NEXT GENERATION INTERNET

Open Call 5

Open and Resilient Radio Access Network for Next Generation Wireless Backhauls (OR2AN2G)

Deliverable 3: Experiment Results and Final Report

Authors	**Leonardo Maccari, Gabriele Gemmi, Andrea Marin** leonardo.maccari@unive.it gabriele.gemmi@unive.it marin@unive.it Ca' Foscari University of Venice
Due Date	7/12/2022 (note: due at end of the project)
Submission Date	7/12/2022
Keywords	5G, IAB, topology, Colosseum

Deliverable 3: Part I

Analysis, results, and wider impact

1 Abstract

Possibly, the greatest innovations introduced by 5G will be invisible to the users and reside in the open, dynamic, intelligent architecture that runs the network. However, we can only imagine a tiny bit of the complexity of future networks, not enough to base current research on realistic assumptions. This project aims at merging the data-based network design carried on by the University of Venice (Italy) with the advanced experimental facilities by Northeastern University (USA) to study the backhaul network of 5G and beyond, which will radically change compared to previous generations of mobile networks. We will produce scientific results, open source code, and open data that will contribute to shape this research field in the years to come.

2 Project Vision

5G was designed to provide an impressive performance step-up compared to its predecessors, but its success depends on a strong increase on the number of base stations, named Next Generation NodeB (gNB), which may pass from 8-10 per km^2, to tens and maybe hundreds of per km^2. gNBs need to be connected to the core network, thus increasing the costs of the infrastructure and making the network affordable only in densely inhabited areas.

To solve this challenge 5G introduced the concept of Integrated Access and Backhaul (IAB) in which "IAB-nodes" collect the traffic from user terminals and have no wired connection, and "IAB-donors" are fiber-connected to the network core. A wireless backhaul network between gNBs needs to be created to route the user traffic from IAB-nodes to the closest IAB-donor.

Replacing the cable connections with a wireless multi-hop network is a key opportunity to make 5G services profitable even in situations where a capillary cabled infrastructure would be unaffordable, and in general to extend the reach of mobile networks. A new challenge for research is to provide performance, reliability, and dependability to those we call Next Generation Wireless Backhaul (NGWB), leveraging on the research background on wireless mesh networks.

This project aims to evaluate the technical feasibility of large-scale NGWB using real world data and taking advantage of the "Colosseum" testbed available at Northeastern University. This testbed, based on 256 software-defined radios, allows the deployment and testing of solutions at scale under different channel conditions, which are emulated through FPGA-based filters.

Our goal is to design networks based on realistic scenarios and assess to what extent NGWB can be used to make 5G more accessible by the end users.

Co-funded by the Horizon 2020
Framework Programme of the European Union

3 Details on participants (both EU and US)

Ca' Foscari University of Venice (UniVe) - Department of Environmental Sciences, Informatics and Statistics (DAIS)

- **Prof. Leonardo Maccari** is the Italian PI of the project. He has been an associate professor at DAIS since 2019. He has more than 60 publications in the networking field, among which several works on mesh networks
- **Prof. Andrea Marin** is an associate professor at DAIS, with more than 110 publications in the networking and performance evaluation field.
- **Gabriele Gemmi** is a third-year Ph.D. Student at DAIS. His doctoral thesis is focused on Wireless Backhaul Networks both for mobile access and Fixed Wireless Access.

Northeastern University (NEU) - Institute for the Wireless Internet of Things (WIoT)

- **Prof. Tommaso Melodia** is the William Lincoln Smith Chair Professor with the Department of Electrical and Computer Engineering at Northeastern. He is also the Founding Director of WIoT and the Director of Research for the PAWR Project Office. He has extensive experience in self-organizing mesh networks.
- **Dr. Michele Polese** is a Principal Research Scientist at WIoT, with several publications on O-RAN and IAB.

Co-funded by the Horizon 2020
Framework Programme of the European Union

4 Results

In this section we briefly review the KPI introduced in the project proposal, and then list the results we obtained during the project. In the next section we will detail all the work done, the scientific results, the open data and the source code published.

KPI: Project description

The KPI that were introduced in the project submission were the following ones (with reference to their milestone):

- M1: *Source code of the IAB solution implemented, to be published on github.*
- M2: *Implementation of a NGWB controller compatible with the O-RAN interface, with source code on github*
- M3: *Initial results on the characterization of the backhaul performance. Open data to be published on the Zenodo platform for at least 2 realistic backhaul scenarios.*
- M4: *A technical report as the base of a scientific publication to be submitted in a Q1 journal or top conference*

The main results of the project can be summarised as follows, for the details of the availability of the source code and open data, please see the relevant subsections and the tables in sect. 11:

- The extension and testing of an IAB-like interface to realise a multi-hop path in a IAB backhaul for 5G networks, with open source code published on github (M1).
- The realisation of scenarios now included in the Colosseum emulator based on open data, for the 4 cities of Florence, Milan, Barcelona, Luxemburg. The generic open data are now published on an internal repository, but will be published on the Zenodo platform as soon as they are completed. The specific elaborated data to be used as input to Colosseum are now part of the Colosseum environment at NEU. The source code for the scenario generation is published on github (M4, M3)
- The publication on the Colosseum testbed of three ready-to-use system images containing a linux container with the IAB extension of OpenAirInterface, the IAB extension of the core and the non-real time RIC.
- The algorithms for the selection of donor nodes and active links in the backhaul graph obtained from the previously introduced scenarios. The algorithms produce the topologies used by the controller and are explained in one draft paper attached to the report (M2,M4). Source code is on github.
- The realisation of all the code needed to import the physical scenario, the radio scenario, the link channel allocation in the Colosseum emulator (M3). The source code for the extended iab-manager, the toolkit needed to set-up the experiments is also published on github.

Co-funded by the Horizon 2020
Framework Programme of the European Union

- The initial testing of the performance of OAI in the selected scenarios, results are summarised in the deliverable and will be further expanded in the papers (M3).
- The realisation of a controller for the Colosseum emulator that is able to deploy the topology. The controller uses iab-manager as a library and uses standard, reusable interfaces. Source code is available on github (M2).
- The implementation of the O1 interface termination in OAI, used in the IAB node. The interface is part of the controller and it is the subject of one of the submitted publications (M2, M4).
- One accepted short paper to an IEEE/IFIP conference and the preparation of 2 more publications based on the points outlined so far, that will be completed and submitted in Dec 2022 and by mid-2023. Since the publications are already in advanced state, we attach the draft for review instead of a single technical reports (M4).

Of the three publications, one has been already submitted and it was recently accepted for publication, it is present in the bibliography as [Gemmi2023], while the second and the third are to be submitted to two conferences in the first half of 2023. Since the conferences use a double blind submission process, we can not attach the paper drafts to the deliverable or provide a title, or else, the text would be searchable on-line and conference reviewers would find it and invalidate the submission. As such, we leave them out of the deliverable but we attach it to the deliverable email for the internal project review. We refer to them as SUB1 and SUB2 in the text hereafter. The former focuses on the extension of O-RAN and the implementation of the controller, and it will be authored by personnel from UniVE, NEU, and even other universities that participated to the development of the code and realization of the tests. The latter one focuses on the optimization of the backhaul network. In the deliverable we summarise the technical findings, and we refer to the publication drafts for all the details. Gabriele Gemmi will remain at NEU for an extended period, up to March 2023, financed by both institutions with external funds, in order to complete the research work and the publications.

Before we describe in the detail the result, let us simply recall the workflow that has been at the base of our research:

1. Collect and refine open data available to study a real world urban scenario. This produces a usable 3D scenario as explained further on in the deliverable.
2. Use techniques from our prior art research to define the positions of gNBs in the 3D scenario in order to achieve line-of-sight (LoS) coverage of ground points. We also improved the prior art with a new demand model we use in one scenario.
3. Develop and use optimization algorithms to define the position of the donors and the backhaul links that are going to be activated to create the multi-hop paths to the core network.
4. Extend and test an IAB-like solution that can be the building block to create a multi-hop backhaul network using available open source code.

