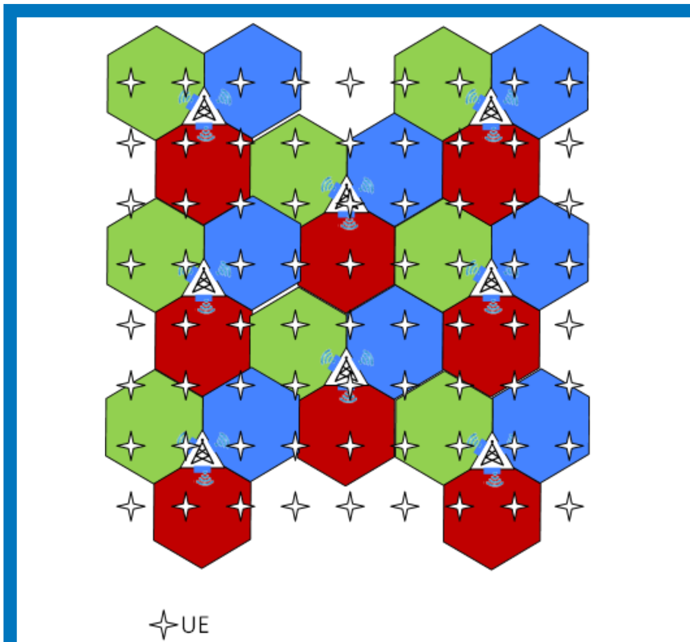

Figure 9. Smaller sectors after cell splitting: fewer devices per sector

Sector size and cell splitting can be managed by a combination of RF power control and antenna downtilt. The optimum solution for performance and maximizing capacity depends on the area topography and its effect on RF propagation as well as user density and position.

Service Provisioning and Management

In 3GPP networks (LTE and 5G), the network carefully manages the use of the RF spectrum resources (time and frequency). Transmissions on the uplink and downlink are scheduled and prioritized. The network can support multiple applications and services at a single device with varying performance requirements. For example, a cellular device at a substation might be provisioned to carry supervisory control and data acquisition (SCADA) data, voice, and IT data for email and work management. Each service with unique requirements for performance and reliability can be configured to use a logical data "pipe" through the network. This is called a bearer.

Bearers are established when there are data to communicate. The network uses Radio Resource Control (RRC) signaling (also called control plane) to create and terminate bearers to handle data requirements. The number of bearers and their utilization can be optimized to ensure that service requirements are met and to get the most overall throughput from the network. Establishing and terminating bearers takes time and uses network resources, so optimization of their rules and timing for managing bearers is also needed.

Data sets described next will refer to bearers, which are an indicator of user data flow.

Data Set Description

In this section, the contents and taxonomy of the LTE data set is described. A few of the tools used by EPRI to work with and handle this data set are also mentioned. The importance of this section is to provide a high-level overview of the data set so that data analysts, scientists, and database engineers are aware of the challenges of handling big data.

The data set was originally provided to EPRI as *.gz compressed files. A total of 160 files makes up the one year's worth of operational data, and the total file storage size in the compressed state is 166 GB. EPRI started by decompressing the data set to allow more control over accessing and reading the data files. In the uncompressed state, the total file size is 1.68 TB. A data set of this magnitude does require consideration of the hardware and software tools used in its analysis and processing. The data cover the one-year period from April 2021 to April 2022.

The data set contains the configuration settings and performance indicators of eNodeB (Evolved Node B) and EUTRAN (Evolved Universal Terrestrial Radio Access Network) by cell. The data for configuration management are typically recorded once every 24 hours or day, and performance management is recorded at an interval of every 15 minutes throughout a day.

Digital Infrastructure Considerations

The enormous amount of data that was involved in this project took our data scientists effort to undergo various software and hardware challenges. In particular, the data set is much larger than RAM that needed to have better configuration devices such as servers, which increases high upfront capital costs or requires cloud-based solutions, which has its own challenges. On the other hand, various other software is particularly designed to handle large amounts of data. New software comes with a learning curve and getting the software approved to use in the organization is a time-consuming process. Other challenges include data governance that investigates tools and legalities, network considerations for the data transfer stage, data backups in case of data loss, data security, and staff expertise such as data engineers, IT, researchers/scientists/data scientists.

