April 2022

VRAN CAPACITY ENHANCEMENT SOLUTION BY RADIO UNIT SELF-DISCOVERY AND SELF-MANAGEMENT

Fred Anderson

Shailender Potharaju

VRAN CAPACITY ENHANCEMENT SOLUTION BY RADIO UNIT SELF-DISCOVERY AND SELF-MANAGEMENT

AUTHORS:
Fred Anderson
Shailender Potharaju

ABSTRACT

In a 3rd Generation Partnership Project (3GPP) fifth-generation (5G) virtualized radio access network (vRAN) architecture, a shared cell consists of one distributed unit (DU) instance and a group of radio units (RUs) sharing the same radio frequency (RF) parameters. Such a group of RUs act as a single cell, covering a large contiguous area in a way that is not possible with a single RU thus significantly improving the signal strength and the signal-to-noise ratio (SNR) for the users that are in the overlapped coverage regions of the multiple RUs. While a shared cell approach addresses coverage problems with limited frequency spectrum, it suffers from a capacity issue. For example, using 3GPP-specified Synchronization Signal Block (SSB) signaling only eight RUs may be defined within a shared cell. Such a restriction can be very limiting in venues where greater than eight RUs are required to provide handover-free coverage areas. To address the types of challenges that were described above, techniques are presented herein that support a novel shared cell architecture that, for example, is suitable for enterprise use cases. Aspects of the presented techniques comprise self-discovery and organization, an unconstrained number of RUs, and low overhead. Further aspects of the presented techniques encompass the self-discovery of user equipment (UE) and RU relative positions and an exploitation of the measured RF isolation between those entities. Additionally, a smart scheduling algorithm may then be applied which multiplies cell capacity in proportion to the RF isolation.

DETAILED DESCRIPTION

In a 3rd Generation Partnership Project (3GPP) fifth-generation (5G) virtualized radio access network (vRAN) architecture, a shared cell consists of one distributed unit (DU) instance and a group of radio units (RUs) sharing the same radio frequency (RF) parameters such as Physical Cell ID (PCI) and center frequency. Under such an arrangement, a transmission from one DU is conveyed over multiple RUs. Such a group of RUs act as a single cell, covering a large contiguous area in a way that is not possible with a single RU. Such a mechanism significantly improves the signal strength and the signal-to-noise ratio (SNR) for the users that are in the overlapped coverage regions of the multiple RUs. In such a manner, most of the users in the coverage area of the shared cell experience uniform signal strengths no matter where they are in the network. This significantly improves per-user throughput and reduces call drops due to coverage holes. Another advantage of a shared cell is that handovers are completely eliminated between the RUs. A user that is moving from one RU to another RU is considered to be intra-cell, so no additional handover overhead is involved.

While a shared cell approach addresses coverage problems with limited frequency spectrum, it suffers from a capacity issue (in terms of throughput and the number of users) when it is compared to single-cell deployments with a least interference between the cells. There are ways to increase the capacity of shared cell deployments using techniques similar to multi-user, multiple input, multiple output (MU-MIMO) technologies. Downlink (DL) and uplink (UL) transmissions for more than one user equipment (UE) may be scheduled on overlapped Resource Blocks (RBs) as long as they have sufficient isolation to avoid crosstalk or interference between the users. This can increase the capacity of a shared cell while limiting the interference.

In order to implement the technique that was described above, the Layer 2 (L2) scheduler must have sufficient knowledge about the position of the users to estimate the interference between them.

Currently, in order to implement shared cell mechanisms, UEs measure separate, distinct reference signals from each RU (e.g., Synchronization Signal Block (SSB) beam reference signals). The measurements are used to estimate a SNR for UL and DL transmissions to those respective RUs so the scheduler can localize the RU source and destination of Physical Resource Blocks (PRBs). Limitations of this approach include increased overhead (e.g., non-overlapping reference signal scheduling, increased physical downlink control channel (PDCCH) traffic, etc.), and most importantly, restrictions on how many RUs may be included within a shared cell. For example, using 3GPP-specified SSB signaling only eight RUs may be defined within a shared cell. This can be very limiting in venues where greater than eight RUs are required to provide handover-free coverage areas.

Consequently, a shared cell implementation method is needed which adds very little overhead to cell management and has no limitation on the number of RUs within a cell (and therefore the number of RUs that are associated with a single distributed unit (DU)).