Co-funded by the Horizon 2020
Framework Programme of the European Union

5. Convert the data in the input format of Colosseum.
6. Implement all the required code to run the network on Colosseum (iab-manager extension).
7. Make experiments to functionally validate all the components, so that we verify the backhaul is functional, the IAB-like solution supports complex networks deployments.
8. Make experiments to test the performance of the network on Colosseum. Results from this point also feedback to the design of the optimization algorithms in 3.
9. Produce open data from every step.
10. Finally, once the topologies are implemented, we developed and tested an rApp (a software application using the O-RAN architecture) that in a real working condition can set-up the network, monitor the network state and is able to implement a new topology if something changes. The rApp uses iab-manager as a library.

The rest of the document is organised as follows, with reference to the points described before:

- We first describe the whole process of open data creation, refinement and release, which implements point 1, 5, 9.
- The description of a new demand model we used for point 2
- The description of the optimization algorithms we proposed and their performance evaluation (point 3)
- The description of the IAB-like solution and the tests we performed (point 4, 7, 8)
- The description of the rApp and its use (point 10).

For each step we outline the advancement from the state of the art as the novel contribution of the project.

Co-funded by the Horizon 2020
Framework Programme of the European Union

4.1 Discussion and Analysis on Results

Open Data Creation and Release

The creation and publication of open data is particularly complex and follows the block diagram in Fig. 1. We will describe each block, some of them (the output block labelled with a number) produce data that can be used by other researchers, with an increasing level of elaboration. Some blocks (labelled with a letter in the figure) produce open source code.

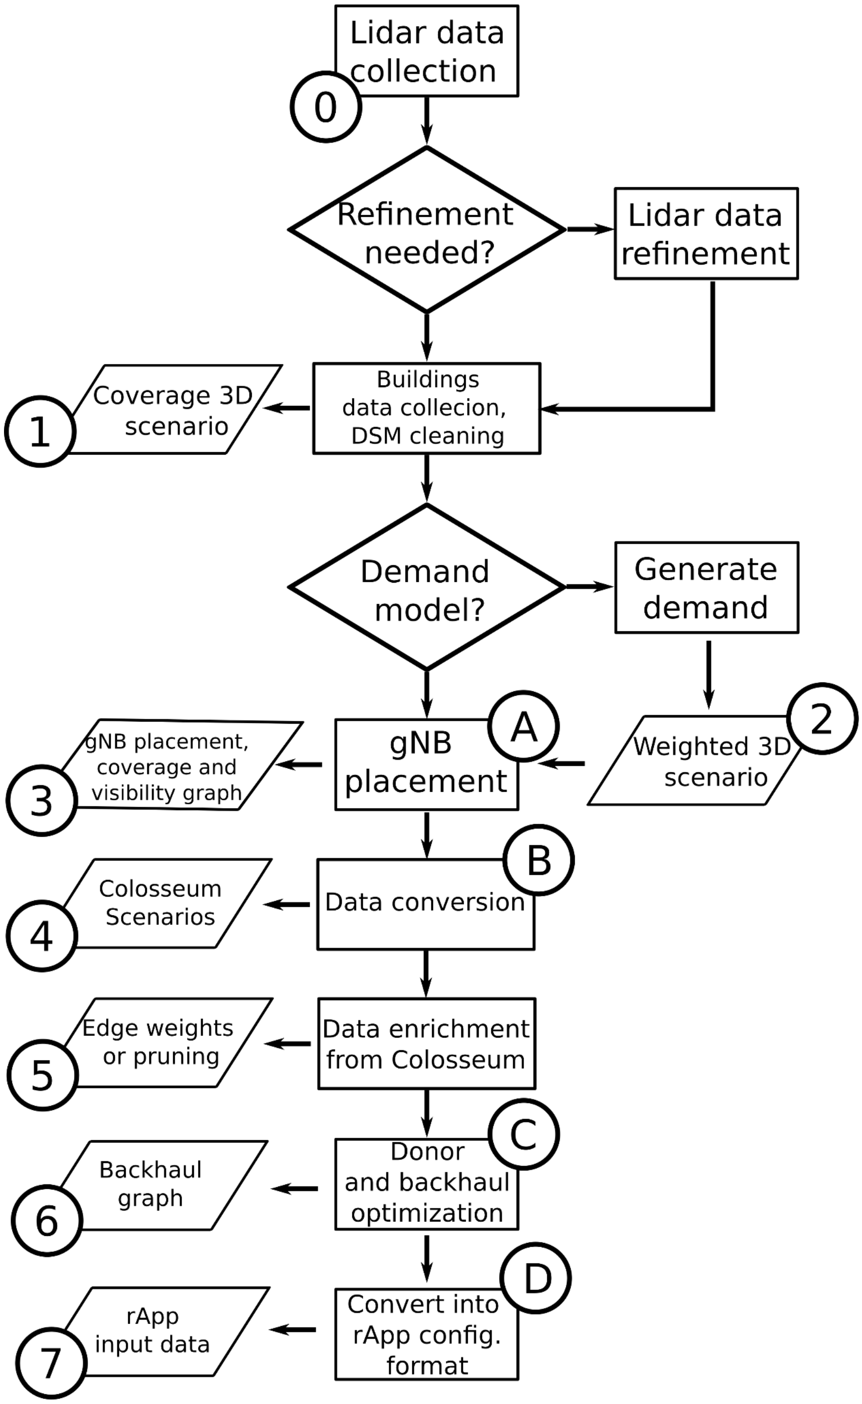

Fig. 1: The Open Data creation process.

Co-funded by the Horizon 2020 Framework Programme of the European Union

As a first step (*Lidar Data Collection* block) we had to obtain open data from public administrations about ground altitude, extracted from Lidar campaigns. These are generally made available with an open licence, and can be of two kinds, point clouds or Digital Surface Models (note, all links to the data we use, publish, and to the open source code are reported in sect. 11).

A point cloud is a matrix of scattered points that is the result of the Lidar measurement, while a DSM is the elaboration of a point cloud that produces a tessellation of the surface, with a certain sampling, after the application of some smoothing algorithm.

If the provided data is not a DSM, we need to refine the data (*Lidar Data Refinement* block) producing a DSM to make the surface usable. Even if a DSM is provided, we may need to convert data from one of the various formats used to release surface data, and cut out an area that is large enough to perform meaningful experiments while being still small enough to fit the 16GB of RAM our GPU has, as line-of-sight is checked using CUDA libraries on GPUs. Fig. 2 shows a detail of a raster image representing a DSM with buildings in Barcelona.

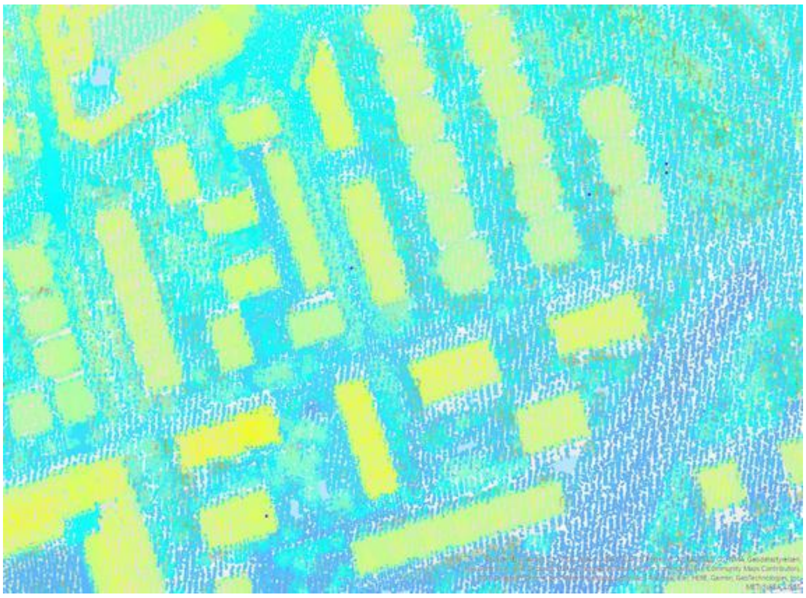

Fix 2: An example Lidar data, lighter colour means higher elevation.

The next step is the isolation of buildings inside the area, which can be done starting from two data sources, the Open Street Map (OSM) building database or other public administration databases (*Building Data Collection, DSM Cleaning* block).