Exploratory Data Analysis

The configuration management data for EUTRANFDD consist of 1,523,823 samples of data with 524 attributes recorded at an interval of 24 hours for a year. Similarly, performance management data for EUTRANFDD consists of 122,443,207 samples of data with 1,310 attributes recorded at an interval of 15 minutes throughout a year. Due to the huge amount of data, performing data analysis required a great deal of consideration for digital infrastructure.

Exploratory analysis was performed by subject matter experts (SMEs) to identify the sensitive data; the identified data were masked with the appropriate map. During this process, some of the attributes marked as valuable. In addition, the data scientists performed analysis on configuration and performance management data to preprocess, clean, and filter out the attributes they may not use in performance optimization for reasons such as empty columns, features with all zeros, or non-varying values with a standard deviation of zero. Therefore, the data set size was reduced from 524 to 85 attributes for configuration management and from 1310 attributes to 465 for performance management. The curated data will be available as compressed csv files when the data is released.

Data Statistics: EUTRANFDD

In the data set we received, there are a total of 1,287 eNodeBs and 3,843 cells with around 3 cells/eNodeB. This is expected based on the most common configuration of three cell sectors per tower site as described above. A few sites have one or two sector antennas on a tower—for example, when the tower is located on the coastline or at the edge of the service territory, the sector facing the unserved area will not be installed.

Figure 10 shows the number of active data radio bearers in the downlink direction and uplink direction for each cell on each eNodeB. An active bearer is a proxy measurement for utilization. A bearer is like a pipe that is set up to carry data traffic. Having an active bearer does not mean that data are flowing and does not relate to data throughput or volume. However, if user equipment (UE) does not have an active bearer, no user plane data are possible. If data need to move in either direction, a bearer must be active, which is enabled by control plane (RRC) signaling. Networks are typically configured to move a UE to the RRC idle state and deactivate bearers after a timeout period with no data traffic. Because it takes time to reconnect and activate a bearer, a private LTE network operator can extend the timeout for UEs that are serving latency-sensitive use cases. Because the RRC signaling to move to the RRC connected state and establish bearers requires network bandwidth and resources, one optimization based on the active bearers is to look for UEs that are frequently changing between connected and idle states. This condition indicates that the periodicity of the user data is not aligned with the timeout. Extending the timeout may allow the UE to remain connected with the active bearer, which will reduce application latency as well as unnecessary RRC signaling.

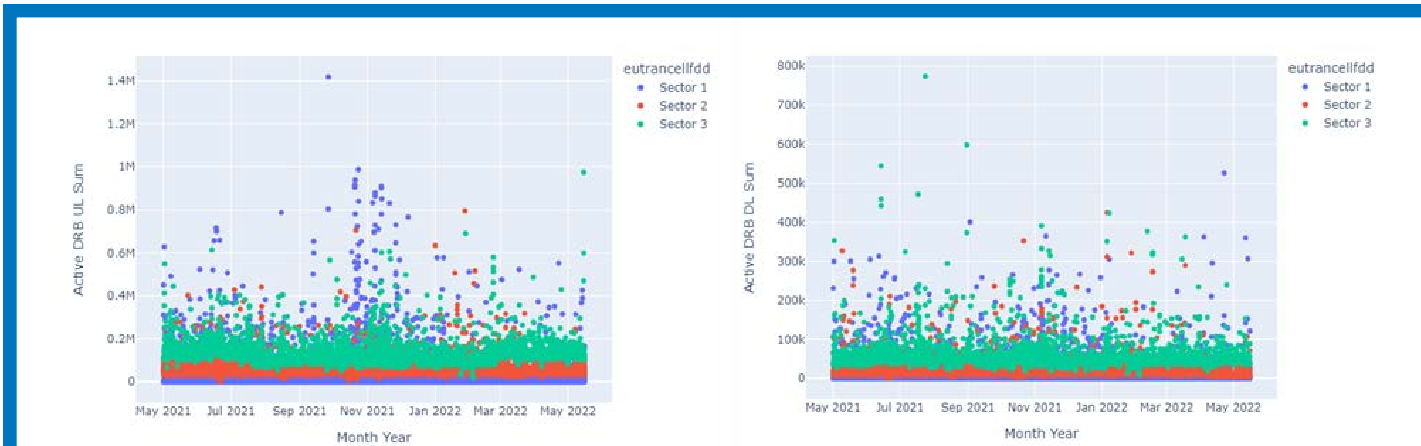

Figure 10. Number of active data radio bearers in the downlink direction and uplink direction for each sector on an eNodeB

A UE will have a minimum of one active bearer to transfer user data, but in some cases multiple bearers may be used from the same UE. Bearers can be established with different performance parameters and priorities (referred to as q***uality of service [QoS] Class Identifiers*** or ***QCI***). For example, a single UE in a utility's distribution control system could serve a protection relay and a cap bank controller. The protection relay's data could be assigned to a bearer with high-priority, low-latency QCI settings, while the cap bank controller could use a bearer with a QCI providing low-priority "best effort" service.

