TELECOM INFRA PROJECT

Automation and Optimisation for OpenRAN

TIP White Paper

January 2022

Table of Contents

1 Executive Summary 3
2 Motivation 5
  2.1 The business case 5
  2.2 Deployment 7
3 System Architecture 9
  3.1 Project Scope 9
  3.2 Introduction to dRAX 9
  3.3 The dRAX nRT-RIC 10
  3.4 Smart Interference Management 11
  3.5 Smart Interference Management Results 13
  3.5.1 Introduction 13
  3.5.2 Baseline results 14
  3.5.3 Heuristics 16
  3.5.4 Analytic Model 17
  3.5.5 Testbed evaluation 18
4 Conclusion 19
5 Contact Info 21
6 References 22
7 Abbreviations 23

1 Executive Summary

The RAN Intelligence & Automation (RIA) project is a subgroup of the TIP OpenRAN project which aims to help commercialise OpenRAN systems, specifically looking into the role of data science and Artificial Intelligence (AI) or Machine Learning (ML) for the purposes of RAN automation and optimisation. The activity therefore focussed on the RAN Intelligent Controller (RIC), as specified by O-RAN Alliance, and its interaction with underlying RAN components. For this project, a productized RIC was used which is part of the dRAX product from Accelleran. This report provides early results to demonstrate some immediate value that can be achieved by the use of the RIC. The Accelleran dRAX, along with the RIC, comes with its own xApps, but also allows 3rd party developers to easily create their own xApps with specific functionalities and algorithms. Therefore, the Accelleran dRAX RIC provides a platform for fast and easy xApp development and creates an ecosystem for innovation and research. In this project, two different xApps on top of the RIC have been developed, one by Accelleran which optimizes the handover process, and one by BT for better interference management. This report identifies the data that is exchanged between the RIC and underlying RAN components that enable these, and many other, features.

The interference management xApp enables coordination of cells with overlapping coverage to limit interference experienced by users in these overlapping areas. We successfully demonstrate that edge-of-cell performance can be significantly improved. This often comes at the expense of overall system capacity – a key trade-off that an operator must make when optimising such algorithms. The key benefit from interference management is more consistent user experience for users as they move around a cluster of cells.

The results from testing in the Accelleran Lab and BT's TIP community lab, as well as some supporting simulations and analytical modelling, show how these relatively straightforward xApps can provide automations that improve user experience. The expectation is that the heuristic models defined for this effort can become the foundation for the machine learning algorithms that will continually optimize system performance in a production network. This is where the TIP RIA project is focussing many of its upcoming activities.
The business justification for a RIC is also crucial to understand. For the interference management feature this includes easier deployment for faster roll-out and quicker densification, more consistent user experience, and improved spectral efficiency. We also summarise some of the key findings from the recent Analysys Mason survey which focussed on operators' expectations for the RIC. Some of the key commercial drivers are based on

cost savings, many of which come from automation and improved asset utilisation.

The deployment of RICs and OpenRAN in general do require significant considerations for an operator. In particular, the near-Real Time control in the RAN driven by ML would suggest a complete change in mindset to support continuous change. The full benefits of the RIC will likely increase gradually, due to both OpenRAN becoming more mature, and the fact that for many operators, OpenRAN will be phased into existing networks. However, the expectation is that long term, the focus will become fully user-centric and be the key way by which we can provide service differentiation.

2 Motivation

To successfully deploy, manage, and optimise a multi-vendor OpenRAN solution, a common control and management plane is required. This is particularly the case for a next-generation RAN, which will be expected to adapt to support a wide range of requirements including high performance, massive connectivity, and mission-critical communications.

This report details the activities carried out by BT and Accelleran as early contributions in the TIP OpenRAN RAN Intelligence & Automation (RIA) project. TIP RIA focuses on the availability and use of data, so that a centralized RAN controller can manage a disaggregated, virtualised and multi-vendor RAN. This central controller, the RAN Intelligent Controller (RIC), plays a crucial role in enabling OpenRAN deployments to satisfy the wide range of use cases envisaged for commercial deployments of the next decade. The RIC, as specified by the O-RAN Alliance, is disaggregated into a near-real-time RIC (nRT RIC), satisfying control loops under 1 second, and non-real-time RIC (NRT RIC), satisfying control loops of 1 second or more. This work is mainly focussed on features implemented in the nRT RIC.

2.1 The Business Case

The RIC is a great enabler for automation, the provision of a wide and diverse range of services, as well as supporting the ability to build and manage an open network that benefits from best-of-breed components from an innovative vendor ecosystem. Equally important is the RIC’s positive impact on the overall business case.

In traditional deployments, access to data is restricted. Moreover, even with a good idea and algorithm, integrating into an existing operation becomes a challenge. New interfaces have to be made, proprietary protocols implemented, etc. The Accelleran dRAX RIC, on the other hand, provides easy access to data and well-documented interfaces. In this way, developers from different backgrounds, such as operators, app vendors, academia researchers, and so on, can start developing xApps for automation and optimization. Providing an accompanying xApp framework, which abstracts the low-level technical details of hooking into the dRAX RIC, opens the way to an ecosystem of innovation. Operators using the ORAN concept with a RIC can create sandboxes for research and development based on real network data, after which such xApps can be moved into production.

This activity has focussed on interference management, which can be used for many scenarios, but are particularly crucial to the dense deployment of small cells.

Interference management is an even greater challenge for dense deployments of small cells. Huge capacity increases can be achieved through spectrum reuse, but that comes with the challenge of greater intercell interference. An effective interference management feature within the RIC can improve the business case in several ways, including:

- **Easier deployment** - Limiting the need for radio planning to minimise cost and delay in rollout. This can also ease the densification of existing cell clusters.
- **Consistent User Performance** - The key issue with intercell interference is the poor performance of edge-of-cell users. A significant difference between cell centre and cell edge performance can make it difficult to maintain QoS. This issue will only become worse as users consume higher capacity services. Interference management can improve this edge-of-cell performance. The improvement to the user is a both, an improvement in the poorest performance they will receive, but also a more consistent level of performance as they move around a cluster of cells.
- **Improved spectral efficiency** –spectrum reuse can significantly improve an operator's use of this finite resource. In many regions spectrum acquisition is one of the major costs of a mobile deployment, so enabling higher spectral efficiency, can give a much better return on investment.