Co-funded by the Horizon 2020
Framework Programme of the European Union

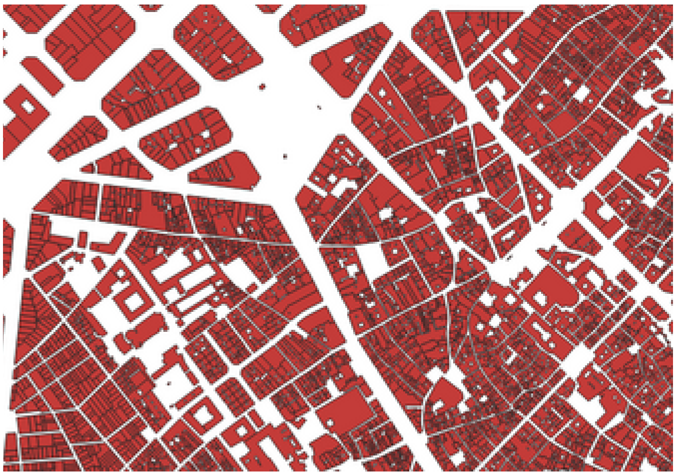

Fig. 3: the building map provided by OSM.

The choice depends on the precision and availability of the second ones. In general, in urban areas, the OSM database contains more building polygons than open data from public administrations, and is better updated, while in rural areas the opposite can happen. In our case, since we studied only urban areas, we opted for OSM databases after reviewing the availability of other sources. Fig. 3 reports the shapes of the buildings, again, in Barcelona.

In the same block we generate a refined 3D coverage scenario in which we remove obstacles from streets and other public areas (we refer to them as "street points" for brevity), which were included in the Lidar measurement, for instance, vehicles in streets (see Fig. 4).

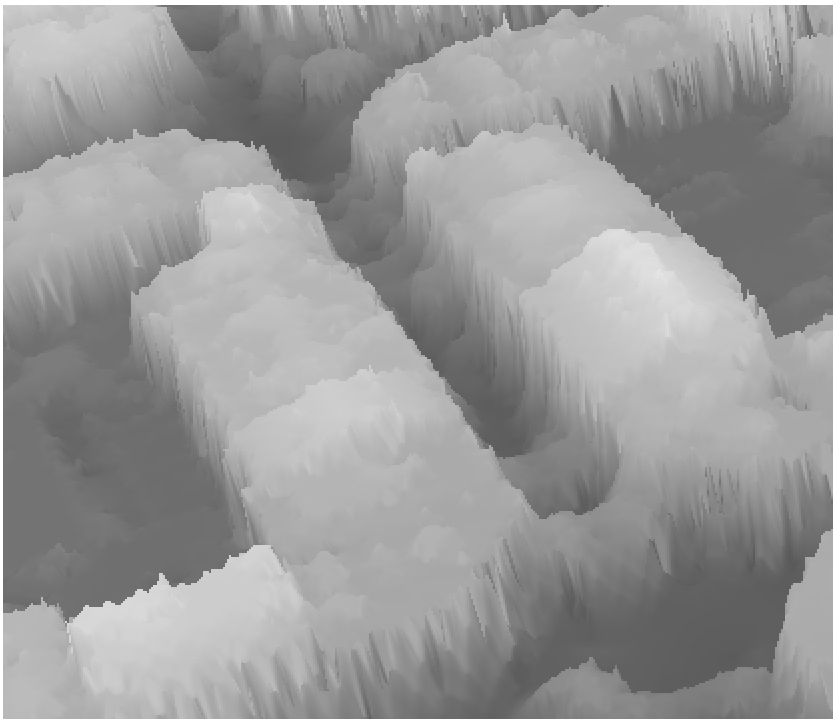

Fig. 4: a detail of the DSM. Between two buildings temporary obstacles (cars and large vehicles) are present. The smoothing process of the DSM may even merge them to building shapes.

Co-funded by the Horizon 2020 Framework Programme of the European Union

This requires isolating streets, computing an average ground altitude and correcting the altitude on the 3D surface, obtaining the results visually represented in Fig. 5.

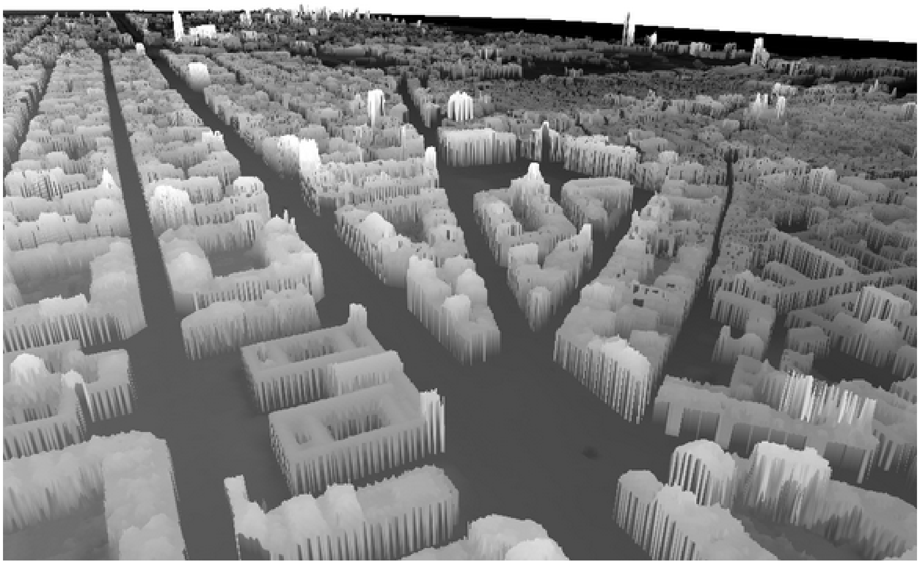

Fig. 5: The final city shape.

The whole process so-far described is time-consuming and only partly automatic, as it requires finding the data sources (with related language difficulties and licence agreements), converting the various data sets in a DSM format, obtaining building shapes, applying the cleaning and merging all the data in a single layer.
The output data format is an nxm matrix where n and m are the sizes of the rectangular area chosen, expressed in metres. Each value of the matrix provides the altitude of the point. Another binary matrix contains zero or one for each coordinate expressing the fact the point belongs to a building or to a street.

This process has been done before the beginning of the project for some Italian cities (Florence among them). During this project we applied it again to also non-italian cities, to extend our coverage. **We generated data-sets for Milan, Barcelona, and Luxembourg city**. We thus published the coverage 3D scenarios for 4 areas between 0.5 and 0.7 square km, of which three were produced during the project (plus Florence). **The output of this phase is a first publishable set of open data made of the 3D shapes that represent some portion of urban area (output block 1 in Fig. 1).**

The data can be used to define an optimal placement of gNBs in the urban area. Ideally, when designing line-of-sight mobile network coverage, every point in the street should be covered as we don't want blind spots. In practice this can not be achieved as in some cases a new gNB should be placed to cover only a few square metres. A reasonable approach is to set a threshold of LoS coverage (e.g. 95%) and cover the remaining part with lower frequencies that can be used also without LoS. This introduces the problem of giving a priority to some ground point over another. Depending on the application there could be 2 approaches, a simple one in which every point has the same weight, and one in which we define a weight function.

Co-funded by the Horizon 2020
Framework Programme of the European Union

A sensitive weight function could be given by the probability of having a person in that point, which requires data that are not available, or mobility patterns that are specific to the urban area which we don't have. A second approach we adopted was to restrict the application to vehicles, for which, precise mobility models are available in the literature. Specifically, for the city of Luxemburg there exists a detailed model implemented in the SUMO simulator (on top of the Omnet++ discrete events simulator), that uses precise vehicle measurement campaigns and can reconstruct vehicle flow in the city using OSM maps.

During the project we have realised a publication submitted as a short paper to the IEEE/IFIP Wireless On-Demand Network Systems and Services (WONS) conference in which we run SUMO to simulate realistic patterns of traffic for 24 hours, we overlay the map to the 3D scenario we realised, and for each street point in the 3D scenario we sample the presence of vehicles once per second [Gemmi2023]. The paper was accepted on Dec. 5th. At the end of the simulation we have an nxm matrix of integer points that assigns a weight to each square metre in the map. The whole process and its application is described in more detail in the dection "*Demand Model for Vehicular Networks*" later on in the deliverable.

The *Generate Demand* block produces, for the Luxemburg scenario, a new matrix of weights that we publish as open data (output block 2 in Fig. 1).

The next step is the application of the optimization algorithms to define the best positions in the 3D area for gNBs, which follows the base heuristic approach we introduced in [Gemmi2022] in the unweighted case, and the weighted heuristic presented in [Gemmi2023] for the case of Luxemburg. This has been repeated with an increasing density of gNBs per square kilometre obtaining the heuristically-optimal placement with increasing percent ground coverage.