To address the need that was described above, techniques are presented herein that support a novel shared cell architecture. Inspiration for the presented techniques comes from the need to offer a simple, effective, and self-deploying solution for the Layer 1 (L1) and L2 collaboration of RUs in a vRAN system. Such a facility is particularly important to network equipment vendors as it offers a product solution for enterprise or commercial private 5G environments where an ease of deployment, along with minimal technical expertise required for installation, are not only a selling point but a key to widespread adoption.

Aspects of the presented techniques enable a self-managing Radio Access Network (RAN) solution that achieves capacity and coverage for the most sophisticated shared cell deployments, without the need for added sophistication or cost in RU hardware. RUs may self-discover any coverage and interference overlap, and therefore the number of RUs that are deployed is not limited by a discrete (and small) number of beams and/or reference signals that are utilized by conventional shared cell solutions. Similar to self-managed WiFi networks, site survey requirements are drastically reduced and the adaptive nature of the presented techniques guarantees that UEs receive the coverage and capacity that they need while the system efficiently reuses the spectrum.

Aspects of the techniques presented herein support a low-overhead method which may be used to approximate UE position and RU position, thereby allowing for the prediction and the exploitation of the RF isolation between those entities to implement a smart scheduling algorithm which multiplies cell capacity.

In a shared cell network, it is possible to reuse the resource blocks (RBs) that are scheduled for one user on one RU for another user on a different RU as long as the RF isolation is sufficient between the RUs and the users. Such reuse will increase the system capacity by some multiplier depending upon how many users may be stacked on an RB. While such an approach is similar to a MU-MIMO layering technique, it does not use the spatial multiplexing concept or beamforming to further enhance SNR. However, SSB beam forming techniques that are available in a 5G environment may be leveraged along with aspects of the techniques presented herein to further increase performance.

An exemplary shared cell network deployment is depicted in Figure 1, below.

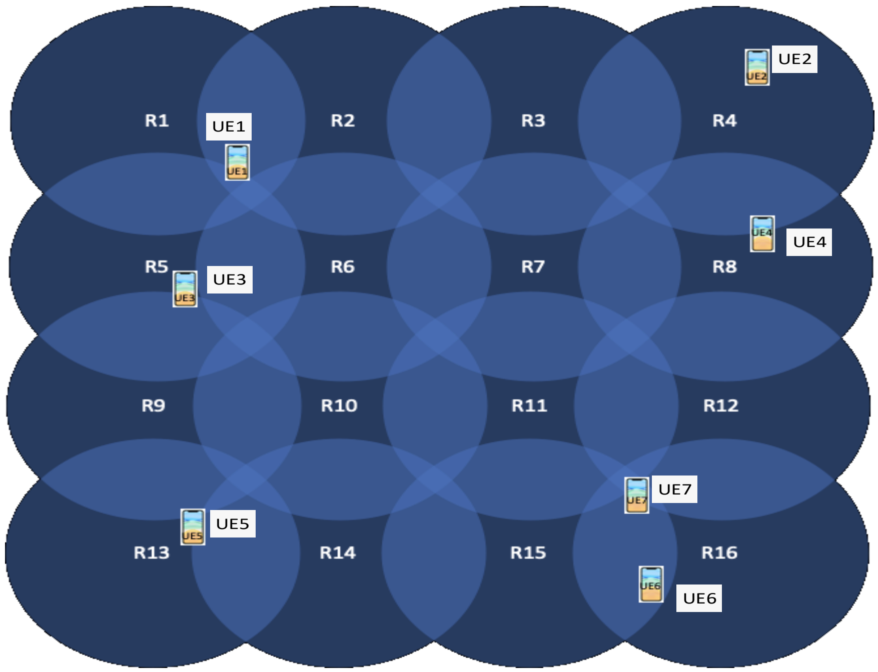

Figure 1: Shared Cell Deployment with Uniform Coverage

As depicted in Figure 1, above, the position of the RUs R1, R4, R13, and R16 are sufficiently far apart and the UEs UE1, UE2, UE5, and UE6 have sufficient isolation from each other. Accordingly, these may be scheduled at the same time on the same set of RBs at different RUs, providing four times the capacity on those RBs. In contrast, the same set of RBs at RUs R1 and R5 may not be scheduled at the same time unless those RBs are allocated to the same UE, which is in their overlapped coverage zone, thus significantly improving the SNR of the UE. Otherwise, both DL transmissions and UL transmissions on those overlapping RBs would interfere.

However, those pairs of RUs may be scheduled in one slot with different RBs as shown Figure 2, below.