The TIP RIA project has identified eleven key use-cases where they believe the ecosystem should focus in order to commercialise the RIC. Interference management is one of those key use cases, showing that operators are already highlighting that this is a real challenge, but one that has strong commercial value.

A recent report from Analysys Mason provides survey results from 35 Tier 1 and 2 MNOs [2]. The report is focussed on the expectations for a RIC to control and manage future RAN deployments. We have selected a few key findings here as further evidence supporting the RIC's positive impact on the OpenRAN business case:

- MNOs were asked to identify the key commercial driver for deploying RIC. The top answer was "Reduce TCO of the RAN" (34%),
- 40% of respondents said cost reduction was a key business benefit of deploying nRT-RIC. Other key benefits selected were based around automation (predictive RAN operations (26%) and Autonomous RAN (9%) which would lead to OPEX savings. So combining these answers, the majority of MNOs would expect the key business benefit of nRT-RIC to be associated with cost savings.
- In terms of the improvement to asset utilisation efficiency by implementing the RIC, expectations vary from less than 10% up to 50-70%, however the majority of respondents (74%) expect to see an increase in asset use efficiency by 10%-30% by deploying RIC".

2.2 Deployment

This project illustrates some of the potential benefits of intelligent automation in the RAN using the near-Real Time RAN Intelligent Controller (nRT RIC). One aspect is in developing algorithms for automation and optimization, such as the two algorithms described in this paper. A second aspect is the xApp development platform the RIC represents, and the ease of access to the data and commands it brings. There are, however, open questions to resolve before this can be widely deployed in operational networks.

Mobile Network Operators (MNOs) may need to reconsider the organisation and skillset of their operational teams in order to fully realise the benefits of OpenRAN. Traditionally MNOs are highly resistant to changes on the network and have lengthy approval processes. To successfully adopt near-Real Time control in the RAN driven by ML this mindset must evolve to embrace a new paradigm of continuous change.

Whilst the full benefits of OpenRAN and the RIC may only be delivered when this architecture is widespread across a whole network, some level of benefit can be realised on a partially OpenRAN network. A Heterogenous Network (Hetnet) topology with a traditional Macro RAN layer and an OpenRAN Small Cell layer could utilise advanced interference management techniques such as those described in this paper to manage in-band interference within and between layers.

As MNOs consider deploying OpenRAN solutions at scale many open questions remain. In terms of the architecture, we need to understand how it will effectively scale. This means understanding the optimal RIC topology and relationship between the various disaggregated RAN components. The latency between components can also have a particularly big impact on the optimal locations to deploy nRT RIC. Using the cloud-native architecture, dRAX enables flexible deployment of time-critical xApps to the network edge. When considering the relationships of the NRT RIC we are often thinking of a longer control loop, but benefiting from a more global view of the network. The NRT RIC should benefit the most from AI/ML techniques, but to what extent is still to be determined.

The expectation is that xApps deployed in the nRT RIC can come from third parties and can be deployed across many systems (xApp portability), so we need to ensure that it is possible to manage and police these potentially large collections of xApp features. These xApps may consume many compute resources. The complex interactions of different xApps may attempt conflicting updates or combine to realise unexpected outcomes. xApps may be better trained or designed to work with components from particular vendors. Overall, this means the operator needs to take a global view of the characteristics of its suite of xApps to ensure desired outcomes are achieved.

This also highlights an area where AI or ML can play an important role to enable the most effective use of OpenRAN components.

The granularity of control offered by the nRT RIC can shift the focus of Network Optimisation from minimising load on the network elements to delivering the best user experience for customers. Targeting network control algorithms in this customer-centric mode has the potential to transform MNOs' capability for service differentiation in the RAN.

3 System Architecture

The report highlights the types of data that should be made available over a common API and how it can be used by a vendor or operator to manage the RAN. We provide example data to support the wider industry development of AI techniques in the context of vRAN. While this is part of the TIP RIA project this supports all TIP OpenRAN projects and the ecosystem in general.

Activities reported in this report are based on the Accelleran dRAX which contains a nRT RIC. The dRAX nRT RIC is a pre-standard, OpenRAN aligned nRT RIC. As O-RAN alliance specifications mature, the dRAX platform will adopt these open specifications. As the focus of this activity is on the value of functionality within the RIC and the data that it has available from RAN components, the existing proprietary interfaces are sufficient.

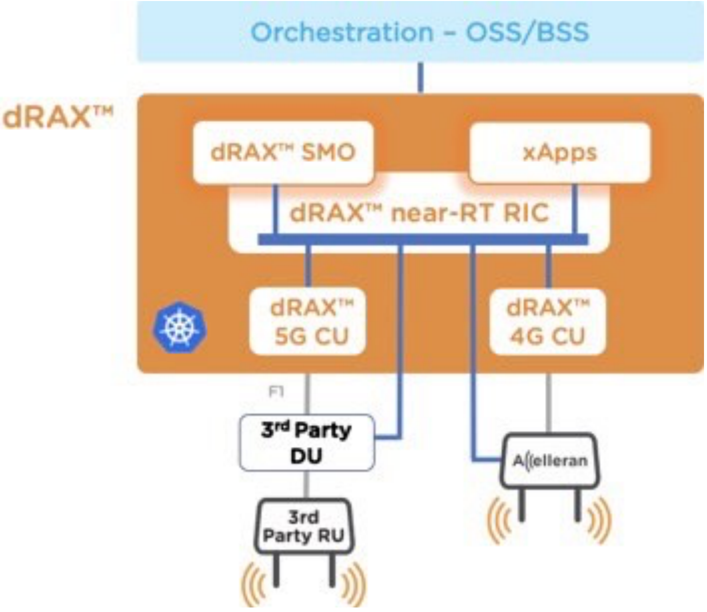

Figure 3-1 – Accelleran RIC platform.