The output block 3 in Fig. 1 then produces two sets of open data for each area and for each density of gNBs. The first is the set of coordinates of the position of gNB in the nxm grid. The second is the set of binary matrices that specifies for each gNB what is the subset of points in the street that the gNB covers. In the case of Luxemburg we have two data-sets, one for the unweighted version and another one for the weighted version.

For each scenario we also evaluate the presence of LoS between each couple of gNBs, and create what we call the visibility graph. This is saved in the standard .graphml format and published, for each scenario and density. Note that running the heuristic, depending on the configuration parameters, needs tens of minutes per area, on a 16 cores server (Intel Xeon Gold 6342 CPU @2.80GHz), with 64GB of RAM and an NVIDIA GPU (Tesla T4). **Again, this is part of the output block 3, in which we provide 4 different densities (45, 60, 75, 90 gNB/km^2) for four urban areas.**

The following step is to use data from Colosseum Emulation to enrich the visibility graph with performance metrics. Ideally, we would have used Colosseum to implement the whole scenario and test the capacity of each link connecting every couple of gNB in LoS, this would have generated new information per each link with the maximum bit-rate and probably

Co-funded by the Horizon 2020 Framework Programme of the European Union

removed some of the edges that were introduced due to their very low performance. As explained in the next sections this was not yet possible. We were instead capable of verifying that the potential links in the graph were usable, meaning that their length was short enough to allow the 5G stack to negotiate a link and synchronise.
In the follow-up of the cooperation with NEU we will try to address the performance issues appearing with OAI and improve the generated open data produced at output block 5 in Fig.1, including the estimated bit-rate per link.

We then run the optimization algorithms proposed during the project to create the backhaul graph. The optimizations models were briefly described in D2, their details are reported in SUB2, and they are summarised later on. The algorithms can be purely topological or flow-based. They could be successfully executed and produced feasible backhaul graphs and the position of the donors in the network, with or without arch redundancy. This was included in the .graphml output setting an edge label to be "off" in the backhaul topology if unused, and for each node a label of the values “donor” or “gNB”. The optimization process is time-consuming, we report the CPU time needed in the rest of the deliverable. **The open data we produced at output block 6 contain two series of .graphml files marked as “single” or “multitree” depending on the model used**.

Finally, we converted the provided topology in a format that can be used by an rApp to implement the whole backhaul network. This requires a detailed matching between the generic graph format we used for open data and the implementation of IAB chains using real radios, links configuration and the set-up of the correct interfaces as described in Fig. 8. **Output block 7 in Fig. 1 produces another set of .graphml files in the format the rApp can use.**

> To the best of our knowledge the open data we published represent the first-of-its kind data-set merging realistic data for 5G ground coverage in LoS and backhaul formation derived for 4 different urban areas and 4 different gNB density. The data-set enables new research in the many topics related to this area.

Demand model for Vehicular Networks

The problem of defining a realistic demand model to be used as an input to the optimization procedure is a challenging one. In this section we describe the work submitted and accepted to the IEEE/IFIP WONS conference [Gemmi2023] in which we tackle the problem using a vehicular network simulator, restricting the problem to vehicles only. The paper was jointly authored by Gabriele Gemmi and Leonardo Maccari, and Michele Segata (university of Trento) as an expert of vehicular network simulation. The paper starts from one of the open-data sets produced during the project and introduces:

1. A demand model
2. A dedicated heuristic to optimise gNB placement considering the demand

Co-funded by the Horizon 2020 Framework Programme of the European Union

Here we describe only the first point that is the most relevant to the project, the rest of the details can be found in the paper.

We use the urban traffic simulator SUMO to generate realistic mobility traces of the city of Luxembourg. In particular, we make use of the LuST scenario, a publicly available scenario generated from traffic data provided by the Luxembourg government which includes both public and private transportation over a period of 24 h.
To obtain traffic traces, we run the scenario over the full 24 h for a total of 286215 vehicles moving on the streets. The simulation step is set to 1 s and, at the same frequency, we log the positions of the vehicles in the area of the city shown in Fig. 6, corresponding to roughly 4 square km (a superset of the area we publish as open data in the project). We collect traces using GPS (latitude/longitude) coordinates and then convert them to a .gpx file for later processing.

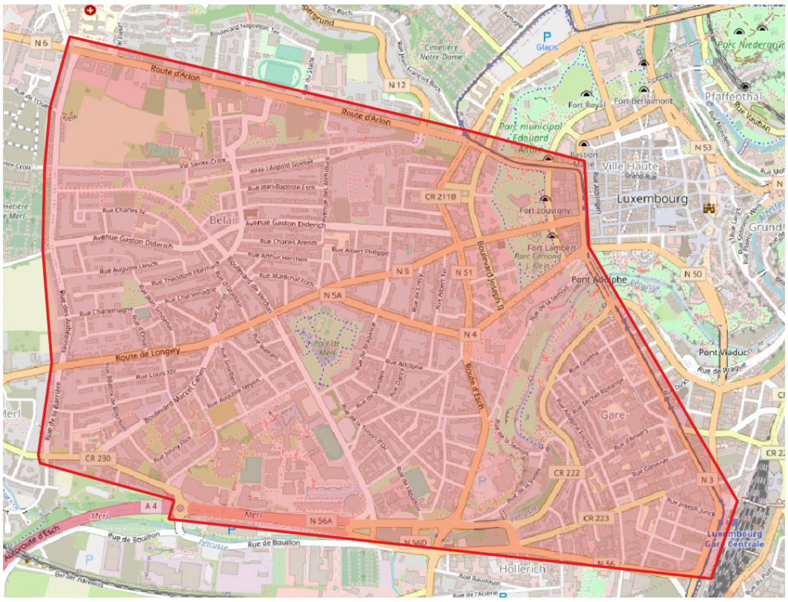

Fig. 6: The area selected for the simulation in Luxembourg city.

To generate a matrix of weights from these traces, we first rasterize the traces mapping each logged position to a cell of a matrix, where a matrix cell represents one square metre. We obtain a matrix τ^* in which every (x,y) element takes value *n* if *n* vehicles have passed from the cell (*x*,*y*) during the whole simulation. For the sake of readability we rescale it to the number of passages per minute, fig. 7 shows the empirical PDF of the values of the cells with non zero value, binned with bins of size 0.125 passages/minute. It can be seen that the majority of the cells have less than one passage per minute, with the 95th percentile roughly at 0.55. The distribution is pretty skewed, with about 5 orders of magnitude between the largest and the lowest frequency.

Co-funded by the Horizon 2020 Framework Programme of the European Union

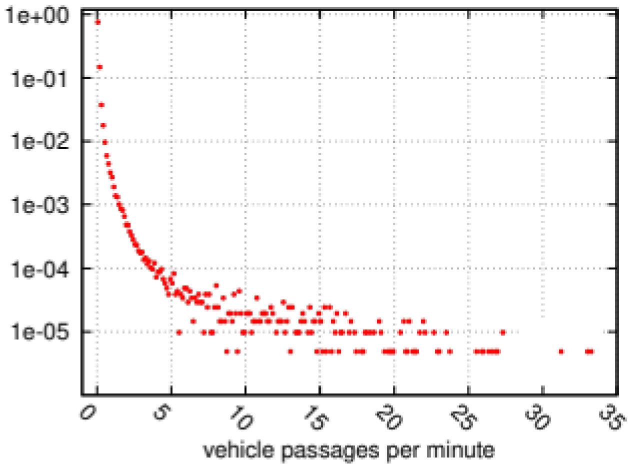

Fig. 7: PDF of the passages per minute.

When we need to use this score in an optimization or heuristic algorithm we need to map the values to some integer range that is representable in the problem formulation. We chose an `uint8` type that works as follows: The values up to the 95th percentile (0.55) have been linearly mapped in the interval [0,64), while the rest of the values have been linearly mapped to the range [65, 255] as follows:

$$\boldsymbol{\tau}_{x,y} = \begin{cases} \lceil 116 \cdot \boldsymbol{\tau}^{\star}_{x,y} \rceil & \text{if } \boldsymbol{\tau}^{\star}_{x,y} \leq 0.55 \\ \lceil 9.34 \cdot \boldsymbol{\tau}^{\star}_{x,y} + 58.77 \rceil & \text{if } 0.55 \leq \boldsymbol{\tau}^{\star}_{x,y} \leq 21 \\ 255 & otherwise \end{cases}$$

Note that we have considered the values greater than 21 as outliers and thus they are all equally mapped to 255. We finally have an integer weight that can be used to rank each cell of position (*x*,*y*) we want to cover.