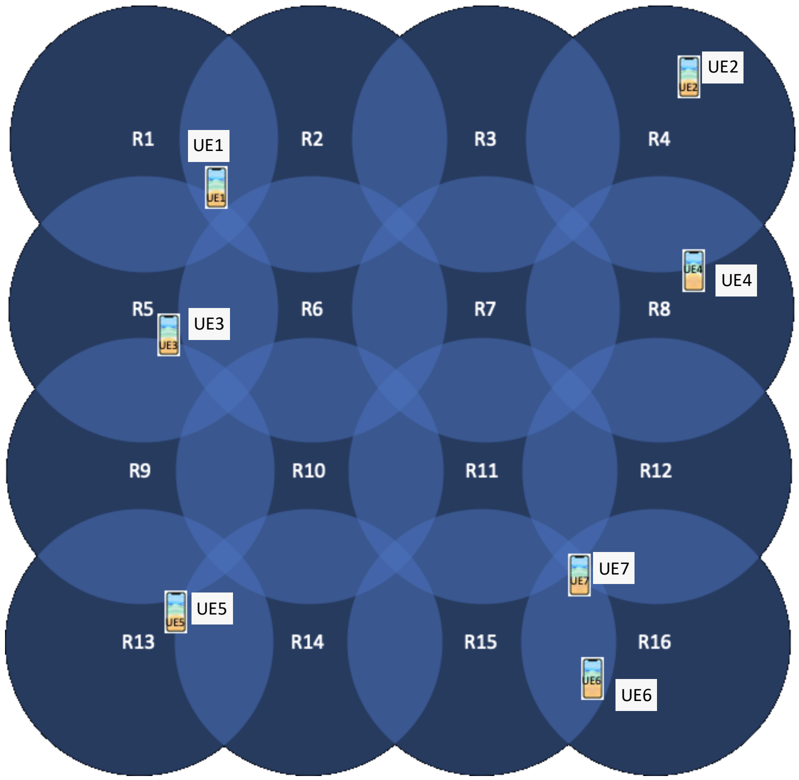

Figure 2: Overlapped RB Allocation with Shared Cell Showing Capacity Management

Similarly, as depicted in Figure 2, above, the RUs R4 and R8 and the RUs R15 and R16 may also form RU pairs for the UEs UE4, UE6, and UE7.

Under the techniques presented herein, various scheduler enhancements rely on the relative position of RUs and the relative position of a UE. In support thereof, aspects of the presented techniques support a novel way of determining those positions and proximities.

Regarding the relative position of RUs, a radio transmission on each RU, one at time, may be enabled for a short period (e.g., approximately ten milliseconds (ms)). During that time the other RUs may be placed into a network listen mode to carry out measurements. Those RUs may measure Reference Signal Received Power (RSRP) and SNR and report the same back to a DU. The DU may then compute the radio distance and the path loss from the source RU to the RU that is reporting the measurements. A path loss is the difference between a Reference Signal Transmit Power (RSTP) and the RSRP that is reported by the receiver. Such a measurement helps a DU very accurately estimate the interference from each RU to the other RUs in a network. Worst case interference will be experienced by those users that are in the middle of two RUs. If the interference is lower than some acceptable level, technically all of the RBs may be reused while scheduling transmissions over two RUs.

As is clear from the above narrative, the techniques presented herein offer a number of points of novelty and efficacy. Under one point, aspects of the presented techniques aid in deploying RUs with no, or very little, RF planning. Under another point, other aspects of the presented techniques need to be executed only at the commissioning time of a network and whenever the network is updated by adding newer RUs to the network. Consequently, ongoing measurement and computation overhead is avoided.

Regarding the relative position of a UE with respect to RUs, each UE may be in the overlapped coverage region of more than one RU. Identifying such a coverage overlap helps in estimating the interference at the UE due to transmissions from multiple RUs. Interference estimates from each RU helps in creating two sets of RU groups for a specific UE – one group with a signal strength that is higher than a required minimum (meaning the UE is closer to the RU) and the other group with a signal strength that is lower than a required minimum (meaning the RU is farther away from the UE). Technically, interference should be minimal if the same set of RBs that are allocated to the UE are reused on other UEs that are surrounded by the RUs from the minimum interference group. The first group of RUs is used to increase the SNR at the UE by transmitting the same stream of data, which helps to increase throughput (by maximizing Modulation Coding Scheme (MCS)) in both DL and UL transmissions. The second group of RUs are used to reuse the RBs for a different set of UE. Such an approach helps to increase the overall capacity of the system.