3.1 Project Scope

This activity assessed interference management, based on features implemented in the 4G dRAX nRT-RIC.

3.2 Introduction to dRAX

Accelleran's dRAX™ delivers proven virtualized software components to enable real-world deployment of multi-vendor, disaggregated Open RAN, in alignment with the open standards such as the O-RAN Alliance. These Cloud-Native components deliver reliable, cost-effective, and scalable solutions for both 4G and 5G networks. dRAX components cover the key control and resource management functions of the RAN (Service Orchestration, RIC, CU-

CP, CU-UP) and are pre-integrated with a range of DU and RU solutions from our partners and other vendors in the RAN ecosystem. Accelleran's dRAX™-CU provides a fully standards-compliant, Cloud-Native implementation of the Central Unit – User Plane (CU-UP) and Central Unit – Control Plane (CU-CP) as defined by 3GPP. Containerized and ready for deployments on 21st-century network infrastructure, dRAX is fully integrated with popular orchestration platforms such as Kubernetes. Communication between containers is achieved through the messaging and shared data services based on cloud-native principles. One of the key benefits of the architecture is location transparency for all microservices, which brings enormous flexibility in the operational network for placement of functionality and the design of a solution for scaling and resilience.

In this whitepaper, we will be using the Accelleran dRAX with 4G. The following sections indicate the RIC details, 4G data available, commands and APIs which are used in the 4G context. It should be noted that the same functionalities are available for 5G dRAX as well. 4G dRAX supports the OpenRAN split of the Central Unit (CU). The fully 3GPP standards-compliant 4G Central Unit – Control Plane (CU-CP) is centralized and built using a microservice oriented, Cloud-Native architecture. The Central Unit – User Plane (CU-UP) resides on the eNB. A messaging databus is used for communication between the RAN components, as well as communication towards the nRT-RIC.

3.3 The dRAX nRT-RIC

Accelleran dRAX™-RIC delivers a truly Cloud-Native, pre-standard, OpenRAN aligned near real-time RAN Intelligent Controller (nRT-RIC) that enables near real-time control and optimization of Open RAN elements and resources under its control. At the basis of the dRAX™-RIC lies a framework that provides all the necessary functions for onboarding and life cycle management of xApps. The dRAX RIC is a productized platform where any 3rd party developers can develop their own xApps. It is not a proprietary platform, but using well documented APIs and interfaces, it gives access to RAN data and provides a friendly experience to xApp developers.
In summary, the Accelleran dRAX RIC provides a number of services:

- xApps onboarding and lifecycle management
- access to real-time RAN measurements and events
- configuration of RAN components
- real-time commands to direct RAN behaviour (e.g. force a handover, sub-band masking)
- real-time state database
- Inter xApps communication
- API-driven xApps configuration and policy management

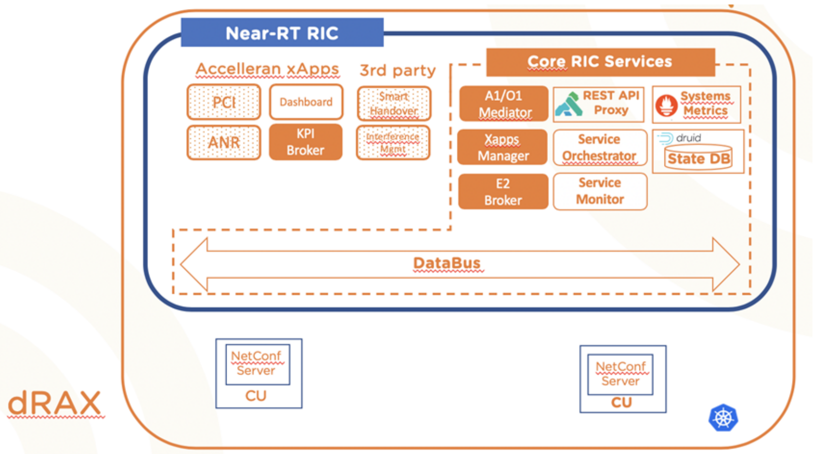

Figure 3-2. dRAX nRT-RIC

This RIC platform can be leveraged in multiple ways and for multiple market segments:

- Providing plug-and-play capabilities when small cells are added to a cluster, reducing the time and complexity of cellular installations. This is not only critical in a service-provider environment but also in Private Cellular deployments where the 4G/5G skill levels of Enterprise IT personnel are limited.
- RAN Automation.
- Implementation of AI/ML-based algorithms to augment the intelligence of the RAN. The developer can focus on the algorithm rather than the complexity of the integration and access to the data.
- Leverage an ecosystem of complementary xApps communicating with each other.
- Can run on both Intel and ARM architectures.
- Is a distributed platform, which can run on premise, at the edge or in the cloud. Using a messaging framework as the dRAX Databus, the platform can combine different levels of distribution, for both near-real-time and non-real-time use cases

The RIC is a key investment area for Accelleran as it sits at the intersection of RAN expertise, cloudification and AI/ML. The ability to offer a productized, independent, and open development platform to increase the RAN intelligence and programmability is critical for the open and disaggregated RAN ecosystem.

3.4 Smart Interference Management

All cellular radio systems have the potential to benefit from various methods of interference management, which are intended in particular to ensure improved edge-of-cell performance, for connected devices situated in areas suffering from the highest levels of interference ICIC (inter-cell interference coordination) is a common form of interference management feature, and it is this feature that we have investigated for this activity.

ICIC, as defined by 3GPP, has been available for LTE since release 8. It allows for neighbouring cells to coordinate resources in the frequency domain in order to reduce intercell interference. The expected behaviour is that a cell which is causing interference on a particular frequency block will aim to reduce its transmission power. This reduced interference will allow a serving cell to then provide a better service to users who now experience a better signal-to-noise ratio (SNR). ICIC is therefore mainly aimed at improving service for cell edge users as they tend to have the weakest serving signal and the strongest interfering signals.