<u>Practical Applications</u>

The result of this process is that every point in the ground now has a weight, which can be used in two ways. In the accepted paper this strategy is used to optimise the placement of gNBs in order to cover 95% of the car passages (and not of the ground). The paper shows that prioritising ground points based on their weight influences the optimization in a critical way. **A second use, that is the next step in this flow of research, which we drafted in SUB2, is to use this score for backhaul optimization**.

The flow-based optimization model, in fact, requires a demand value for each gNB, which represents the aggregated traffic demand from its nodes, which will be used, together with the bit-rate attribute for each edge in the flow-based optimization to optimise the backhaul topology. The demand for a gNB can be derived by summing the weights of the cells covered by the gNB, expressed as an integer. The sum of all elements of the matrix is a unit-less value proportional to the probability of having user equipment (UE) (vehicles, in this case) in the

Co-funded by the Horizon 2020
Framework Programme of the European Union

area, but does not take into account the possibility of overlapping areas. A more elaborated metric is described in SUB2, it takes into consideration the overlapping areas and transforms the unit-less value into a bit-rate.

Testing on chain topologies

We can now discuss the results of the experiments we performed on the multi-hop paths generated using the IAB-like solution we described in more detail in D2.

We performed several experiments to validate the approach we had proposed in the previous deliverables, all the experiments were set-up using the code we realised to create multi-hop network topologies using the IAB-like interface. The goals of the experiments were to:

1. Validate empirically the design of the IAB-like solution, verifying that we had correctly used and implemented the O-RAN interfaces, so that the IAB-like solution and iab-manager are functioning.
2. Test the performance of the backhaul generated with the topology-based optimization models.
3. Provide measurement we could feedback to the optimization models, in particular, estimate the maximum path length and model the available bit-rate in generic conditions in order to feed the flow-based model and produce open data in block 5 of Fig. 1.

Two versions of this test were performed, the first in a toy scenario in which the nodes are artificially placed in a straight line, and the second in which the paths are realised in the realistic scenario of one of the cities we modelled.

The results outlined some limitations of the OAI code base, that is currently under heavy development and thus, the overall performances are yet not close to the expected ones in a real scenario. Let us recall the configuration of an IAB chain as in figure 8:

Co-funded by the Horizon 2020
Framework Programme of the European Union

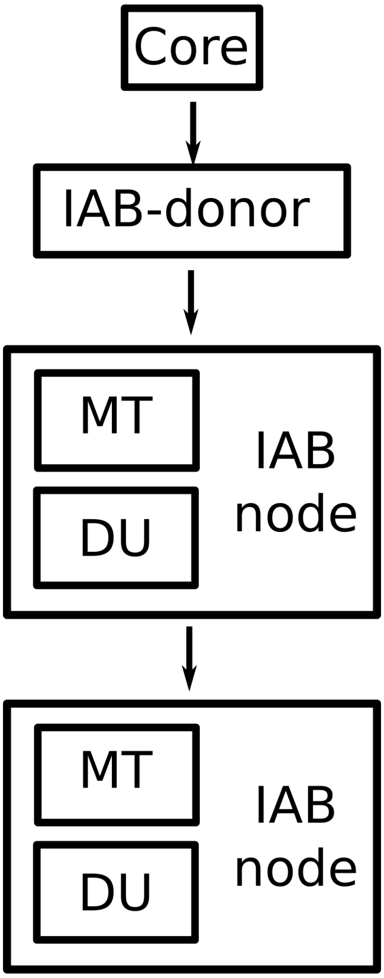

Fig. 8: configuration of an IAB chain made of one IAB-donor and two IAB-nodes

The chain is made of a connection from the core network to the IAB-donor, that is then connected to a sequence of IAB-nodes. Each IAB-node is implemented as two modules, a Mobile Terminal (MT) and a Distributed Unit (DU), each one has a different radio module. in order to avoid spurious interference we used a different channel per link, which simplifies the emulation, using two different 40 MHz channels in the C-band. Each MT and DU are running in separate linux containers and are connected to the software defined radios that are part of the Colosseum testbed. Each connection passes through the MCHEM channel emulator, each MT receives signal only by its father DU with an RSSI of -54dBm, enough to support the highest encoding scheme (see the table at the end of this section).

We ran two kinds of tests, one with the ping application in order to measure the round trip time between every step in the chain and the core network, and a second one in which we tested with the iperf tool the available bit-rate on the chain.

Co-funded by the Horizon 2020
Framework Programme of the European Union

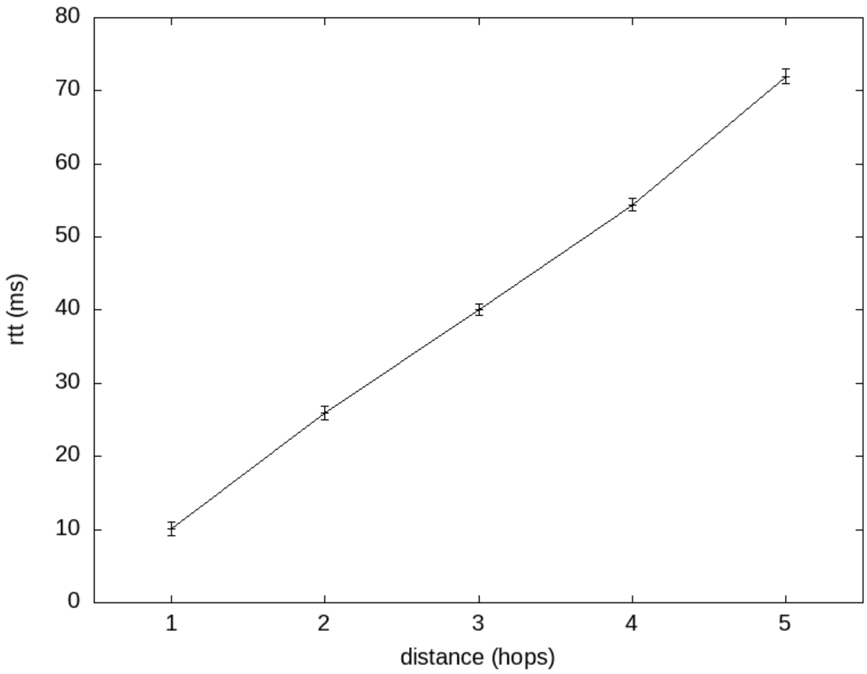

Fig. 9: measured RTT in a multi-hop communication.

Fig. 9 reports the measurement of the round-trip-time on the chain. Each point is realised as the average (with 95%-confidence intervals) over a 40-second communication with one 64B packet per second and zero packet loss. The RTT at 1 hop reports the time to ping from the MT of the first IAB node to the core (and back) using twice the wireless link between the IAB-donor and the IAB-node. At 2 hops we are measuring the RTT from the MT of the second IAB-node and so forth. The link between the donor and the core introduces negligible (sub-ms) delay.

The results show that on a multi-hop path, we can achieve reasonable RTT for real time applications, even including all the real-world constraints like encoding/decoding and routing time. Roughly, at every hop, 7 ms are added in each direction. Consider that [TR38913] introduces a requirement of 10ms one-way delay at the control layer to send the first frame, and a 4ms delay on the link for each direction, so a target RTT of 18ms (for broadband applications). This is the delay introduced by the network, while under a user perspective, this delay must be incremented by the time that upper layers require to handle the packets in the operating system at source and destination.

What we measure in our tests is application-to-application traffic, thus it includes every kind of delay, yet up to two hops we are close to the requirement (about 25ms).

It is of course fundamental to note that in Colosseum the whole infrastructure is virtualized on linux containers using a software router without any hardware acceleration. So a large part of the delay is due to the multiple encapsulation and routing inside the virtualized software chain. It is impossible to separate the sources of delay, so we can not decouple the real delay from the delay introduced by the virtualization of all components. Recent works

Co-funded by the Horizon 2020
Framework Programme of the European Union

based on simulations measure a per-hop delay of 4-6 ms [Polyakov2022], not far from what we measured.