Due to enormous amount of data when plotted for the whole year, it is difficult to infer or look at the plots. Therefore, the data are sampled to per weekly from per year by averaging the data points over every week as shown in Figure 11 for an example eNodeB. In the data set the masked name is G-0001.

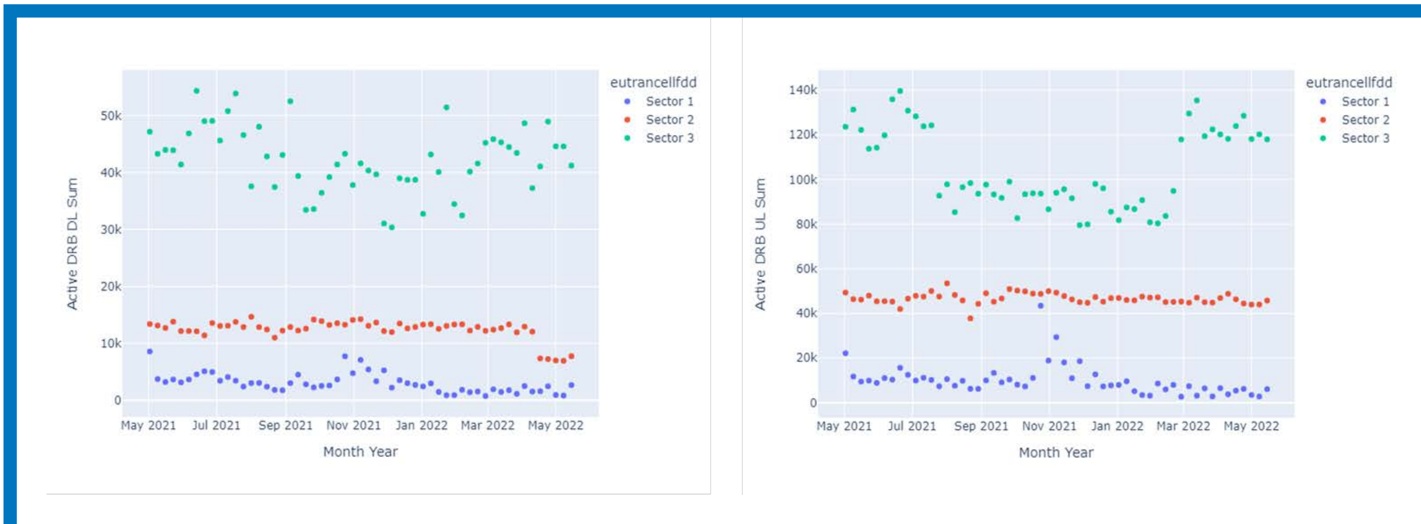

Figure 11. Number of active data radio bearers in the downlink direction and uplink direction for each sector on an eNodeB (Sampled Weekly and Averaged)

The figures show three sectors at same location separated by an angle. The number of active bearers in each sector is related to the number of active UEs in that sector, as described above. Although the number of active bearers can change due to active/idle transitions of UEs, the count can also change due to handover of UEs between sectors. Handovers may be caused by changing signal strength conditions (the UE is handed over to the sector with the best signal strength) or to manage capacity and loading of sectors. Signal strength changes leading to handover often result from a mobile UE moving out of coverage from one sector into another. In a private utility network, many UEs are devices in fixed positions, on poles or other equipment. They still exhibit some level of change in signal strength due to weather or seasonal foliage changes.

Changes in active bearers may also be a result of network-initiated handover to reduce congestion. If a particular sector is overloaded with users or data, the network can look at UE signal measurements and determine UEs that could be successfully handed over to a neighboring sector with lower loading.