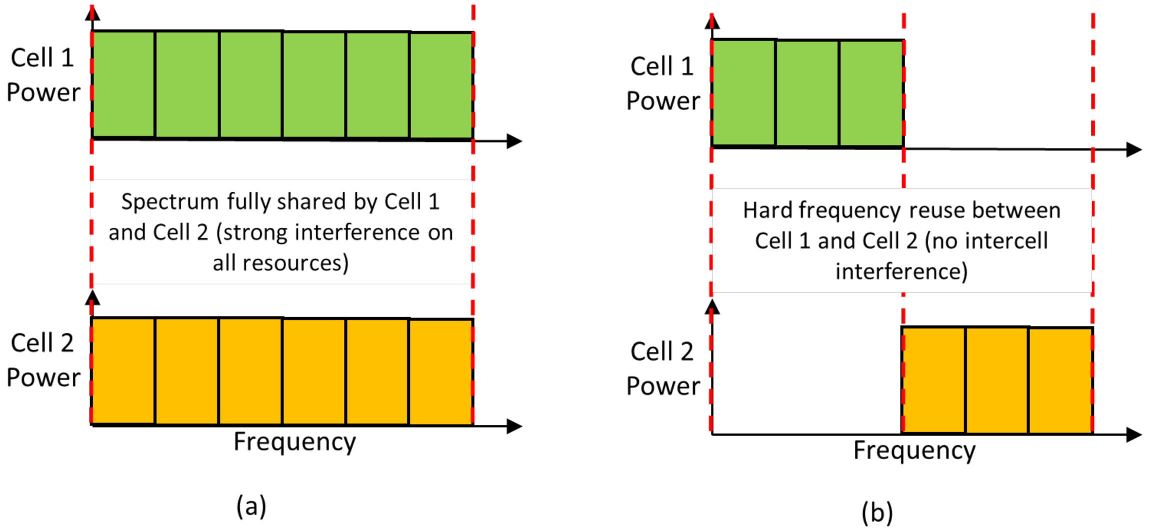

Figure 3-3 ICIC example. (a) full spectrum reuse (no ICIC), (b) ICIC implemented with hard frequency reuse

Figure 3-3 provides a simple example for two cells with overlapping coverage of how ICIC can be implemented and the trade-off it provides. Part (a) of the figure shows the situation with no ICIC implemented. Both cells can occupy all frequency blocks at full power. This enables maximum spectrum reuse, which can enable the highest network capacity. However, this means that cell edge users will have to suffer a poor SNR. Part (b) of the figure shows the situation where the two cells have coordinated so that they each take half of the frequency blocks for themselves, while the other cell does not transmit in those frequency blocks. This means that cell edge users have a much-improved SNR and can experience much better performance. However, each cell now only has access to half as many frequency blocks as in part (a), thus network capacity has dropped. This highlights the main trade-off of ICIC. We can make user experience more consistent (by treating cell edge users better), but overall capacity may reduce. There are various ICIC options that can also allow you to adjust the trade-off. The frequency blocks can be reconfigured dynamically based on demand, soft frequency reuse can be used where interfering cells simply reduce power in a frequency block rather than

transmit nothing, some resources can remain fully reused. There are a huge number of variations which becomes much more complicated when we are dealing with a larger cluster of interfering cells.
This work has focused on the use of the RIC for interference management. The nRT RIC is in an ideal position to make the best decisions for ICIC, as it can monitor the situation for an entire cluster of interfering cells and make the best overall decision on how to assign frequency blocks to each cell in that cluster. Secondly, in order to use AI techniques, we need the more centralized access to system performance data which the RIC provides. Thus, in the present work we are tackling the challenges of making best use of the new RIC open architecture. In particular, we want to perform smart interference management by the use of sub-bands, with hard frequency reuse, and dynamically re-allocate these sub-bands between neighbouring cells in response to user behaviour. In the system under test we get radio measurements and make decisions on a per sub-band basis. This system has 13 sub-bands in total.

We investigate several ICIC algorithms for the RIC, but we will focus here on the main algorithm used in the testbed at the TIP Community Lab at Adastral Park. Each loop of the algorithm tries focuses on improving the situation for the worst served user in the cluster of cells. It will then negotiate assigning an extra sub-band to it and/or turning off a sub-band of an interfering cell. This algorithm is implemented as an xApp on Accelleran's dRAX/RIC platform.

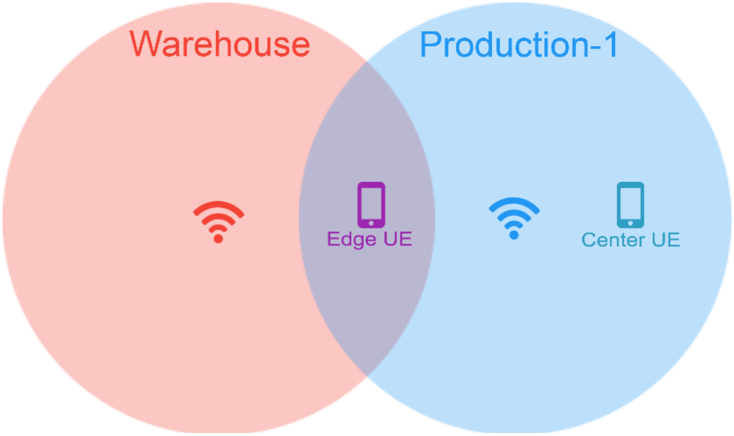

Figure 3-4 Interference Management scenario setup

3.5 Smart Interference Management Results

3.5.1 Introduction

The experiments for the Smart Interference Management xApp are done in the TIP Community lab at Adastral Park. We conduct multiple experiments:

1. Baseline experiments – Here we manually apply sub-band masking to determine its effect on the performance of the network under interference.
2. Analytical model – Here we simply highlight the trade-off between consistent user experience and network capacity.
3. Smart Interference Management xApp results – the xApp implements ICIC on the testbed.
4. Simulator experiments – Two different simulations, each using their own ICIC algorithms, are provided to demonstrate the effectiveness of the RIC to coordinate ICIC in larger clusters.