As a first key feedback from Colosseum we then decide to limit the depth of the backhaul tree network to a maximum of three hops and use this value as an input to the topology formation optimization. While 2 hops would have been more conservative, three hops enable more complicated topologies and put our optimization algorithms under higher stress. Moreover, typical topologies discussed in the literature and in discussions with industry consider 3 hops as a practical bound on the number of IAB hops.

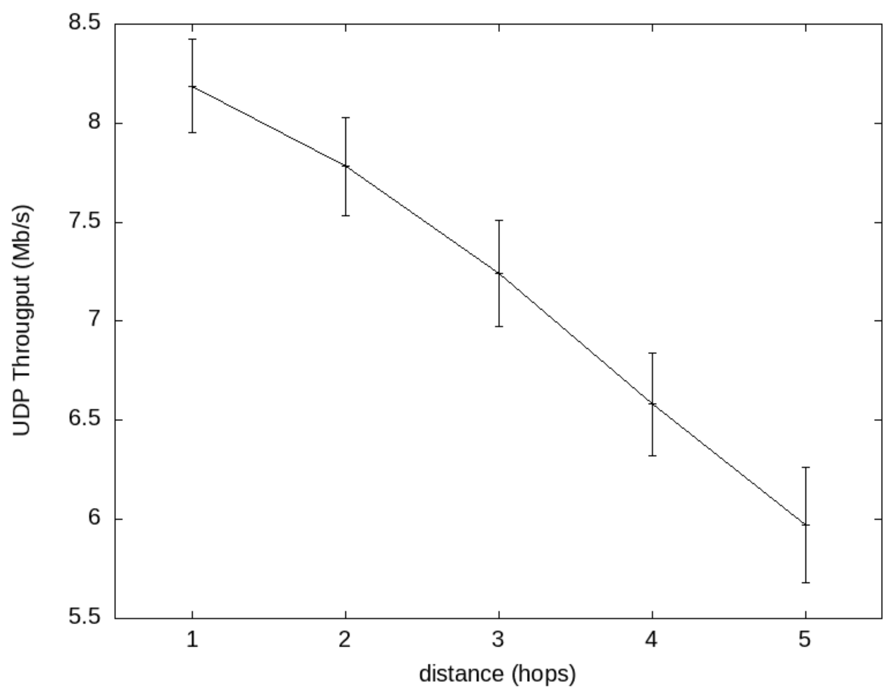

Fig. 10: measured throughput over an UDP uplink connection.

The follow-up experiments we have made describe instead the expected capacity of the wireless links on multi-hop path. In this case, the scenarios present anomalous behaviours. On a single link, we were not able to achieve more than 10Mb/s bit-rate in uplink, and even less in downlink, while instead using a 40MHz channel, we were expecting up to theoretically 40 Mb/s (Fig. 10). Adding links, the throughput decreases almost linearly.

The reason for this this limitation was identified in the state of the code of the OAI software UE (softUE), which is a work in progress project, contributed by all the major players in the 5G arena[1] but still far from a stable state. In the next section we document the efforts we - as the OR2AN2G project and also the US partner from NEU, stimulated by the findings

[1] See https://openairinterface.org/osa-members/

Co-funded by the Horizon 2020 Framework Programme of the European Union

emerged during the project - took to contribute to solve this problem. Our work contributed to track down the problem and partially solve it, but not in time to include the results in the deliverable. As said, Gabriele Gemmi will remain at NEU for an extended period, up to March 2023, and we hope that we will be able to contribute to solving the present issues and perform more reliable tests on link capacity, in order to complete SUB2.

We then document the actions we took to improve the open source code of OAI, this forced us to partly deviate from the initial plans and focus on the other parts of the project (data publication, source code publication, demand model and optimization design).

Compatible results were collected when the emulated scenario was set to be a chain of nodes extracted from a complex topology made of tens of nodes on Colosseum, using iab-manager to build the whole network. The set-up of the network was successful, the multi-hop paths could be used to send low-rate traffic, such as ping packets, but the throughput at saturation was still presenting the same limitations.

> The experimentation on the Colosseum testbed on the one hand confirmed that the IAB-like solution and the extensions we developed for iab-manager are successfully working, as we can set-up complex topologies and test that the communication takes place. On the other hand it evidenced the limitations and some instability of the current implementation of OAI, which we tried to identify and improve.

Contributing to the improvement of OAI

We have worked together with the OAI team to improve the code and help the development of this key component of the O-RAN ecosystem. Some of the actions that were taken during the project were:

1. The problem was studied in several Colosseum scenarios to isolate the cause, and exclude that the IAB-like solution or other code we introduced was causing the problem. The issue was identified to be in the softUE.
2. We replicated the tests out of the emulator in a lab test-bed using real devices, in order to confirm the component that was responsible for the anomalous behaviour. This involved the configuration and use of software defined radios (USRP X310) running the gNB software implementation from OAI, and a OnePlus Nord 5G telephone running its own 5G software stack. The MAC and physical layer of the link was configured to provide a throughput of up to 132 Mb/s[2] and with a TCP connection speed test, we could achieve 125Mb/s. This experiment, which was also replicated by other developers in the OAI community, confirmed that the current state of the softUE had some unsolved problems.

[2] 40MHz channel width, MCS 28.

Co-funded by the Horizon 2020
Framework Programme of the European Union

3. Since Eurecom (the organisation leading the development of OAI) holds weekly meetings where all the parties interested in the development of OAI, including NEU, discuss future developments and report ongoing issues, we participated in the meetings. Gabriele Gemmi (UniVe) has attended these meetings and helped identify the throughput problems affecting the softUE on the basis of the previous point. An issue was officially opened in the issue tracker of OAI[3].

In order to improve OAI, NEU has recently launched the Open6G centre to contribute to the improvement and development of OAI, including the softUE[4].

By the end of the project, the developers at NEU were able to track down some of the problems in the source code which led to partial improvements of bit-rate, both in uplink and in downlink. The following links show the discussion in the OAI open source community, initiated by NEU developers:

https://gitlab.eurecom.fr/oai/openairinterface5g/-/merge_requests/1815
https://gitlab.eurecom.fr/oai/openairinterface5g/-/merge_requests/1811

The proposed patches implement parallel decoding of the frames, improving the final performance. We helped testing the performance of the new version, here we report the results of a single experiment obtained on a one-hop link in Colosseum:

Received power (dBm)	UL speed (Mb/s)	Packet Loss	DL speed (Mb/s)	Packet Loss	UL MCS	DL MCS	Max. Bitrate UL/DL Mb/s
-54	22.5	10.00%	28.58	5.30%	24	28	32/132
-57	21.7	7.70%	29.3	3.70%	23	24	30/108
-60	17.8	14.00%	28.3	5.00%	22	22	28/94
-63	19.2	10.00%	28.12	7.50%	22	19	28/72
-66	16.9	20.00%	28.34	5.90%	22	17	28/62
-68	17.927	15.30%	25.87	14.30%	22	17	28/62

These measurements were taken using the UDP transport protocol, channel width of 40MHz[5], with 7 Downlink slots, 2 uplink slots and 1 Shared slot, the transmission power was

[3] See https://gitlab.eurecom.fr/oai/openairinterface5g/-/issues/573
[4] See Open6G.wiot.northeastern.edu
[5] 106 PRB and 30kHz of subcarrier spacing

Co-funded by the Horizon 2020
Framework Programme of the European Union

set to 30dBm, the antenna gains were set to 10dBi and the configured distances produced the RSSI reported in the table. The table also reports the expected maximum performance in Uplink and Downlink. Even if we measured an improvement compared to previous experiments for the Uplink, Downlink is still far from what was expected, and we measured a packet loss changing in an inconsistent way with the received signal, suggesting that the performance still needs further investigation. This, plus the fact that these results came at the very end of the project did not allow us to use the data.

We then focused on the experimentation of two issues of key scientific and industrial relevance for the success of OR2AN2G and 5G in general:

- Refining and testing the usability of the optimization models, and showing that in real conditions they can be solved;
- Proposing the already mentioned demand model in [Gemmi2023], which solves one of the limitations we outlined in deliverable 2.
- Developing all the extensions to OAI and the rApp that can set-up a network in a real scenario, with the configuration provided by the optimization solution.

As said, Gabriele Gemmi will remain at NEU for an extended period, up to March 2023, and we hope that we will be able to contribute to solving the present issues and perform more reliable tests on link capacity, in order to complete SUB2.