Optimizing the handover process and the distribution of UEs across sectors is most effective when the geographic location of the UEs is known. A UE may attach to two adjacent sectors on a single eNodeB or may attach to sectors on a different, nearby eNodeB. Southern Linc was not able to include UE-specific data in the data set analyzed because its network includes third-party customers as well as devices owned by Southern Company and disclosing data of the third-party customers was not possible due to confidentiality requirements. For anyone analyzing data sets for their own private LTE or 5G network, the data from UEs—including their location and attached eNodeB—can be cross-referenced with the eNodeB data on active bearers to provide a more complete picture of handover behavior and improve the ability to optimize cell edge performance.

LTE Key Performance Indicators

Key performance indicators (KPIs) are used to measure and monitor the network's performance. Monitoring KPI values plays a major role in the optimization processes that in turn enables better use of resources and provides better subscriber quality. KPIs are used to detect and troubleshoot issues in the cell or network. Generally, KPIs are divided into six categories:

1. **Accessibility KPI:** measures the success of devices in accessing the network when they attempt to establish a connection. This includes the random access metric (the first step for a UE establishing a connection) and the session setup success rate, which measures the overall success of the control plane signaling that initiates a connection.
2. **Retainability KPI:** measures the success of devices maintaining connection and access to the network's services.
3. **Mobility KPI:** measures the ability of the network to handle the movement of users and still maintain service for the user, such as handover.
4. **Integrity KPI:** a measurement of whether the network is meeting service-level agreements for the user, such as the throughput and latency that users were provided.
5. **Availability KPI:** measure the extent to which the network is suitable or ready for users to access services.
6. **Utilization KPI:** measures network loading, or whether the network capacity has reached its limits due to spectrum or other limited resources.

KPIs are usually directly measured or calculated based on use case and utility interest. Some of the KPIs of utility interest are discussed next.

- **UE session time:** Defined by the sum of session time taken by all user equipment in particular cell. These data are recorded in seconds. The data or time is recorded only if any data transfer occurs or lasts for more than 100 milliseconds, which could either be an uplink or downlink process. UE session time provides a more representative measurement of UEs that are actively transferring data in the sector. Although an active bearer indicates a UE that is prepared to transfer data, UE session time measures time periods of active data transfer. It does not measure volume of data but a duty cycle of activity, in aggregate, for the sector. (Figure 12).

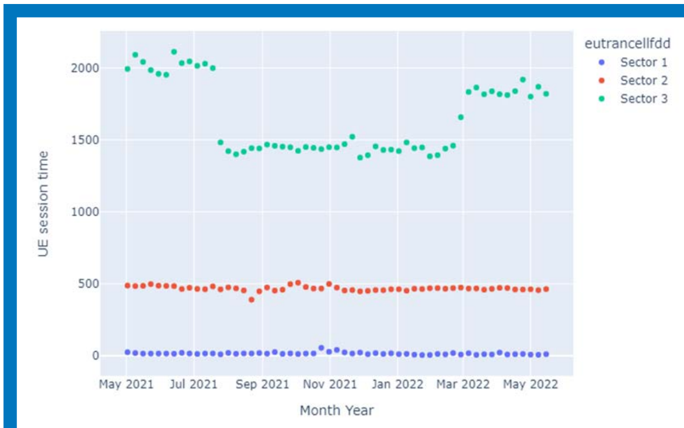

Figure 12. UE session time per sector on an eNodeB

The data show Sector 3 with significantly higher UE session time overall and an extended reduction lasting for over seven months. The data are for active seconds in 15 minute (900 second) measurement intervals. If a sector contained a single UE that was active 100% of the time, the UE session time would be 900. To measure the average activity factor for all UEs in the sector, the UE session time should be divided by the number of active UEs in the sector. The value pmRrcConnMax can be used because it measures the number of active RRC connections during a measurement interval. That value could clarify whether the drop was due to a drop in the total number of UEs (possibly due to their being connected in another sector) or if the UEs remained in the sector but were actually less active in terms of transferring data. The data in the next section show that the dip is indeed correlated to a similar reduction in connected UEs.

This measurement gives a picture of activity but does not directly yield the capacity because a UE session will be recorded as active for a given second if there is a single 100 ms data transfer or the bearer is continually transferring data.

- Average number of radio resource control (RRC) connected users (UEs): Defined by ratio of number of UEs that are connected in radio resource connected mode to the total number of times the corresponding counter has received a new sample. As noted previously, when a UE is in the RRC connected state, it can have an active bearer and transfer data. After a timeout due to inactivity, the network moves the UE to the RRC idle state and releases bearers and other resources (Figure 13).