3.5.2 Baseline results

To determine how sub-band masking works and how it influences the performance of the network under interference conditions, we conduct the initial baseline experiments. In these experiments, we play around with the programable attenuator settings to create interference. As shown on Figure 3-4 **Error! Reference source not found.**, we position one UE at the edge of the two cells Warehouse and Production-1, while the second UE is in the center of coverage of cell Production-1. We then manually apply sub-band masking to determine its effects in such a scenario. In this example we start with all radio resources being made available to each cell. Then we simultaneously mask the first half of the sub-bands on cell Warehouse and the second half of the sub-bands on cell Production-1. This way we give each cell half of their resources, but the sub-bands not experiencing inter-cell interference.

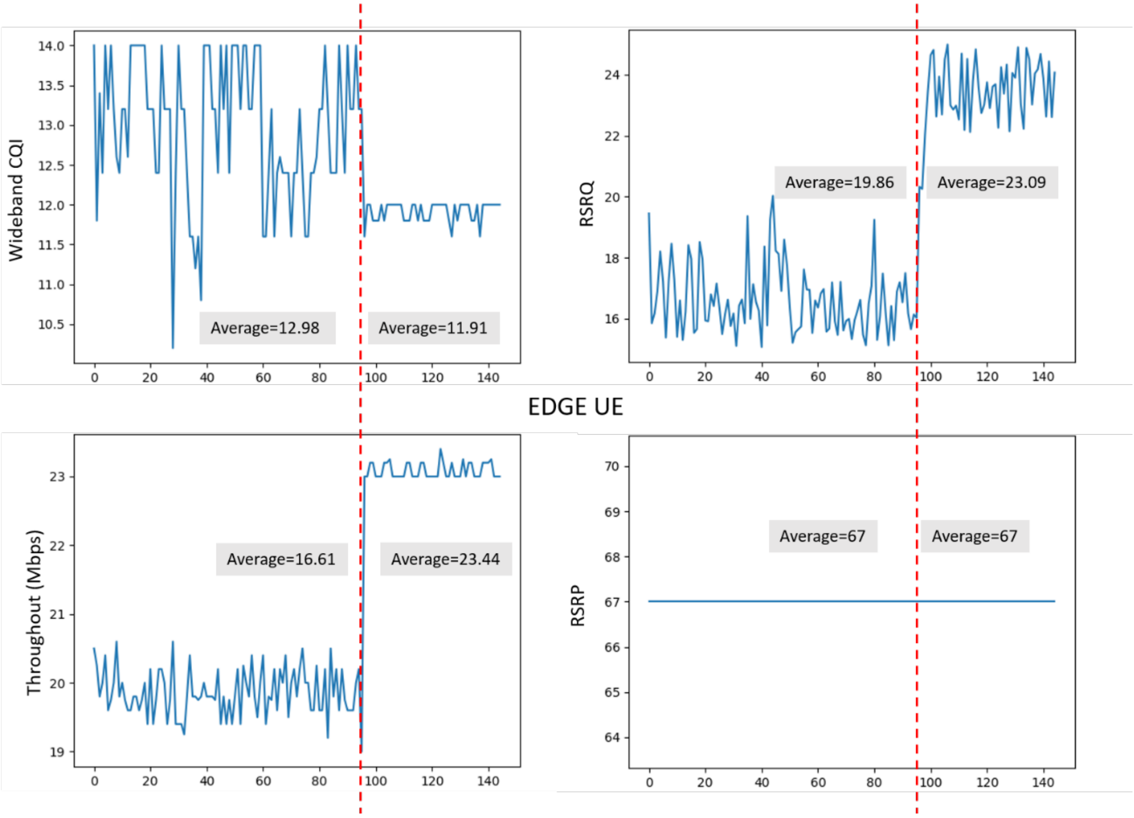

Figure 3-5. Edge UE results showing the throughput, RSRP, RSRQ and wideband CQI time series

Figure 3-5 show the results of the experiment for the Edge UE shown on Figure 3-4. It shows the time series of the wideband CQI, RSRQ, throughput and RSRP. The red dotted line indicates when the sub-band masking was applied. Before the red-dotted line, there was no sub-band masking, while after it the sub-band masking was applied. We also show the average values of each metric before and after the sub-band masking is applied.
From the wideband CQI graph in Figure 3-5 we can see that while all sub-bands are in use on both cells, the wideband CQI value is very noisy. It drops as low as 10, and also reaches the highest value of 15 at certain points in time. However, once the sub-band masking is applied, the wideband CQI for the Edge UE stabilizes to around 12.

Looking at the RSRP, we can see that before and after sub-band masking is applied, it remains mostly unaffected. This is expected, as RSRP describes the reference signal power of an individual cell. Since we are not adding or reducing attenuation dynamically in this experiment, it is expected that the received power of reference signals remains the same.

The RSRQ, on the other hand, tells a different story. Since the RSRQ is the signal quality indicator, it should be affected by interference. We can see that for the Edge UE the RSRQ improves by around 16% after sub-band masking is applied.

Finally, we analyze the throughput graphs of Figure 3-5. For the Edge UE we clearly see an improvement. Due to the interference from the other cell, and attenuation towards its serving cell it was able to achieve an average of 16.61 (Mbps) before sub-band masking was applied. After we apply sub-band masking, its average throughput is 23.44 (Mbps). So, the throughput of the Edge UE is improved by around 41%.