Testing the feasibility of the optimization procedure

The optimization procedure briefly documented in deliverable 2 has been detailed in SUB2, it takes in input a graph representing a realistic network deployment. The graph contains the position of all gNBs, obtained with state-of-the art heuristics, and applied to the open data we generated for the project. In our initial plans, the graphs should have also contained bit-rate per link measured on Colosseum. Each graph is made of :

- a set of nodes, with their geographical position;
- a set of edges that represent the links with line of sight;
- An estimation of the link capacity of each edge (not available at this stage);
- if available, an estimation of the traffic demand of each node, derived from the covered area.

In D2 we introduced three different optimization strategies, whose goal is to define the minimal number of donors, and the active links in the topology in order to have every node connected to a donor with a multi-hop path. The result is a tree forest in which every donor is the root of a tree. In a green-field deployment (in which the operator does not have any existing infrastructure) the optimization is run once before network deployment so that the operator can decide where to place the donors and what links will make the backhaul network. In case of a gray-field deployment some of the gNBs can be already set to be donors and the topology can be conditioned to that constraint. When the network has been realised the positions of the donors are fixed, but the topology can be updated in case the conditions change (i.e. a link fails). We briefly described three different optimization problems that achieve three different results, with increasing complexity:

- A single-tree per node topology, with the only constraints given by the tree depth and the maximum degree of every node. This model creates a tree forest with one donor per tree (with minimum number of trees set to 1) in which each node must belong to at most one tree, and each tree has topological limitations. The model is computationally the simplest one as it uses a number of variables that scales quadratically with the number of nodes.
- A multi-tree per node topology. The model is based on the previous one, but it forces every node to be present in at least *n* trees without overlapping edges. In our experiments we set *n*=2 (*n=3*) so the network topology is resistant to the failure of one (two) edge(s). This model is computationally more complex because it requires a number of variables that scales with up to the third power of the number of nodes.
- A single-tree per node topology that takes into account the estimated link capacity and flow demand of each node. The complexity of this model is the third power of the number of nodes.

Please refer to SUB2 for the exact ILP formulation of all models, the running time on synthetic topologies and the exact data for running time on realistic topologies.

Due to the time-limitation and the absence of the data regarding the bit-rate, we focused our analysis on the first and second one, however, we report also some results on the third one. Compared to D2 we have the following key advancements:

- We completed the multitree model and **produced code that can generalise the optimization to an arbitrary number of disjoint trees**.
- We use the open data-sets from Luxemburg, Milan and Barcelona to generate the feasible positions of gNBs. We started with a density of 45 gNBs/km^2 up to 90 gNBs/km^2 in areas of 05-0.7 km^2. 45 gNBs/km^2 value was estimated to be in average sufficient in one of our previous publications to cover 95% of the city ground [Gemmi2022]. Depending on the specific scenario, this number may largely vary, in some, about half of that density of gNBs is sufficient, in some other cases it may need twice of it. Increasing the density will anyway improve the robustness of the coverage, with gNBs covering overlapping areas.
- We run the optimization software to measure the time to obtain a feasible solution for the backhaul design **for all models and all realistic scenarios**, while in deliverable 2 we tested the optimization only on synthetic graphs.

Co-funded by the Horizon 2020
Framework Programme of the European Union

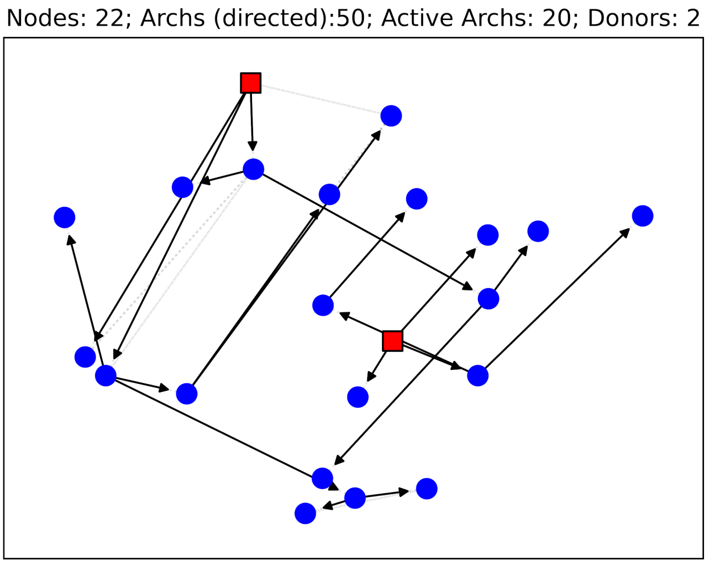

Fig. 11: The topology resulting from the single-tree per node model. Red squares are IAB-donors, blue nodes are IAB-nodes, solid arches are the one composing the backhaul topology, dashed ones are excluded from the backhaul.

We describe qualitatively the results using a small portion of the city of Florence that needs only 22 gNBs (and thus, the result is easier to visualise) and then we summarise results for the other cities and higher numbers of gNBs. Full results are in SUB2. Fig. 11 shows an example result, the position of each node respects the position in the real map. We can see that if we set the maximum tree depth and maximum degree to 3, only two trees are needed to connect 22 nodes, using 20 out of the 50 edges available (counting each edge in its two directions).

Fig. 12 instead shows the solution of the multi-tree per node model. There are two disjoint edge-sets (corresponding to the two colours), 10 donors (and thus 10 sub-trees) and 24 edges. The topology is resistant to the failure of a single edge, as every IAB-node has two incoming edges, one for each edge-set. If one edge fails, one sub-tree is partly disrupted, as some node will not be able to reach the corresponding donor (the root of the tree) but it will be able to reach the core network using the path to another donor in the other edge-set. As a consequence the number of necessary donors grows from 2 to 10.

Co-funded by the Horizon 2020
Framework Programme of the European Union

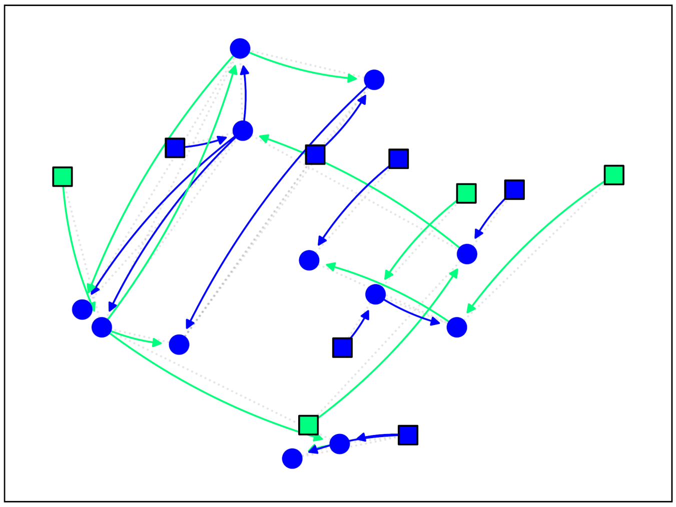

Fig. 12: The topology resulting from the multi-tree per node model. Squares are IAB-donors, blue nodes are IAB-nodes, solid arches are the one composing the backhaul topology, all arches are used in the backhaul. Arches have two different colours because they belong to two disjoint edge-set.

Let us recall that the final goal of the optimization is to find the topology that guarantees the minimal number of IAB-donors, and the links that must be activated to connect all the other IAB-nodes. Two important metrics are the time it takes to find an optimal solution, and the ratio of IAB-donors to gNB nodes, which depends on the specific scenario, as the visibility graph may be more or less dense, and some gNBs may be simply isolated. In fact, the choice of the initial positions of the gNBs is made in order to minimise their number without taking into consideration their connectivity in terms of potential links in the backhaul graph.