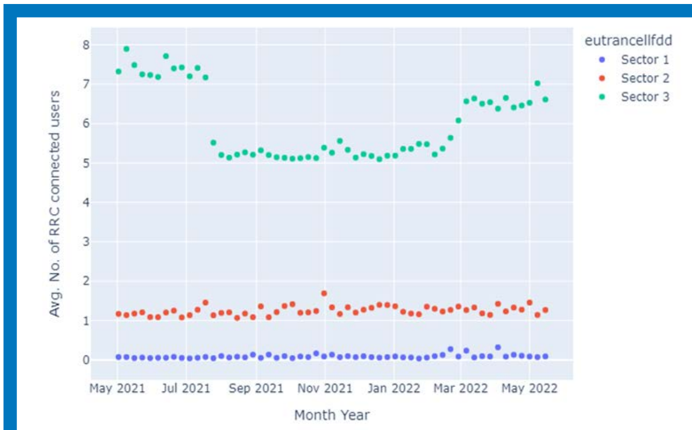

Figure 13. Average number of RRC connected users for each sector of an eNodeB

The data show that Sector 3 for this cell had the most average users, which dropped from 7–8 users to 5–6 users for a period of seven months and then went up to 6–7. This could be due to the connected equipment being out of service or that some UEs were connecting to a different cell during those months.

- **E-RAB attempts:** The KPI for accessibility is the metric for establishing the Evolved Radio Access Bearer (E-RAB), which involved successful signaling and setup between the eNodeB and the Mobility Management Entity (MME) in the core. This value is defined by the sum of initial and added E-RAB establishment attempts. An "added" E-RAB means that an additional bearer is established from a UE after the first bearer has been established. These initial and added E-RABs are those that are present in the S1 message E-RAB setup request. This KPI provides insights on network congestion or resource limitations that can prevent service establishments (Figure 14).

Figure 14. E-RAB attempts

The data on E-RAB attempts show points in time when Sector 2 and Sector 3 experience congestion resulting in much higher E-RAB attempts. Sector 1 is mostly unaffected but, as seen in previous sections, Sector 1 has few active users. There is no clear temporal pattern evident in the E-RAB attempts data at this level, but deeper analysis at a finer time resolution could reveal daily or other periodic patterns related to user behavior or grid events.

- **RRC attempts:** The Radio Resource Control (RRC) protocol is the signaling used in the control plane of an LTE or 5G network. The RRC Attempts KPI (pmRrcConnEstabAtt) is defined by the total number of attempts or requests made for establishing an RRC connection—for example, causing the UE to enter the RRC connected state. It is one of the components of the calculation of E-RAB attempts (described previously). (Figure 15).

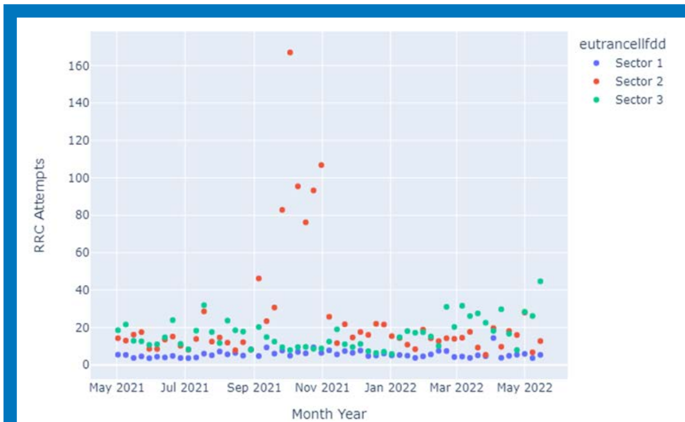

Figure 15. RRC attempts

The data above illustrate that a subset of the E-RAB attempts data above were a result of RRC signaling congestion. Other causes of increases in the E-RAB attempts values are congestion on the S1 interface (network backhaul between eNodeB and core) or capacity limits in the core itself.

- **Peak number of RRC connected users:** Defined as the maximum number of UE in the RRC connected state, excluding the bandwidth-reduced UE such as CAT-M1 devices. (Figure 16).