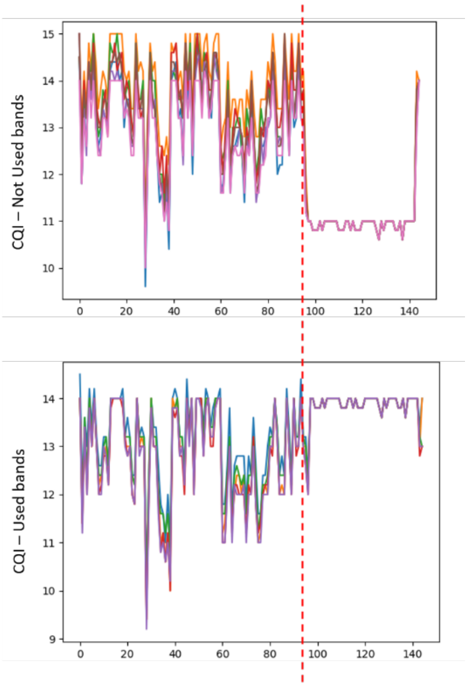

Figure 3-6. CQI results per sub-band for the Edge UE (masked vs un-masked sub-bands)

Figure 3-6 shows the CQI per sub-band for the Edge UE. The upper graph shows the CQI of the sub-bands that are not used by the Edge UE, so the sub-bands that are masked on their serving cell, while the lower graph shows the CQI of the sub-bands that the Edge UE is using. The red dotted line again indicates when the sub-band masking was applied

We can clearly see that the CQIs are noisy before sub-band masking was applied. After the sub-band masking is applied, we can see that the CQIs stabilize. For the Edge UE, the CQI for the sub-bands that it uses is around 14, while the for the sub-bands it does not use the CQI drops to around 11. This can be explained by the fact that the sub-band the Edge UE uses are not used by the other cell, cell Production-1, which means the interference is reduced. These baseline results show that in an interference environment, applying sub-band masking can help minimize its effects. By using one set of sub-bands for one cell and another for the second cell, we were able to improve the throughput of the Edge UE experiencing strong intercell interference by 41%. We also see an improvement of the RSRQ which indicates that the interference was reduced, which resulted in the throughput improvement.

3.5.3 Heuristics

In the previous section we triggered subband masking command in a rather simplistic way, in order to explore whether subband masking helps at all with interference management. We now go on use that knowledge to create a heuristic for a more rigorous approach to the problem.

3.5.4 Analytic Model

Before designing the heuristic, we should understand in a general way how a channel may be shared. The following calculation illustrates the problem. We will compute the exact throughput as a function of the band splitting ratio. The possibilities are illustrated in Figure 3-7, which explores different ways in which a single channel (assumed to have a total bandwidth normalized to 1) is shared between two cells. We allow a three-way choice for each sub-band: either the sub-band is given only to cell 1 (x axis), is shared (y axis), or is only given to cell 2 (1-x-y). It can be seen that the optimal point for the objective of maximizing total throughput is (x,y)=(1,0), which shares the whole band, effectively not using ICIC. However, if we add the constraint that the throughput of the two cells should be equal (thus balancing the loads), the new optimum point is at about (x,y)=(0,0.6), which means that 60% of the band is shared, 40% is given only to cell 2 (with the lower SINR), and cell 1 has no dedicated sub-bands. This is a simplified and theoretical example; it is designed only to make the point that careful choices of objective and constraints will be needed in any heuristics for this type of problem.

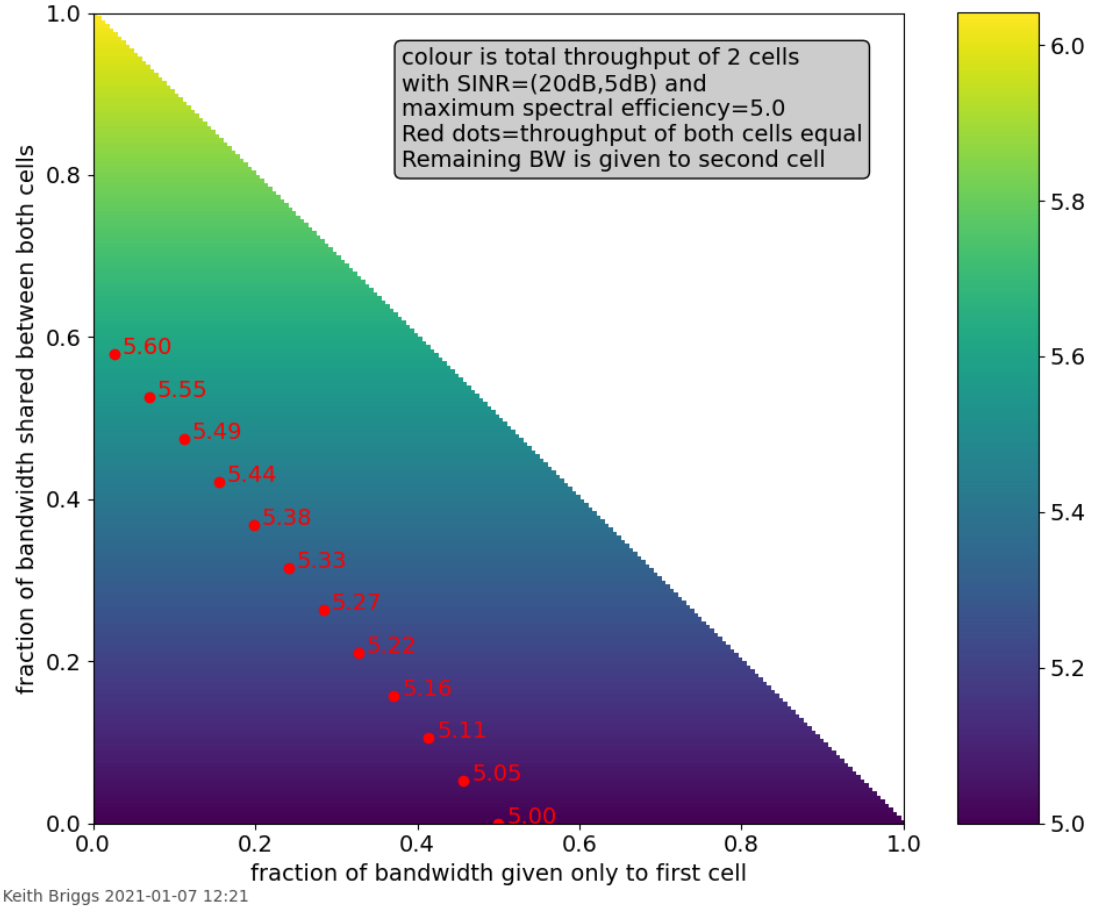

Figure 3-7 Example of how bandwidth sharing can impact per-cell and overall throughput. (two cell simplified analytical example with one edge UE and one cell centre UE)