In other shared cell solutions, there is a challenge in identifying the position of a UE. Since each RU in the shared cell is transmitting the same cell parameters, the UE will not know with which specific RU it is communicating. The only way to obtain that information is based on UL transmissions at an RU assuming the channel is symmetric. However, there is an issue with such an approach. An RU would not process UL transmissions for a specific UE unless it is instructed by a Medium Access Control (MAC). However, unless the location of the UE is known, the RU cannot be instructed. In other words, it is a chicken and egg problem.

Different methods may be implemented to identify the position of a UE. RU Groups 1 and 2 (as described above) for a specific UE may be dynamically updated while the UE is moving.

First, a UE's position may be identified initially when it performs an Attach request or a Scheduling Request (SR). RUs that are involved in processing a Random Access Channel (RACH) may be tagged to that specific UE. Then, a UE's mobility may be tracked by requesting RUs that are adjacent to the initial set of RUs to also process UL transmissions for the UE. The technique for identifying the relative position of RUs that was described previously may be used to dynamically determine the set of adjacent RUs.

Second, a UE's position may be tracked by periodically scheduling a Sounding Reference Signal (SRS). All of the RUs in the deployment may be asked to process such an action. To reduce wastage or a collision of RBs, an SRS may be scheduled in a staggered manner for all the UEs.

As described previously, aspects of the techniques presented herein encompass various scheduler improvements to increase capacity. A scheduler needs to consider a UE's location as an added dimension in scheduling resources to improve the capacity of a shared cell network. For example, the following scheduling procedure may be implemented:

1. For each slot of time a scheduler may select a set of users to be scheduled based on quality of service (QoS) criteria, channel conditions, etc.
2. The relative position of a UE may be checked. RBs may be reserved for the UE on the RUs that fall into the first group that are associated with the UE. Additionally, the relative position of the RUs may also be checked and, if necessary, the resource utilization marked on the RUs that are not captured in the first group.
3. Step 2 may be repeated for each UE that is to be scheduled until either all of the users are fully or partially satisfied or there is no more room.
4. RB utilization may be identified at each RU. In brief, this step identifies the RB positions that may be reused at each RU for a new UE.

5. After identifying the resource holes at each RU, the scheduler may next select a set of users (in the order of priority) that are waiting to be scheduled based on close proximity to an RU.
6. Resources may be allocated for a specific UE utilization and marked on the RUs that fall into the first group.
7. Steps 5 and 6, above, may be repeated until either all of the resources are fully exhausted or all of the users fully satisfied.
8. A PDCCH, a physical downlink shared channel (PDSCH), and a physical uplink control channel (PUCCH) may be prepared for an individual RU and then pushed down to L1 for baseband processing.

The techniques presented herein support a number of specific advantages over conventional solutions. Among other things, aspects of the presented techniques focus primarily on increasing capacity with a multiplicity of low cost RUs that do not have beamforming capabilities.

In contrast, with an Open Radio Access Network (O-RAN) approach a number of alternatives are possible. A first alternative encompasses a shared cell with RUs supporting multiple beams. This case can address a capacity issue as long as each beam for an RU is scheduled independently. Overlapped coverage benefits are seen only if all of the RUs are tracking a UE in the overlapped regions and applying appropriate beamforming, otherwise transmissions would interfere. Among other things, this is an expensive option.

A second alternative encompasses a shared cell with RUs that are not capable of supporting multiple beams. With this case each RU may act as an independent cell and may appear as if it is addressing capacity, but the fundamental benefit of a shared cell is lost due to the fact that a UE will not see SNR gains from multiple transmission as it can only operate on one beam. Since these are not physical beams, transmissions from multiple RUs can interfere severely.

A third alternative encompasses a variation of the second alternative (as described above) where, if coverage of an RU is overlapping with another RU, both of the RUs will be configured with the same beam identifier. For example, if N (which is more than two) RUs have a coverage overlap with at least one RU, then they all need to be configured with the same beam identifier. While such an approach definitely provides SNR gains, the capacity will be 1/N.

Since the techniques presented herein focus on enterprise use cases with a shared cell architecture, those techniques are superior to the second and third alternatives that were mentioned above for RUs that do not have beamforming capabilities.

In summary, techniques have been presented that support a novel shared cell architecture that, for example, is suitable for enterprise use cases. Aspects of the presented techniques comprise self-discovery and organization, an unconstrained number of RUs, and low overhead. Further aspects of the presented techniques encompass the self-discovery of UE and RU relative positions and an exploitation of the measured RF isolation between those entities. Additionally, a smart scheduling algorithm may then be applied which multiplies cell capacity in proportion to the RF isolation.