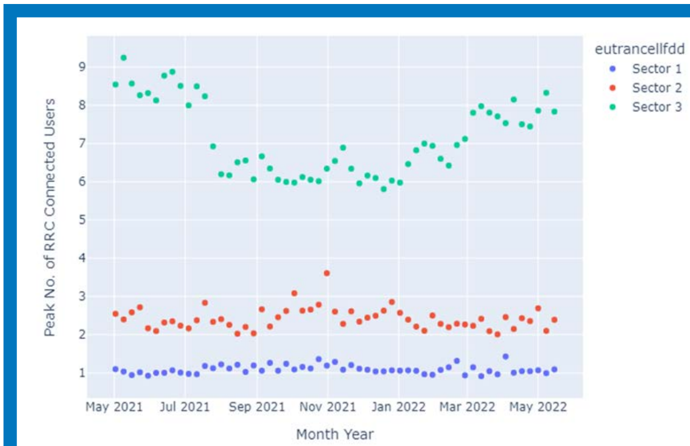

Figure 16. Peak number of RRC connected users

This metric tracks the preceding sections, reflecting the highest number of connected UE on Sector 3, and the drop-off in numbers during the middle seven-month period.

Summary

The preceding sections highlight KPIs that were selected from the data set as having notable variations and representing meaningful real-world performance measurements of the network. When using this analysis with another data set from a different private network in the future, these KPIs are a good place to start. However, more specific opportunities for optimization will be possible if the UE data can be obtained and cross-referenced to the RAN data corresponding to the data set analyzed here.

Applications and Use Cases

LTE SMEs point to remote electronic downtilt (RET) as a key potential ML application. Currently this is a setting that system operators will adjust during initial system test and turnup, but it may also be manually operated in response to events, such as a large crowd of users appearing at a venue. By manually adjusting RET within a cluster of neighboring cell sites, the operator can shrink coverage of some cells and enlarge others to load-balance the traffic. Automating this with an AI/ML routine is leading edge and starting to be offered by equipment vendors. Utility private networks, however, have different use cases that could make use of automated RET. In one such use case for disaster recovery—in which storm damage impacts the PLTE and critical DA devices are unable to connect to their primary (closest) cell sites—automated RET could be used to enlarge surviving cell site coverage area and pick up the disconnect DA devices.

Congestion management is another potential ML application. As mentioned previously, utility PLTE implementations are typically constrained to operate with much smaller bandwidth (3x3 MHz) than commercial public carrier systems (100x100 MHz or larger). There are utilities that plan to use PLTE not only for critical infrastructure machine to machine (M2M) use cases such as teleprotection relaying, reclosers, and capacitor bank control, but also as a replacement for their two-way radio systems. The two-way radio systems may support internal workforce only or in other cases extend service to external entities such as government or other industrial users. In such a situation, the M2M traffic requests would be configured with a higher QCI than the voice users because the latter are able to endure a slight delay in access to a bearer. Ideally the network operator would monitor the system for signs of congestion and proactively manage configuration settings to avoid a degradation in service quality. Skilled and experienced technicians are required to accomplish this task. It would be beneficial if AI/ML could perform this work instead.

Additional Resources

Because the purpose of the project was to enable further development by AI researchers, the data set is being made available. Contact the project manager for access. Please be prepared to describe your proposed use of the data set. The releasable data set will use the column set found in the first months of data and exclude additional columns added after February 2, 2022.

Summary

One year's worth of configuration and performance management data was collected from an operational PLTE network. An exploratory data analysis was performed, resulting in identification of valuable fields that contain varying values and relationships to network KPIs. A compressed .csv file was written to extract the curated data from the full data set. Both the .csv file and access to the full data set may be available to AI researchers and data scientists who have proposals to develop AI/ML applications for PLTE network operators.

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

EPRI PREPARED THIS REPORT.

Note

For further information about EPRI, call the EPRI Customer Assistance Center at 800.313.3774 or e-mail askepri@epri.com.

About EPRI

Founded in 1972, EPRI is the world's preeminent independent, non-profit energy research and development organization, with offices around the world. EPRI's trusted experts collaborate with more than 450 companies in 45 countries, driving innovation to ensure the public has clean, safe, reliable, affordable, and equitable access to electricity across the globe. Together, we are shaping the future of energy.

EPRI RESOURCES

Project Manager

Tim Godfrey
650.855.8584, tgodfrey@epri.com

Principal Investigators

S. Bellary, M. Dunlap, J. Estrada, J. Herman

Technology Innovation

3002023926 December 2022

EPRI
3420 Hillview Avenue, Palo Alto, California 94304-1338 • USA • 800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com