3.5.5 Testbed evaluation

With these points in mind, we can present pseudo-code of a typical heuristic, designed to bring all throughput reports to a common target by allocating or de-allocating sub-bands. One step of this heuristic examines the current average throughput across all sub-bands of the worst-served UE in the network and acquires a new sub-band by de-allocating a sub-band from the interfering cell. Throughput values are computed with an exponentially-weighted sliding window average, to help smooth out oscillations.

This heuristic used in the testbed was implemented by creating an xApp on top of dRAX nRT-RIC in BT's TIP Community Lab testbed, and found to perform very well. See Figure 3-8 below. Until time 200, both cells share a single sub-band (hence the poor throughput). At time 200, the heuristic is switched on. The dramatic rise in downlink throughput (a factor of about 10, as reported by iperf) of both UEs shows that the heuristic is working as intended.

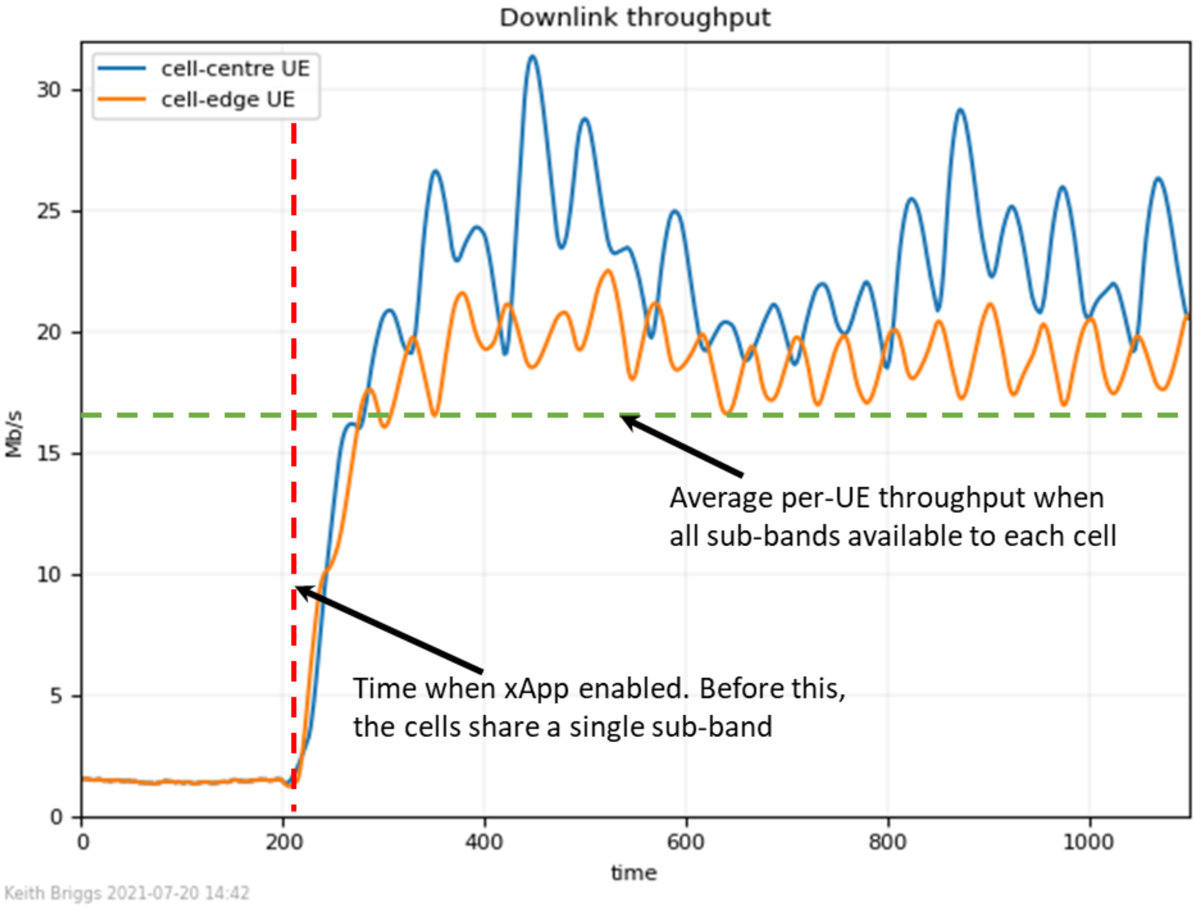

Figure 3-8 Throughput performance of the xApp with two UEs. The time axis is in seconds. The xApp was turned on a time 200.

4 Conclusion

The RIC has a vital role to play in the deployment of OpenRAN solutions. This report provides some early insights into the data exchange between the RIC and its underlying base station elements, and how features within the RIC can provide improved service.

The Accelleran dRAX RIC is a framework that provides easy access to RAN data, as well as APIs to interact with the network elements and change their configuration or apply commands such as handover or sub-band masking. On top of the RIC, xApps can be created by any 3rd party developer to exploit those features and bring additional functionalities. This creates an ecosystem of research an innovation where xApp developers can create automation and optimization algorithms.

The focus area of this paper is interference management. Here the RIC is used to dynamically adapt the sub-band allocations of interfering cells to improve the performance for edge-of-cell users. We have explored a variety of techniques to understand and characterise the problem. First an analytical set of results has demonstrated how there is a balance to be made between the consistency of user experience and the overall capacity of the network. On our testbed in the TIP Community Lab at Adastral Park, Ipswich, sponsored by BT, we have performed manual testing to see the impact of different sub-band configurations. Combining the experience from the analytical results and the manual testing we developed and tested an interference management xApp, which is shown to be effective. To ensure these types of solutions can scale to larger clusters of cells, we have also provided some simulation environments, which show that there are simple heuristic algorithms can deal with larger clusters and can achieve close to optimal results. Should such heuristics be deployed in a nRT RIC, as shown in this work, there is scope for AI to control and tune such algorithms to ensure performance meets the ever-changing needs of the end user in the ever-developing networks.

Within the TIP OpenRAN project, the RAN Intelligence & Automation (RIA) Project looks to help the industry collaborate on the enablement and commercialisation of OpenRAN systems that benefit from the strength of data science and AI/ML. The work from this report acts as an early output from the RIA project to indicate some of the value that can realised from the data available to the RIC and indicate areas where AI/ML has potential to offer further value. The RIA project is focussed on several use cases, some focussed on the NRT RIC and others on nRT RIC where it is believed the industry should focus; Intercell interference coordination, as covered by our interference management activities, is one of those key use cases for nRT RIC.

For the RIC to experience wide deployment, we need a solid business justification and barriers to deployment must be addressed. This report wraps up by highlighting some of the key business case benefits of a RIC where TCO savings and improved network utilisation are key. We also address the remaining deployment challenges, such as managing multi-vendor solutions, competing optimisation features, and scalability. Managing and optimising the RIC in an OpenRAN deployment provides a great opportunity for AI/ML, and this is a focus area for the TIP RIA project's future activities.

5 Contact Info

This work was carried out in collaboration between Accelleran and BT, with the key contacts being Ensar Zeljkovic (ensar.zeljkovic@accelleran.com) and Keith Briggs (keith.briggs@bt.com) respectively.

Thanks also to other contributors. This includes Dimitris Gogos, Gabriele Piscitelli and Arnaud Polster from Accelleran, and Arjun Parekh, Rowan Limb, Arman Shojaeifard, Shipra Kapoor and Richard Mackenzie from BT.

6 References

[1] J. Ḿ rquez-Barja, C. T. Calafate, J.-C. Cano, and P. Manzoni, “Performance trade-offs of a IEEE 802.21-based vertical handover decision algorithm under different network conditions,” in 2011 IEEE 10th International Symposium on Network Computing and Applications. IEEE, aug 2011, doi: http://dx.doi.org/10.1109/nca.2011.50

[2] Analysys Mason, “Near-real-time RIC: enabling AI/ML-driven extreme automation and granular control of Open RAN”, June 2021 [TIP Announcement]()

7 Abbreviations

4G	4th generation
5G	5th generation
AI	Artificial Intelligence
API	Application Programming Interface
ARM	Company name (Advanced RISC Machines)
BLER	Block Error Rate
BT	Company name (British Telecommunications PLC)
CQI	Channel Quality Indicator
CU	Centralised Unit
DU	Distributed Unit
EPC	Enhanced Packet Core
EWM	Exponential Weighted Mean
ICIC	inter-cell interference coordination
LTE	Long Term Evolution
ML	Machine Learning
nRT	near-real time (RIC)
NRT	non-real time (RIC)
OPEX	Operating expenditure
O-RAN	Industry alliance name (O-RAN ALLIANCE)
RAN	Radio Access Network
REST	Representational State Transfer
RIA	RAN Intelligence & Automation
RIC	RAN Intelligent Controller
RSRP	Reference Signal Received Power
RSRQ	Reference Signal Received Quality
RU	Radio Unit
SNR	Signal to noise ratio
SINR	Signal to Interference and Noise ratio
TCO	Total Cost of Ownership
TDD	Time Division Duplexing
TIP	Telecom Infra Project
TTT	Time To trigger
UE	User Equipment
UP	User Plane
URL	Uniform Resource Locator

Copyright 2022 © Telecom Infra Project, Inc. All rights Reserved.

TIP Document License

By using and/or copying this document, or the TIP document from which this statement is linked, you (the licensee) agree that you have read, understood, and will comply with the following terms and conditions:

Permission to copy, display and distribute the contents of this document, or the TIP document from which this statement is linked, in any medium for any purpose and without fee or royalty is hereby granted under the copyrights of TIP and its Contributors, provided that you include the following on ALL copies of the document, or portions thereof, that you use:

1. A link or URL to the original TIP document.
2. The pre-existing copyright notice of the original author, or if it doesn't exist, a notice (hypertext is preferred, but a textual representation is permitted) of the form: "Copyright 2019, TIP and its Contributors. All rights Reserved"
3. When space permits, inclusion of the full text of this License should be provided. We request that authorship attribution be provided in any software, documents, or other items or products that you create pursuant to the implementation of the contents of this document, or any portion thereof.

No right to create modifications or derivatives of TIP documents is granted pursuant to this License. except as follows: To facilitate implementation of software or specifications that may be the subject of this document, anyone may prepare and distribute derivative works and portions of this document in such implementations, in supporting materials accompanying the implementations, PROVIDED that all such materials include the copyright notice above and this License. HOWEVER, the publication of derivative works of this document for any other purpose is expressly prohibited.
For the avoidance of doubt, Software and Specifications, as those terms are defined in TIP's Organizational Documents (which may be accessed at https://telecominfraproject.com/organizational-documents/), and components thereof incorporated into the Document are licensed in accordance with the applicable Organizational Document(s).

Disclaimers
THIS DOCUMENT IS PROVIDED "AS IS," AND TIP MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE DOCUMENT ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.
TIP WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE DOCUMENT OR THE PERFORMANCE OR IMPLEMENTATION OF THE CONTENTS THEREOF.

The name or trademarks of TIP may NOT be used in advertising or publicity pertaining to this document or its contents without specific, written prior permission. Title to copyright in this document will at all times remain with TIP and its Contributors.

This TIP Document License is based, with permission from the W3C, on the W3C Document License which may be found at https://www.w3.org/Consortium/Legal/2015/doc-license.html.