Solutions were obtained for the realistic network topologies generated for the cities of Florence, Milan, Barcelona and Luxemburg, with a target density of 45, 60, 75, and 90 gNB per km^2, the actual area sizes range from 0.5 to 1 square km. The size is currently limited by the GPU memory required to load the data and perform the initial gNB placement.
We summarise here some sample results we obtained for the 4 topologies, with different densities of gNB/km^2 , and two different optimization models:

- When the density of gNB/km^2 is 45, the single (multi) tree model takes at most 8.51 (34.8) seconds to complete. The best ratio of donors/gNB for the single (multiree) model is 12/28 (36/49) and the worst is 17/34 (28/32)
- When the density of gNB/km^2 is 60, the single (multi) tree model takes at most 29.10 (108.23) seconds to complete. The best ratio of donors/gNB for the single (multiree) model is 10/37 (21/37) and the worst is 24/43 (35/43)

Co-funded by the Horizon 2020 Framework Programme of the European Union

- When the density of gNB/km^2 is 75, the single (multi) tree model takes at most 57.81 (251.07) seconds to complete. The best ratio of donors/gNB for the single (multiree) model is 16/82 (18/47) and the worst is 26/58 (33/54)
- When the density of gNB/km^2 is 90, the single (multi) tree model takes at most 119.65 (532.5) seconds to complete. The best ratio of donors/gNB for the single (multiree) model is 15/99 (18/56) and the worst is 23/69 (37/60)

We also ran the optimization using different configuration parameters:
- Tree depth = 2 and degree = 2
- Three disjoint trees instead of 2

We obtained consistent results, reported in SUB2.

As a further experiment we also ran the flow-based optimization model using constant values for the link capacity and the node demand. Since we are not using realistic bit-rate data, the most relevant result we are interested to report is the running time. We ran the optimization for the 45 and 60 density values and obtained optimal solutions in less than 1606 and 20566 seconds respectively, which confirms the feasibility of our approach in realistic conditions.

The fundamental take-away from this batch of experiments is that using areas between 0.5-0.7 km^2 with realistic topology constraints derived from our 3D modelling and our initial experiments on Colosseum to measure RTT, we can now provide insights on the scalability of the algorithms that we propose:
- Under a delay perspective, it is reasonable to have trees of depth at most 3 supporting real time applications.
- With such depth 3 (or 2) our approach is able to optimally allocate the number of donors and the network edges in practical time, with densities that have been shown to be sufficient to cover 95% of the ground. This is true for the single tree model, the multi-tree model (up to 3 disjoint trees) and for the flow-based optimization.

Starting from this result and this methodology, several future research directions are available. As a first point we note that the density of the visibility graph in terms of edges per node is not high enough to always guarantee a strong cut of the needed donors, especially in the multitree version. This is because when we decide the gNB positions we don't take into account their inter-visibility, but only the ground coverage. New coverage heuristics should be proposed to generate denser graphs, that allow to further reduce the number of donors. Note that even creating small trees of depth and degree equal to 2 could produce a donor/gNB ratio of 1/7, without stretching too much the topologies.
We also need precise estimations of the reduced costs using a detailed cost model, which can tell how much the cost grows for every added donor, in order to find a good trade-off

Co-funded by the Horizon 2020
Framework Programme of the European Union

between cost of the infrastructure and performance reduction due to the multi-hop path and scheduling.

Modelling LoS estimation even beyond buildings

In the literature, this problem of determining LoS is generally tackled using a stochastic geometry approach, in which the presence of LoS is determined in a stochastic way. Given our experience with modelling 3D scenarios, in the period UniVe personnel spent at NEU we decided to start a theoretical modelling effort to advance the state of the art. Even if this research is at its initial stage we report it as a promising direction of joint research between UniVe and NEU.

Experimental studies in the literature have derived the probability of LoS in an urban area as a function of the distance between the gNB and the UEs, these models, such as the 3GPP UMI channel model[6] are derived from real measures in urban areas. One limitation of the models is that they assume a radial symmetry, as the probability of being in LoS depends only on the distance. We initially reproduced the dependency between the LoS probability and distance from two points in our 3D topologies and compared it with the 3GPP model.

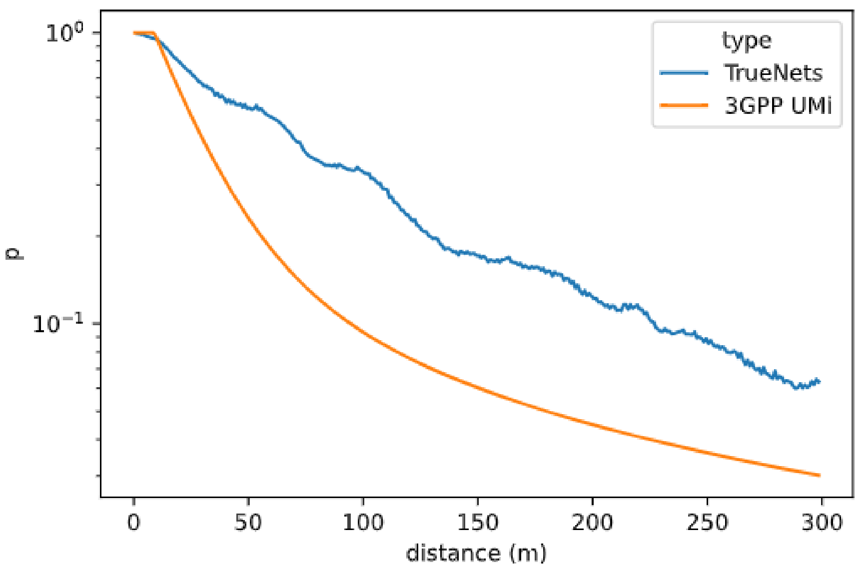

Figure 13: p(d) probability of LoS as a function of the distance between the user and the AP: 3GPP model and empirical data from the Florence scenario with 90gNB/km².

The tests in the Florence area show a substantial disagreement between the data collected with the literature results, which, if confirmed on other areas leaves space for new proposals. Moreover we aim to use stochastic geometry to study the correlation in the LoS probability introduced by the 3D shapes of obstacles, to confirm or confute the assumption of radial symmetry. Let us consider two distinct points X and Y in the area. For simplicity,

[6] TR 38.901

Co-funded by the Horizon 2020
Framework Programme of the European Union

assume they are two gNBs in LoS to the same UE located in point A. Let us call d_{XY} the Euclidean distance between points X and Y. Let L_X (L_Y) be the random variable that assumes value 1 if point X (Y) is in LoS with A. We are interested in the derivation of the joint probability $P(L_X=1, L_Y=1)$. We expect that when d_{XY} is small, we will have $P(L_X=1, L_Y=1)$!= $p(d_{XA})p(d_{YA})$ since the shape of the obstacle introduces correlation, and thus the radial symmetry is not confirmed. This correlation is important to model, for example, the probability of having more than one gNB in LoS with an UE in order to provide reliability.

The follow up of this research consists of three phases:
a- Verify the impact of the correlation between near points in the scenario. The method will pick a random point A in the space and an annulus of radius R and R+i centred in A. We sample K points in the circle and perform an independence statistic test on the probability of being in LoS with an A. As R grows, we expect the correlation to become smaller. The process will allow us to experimentally determine the minimum R such that the correlation between the sampled points and X becomes negligible.
b- Determine the correlation between L_X and L_Y as a function of d_{XY} with the use of stochastic geometry with random shapes.
c- Validate the stochastic model with random sampled points on the dataset collected during the project.

This has two interesting research goals, the first is to produce generic data-based models that can be used to study the performance in areas for which open data are not available.
The second is that this approach could be reproduced at a smaller scale to complement the data-based approach. In fact, one of the intrinsic limitations of our work is that we do not consider the presence of obstacles besides the building shapes, as when we are performing data cleaning, we remove obstacles that are not inside the shape of a building. With a stochastic geometry approach we could add a random component to the LoS determination that 3D analysis produces.

Co-funded by the Horizon 2020
Framework Programme of the European Union

rApp and RIC integration

Next Generation Radio Access Networks (NG-RAN) dramatically increase the complexity of mobile networks with heterogeneous wireless technologies, and self-backhauling capabilities offered by IAB. This poses several challenges to the vertically integrated traditional management and monitoring tools. The O-RAN Alliance introduced a generic architecture and several open standards which allow third-party software solutions to be used to manage NG-RANs. In the context of OR2AN2G, monitoring interfaces such as the O1 O-RAN interface, together with the use or rApps and xApps running in the RAN Intelligent Controllers (RIC) are of particular interest, as they allow the automation and control of different aspects of NG-RANs.

One of the objectives of the project was the study of how these interfaces and RICs can be used to deploy and monitor the backhaul. In the last part of the project we implemented an rApp, running in the non-real time RIC, that can automatically deploy the IAB topology. Moreover, by implementing the O1 interface in both OAI gNB and softEU (used in the IAB node) the rApp is able to collect radio link metrics from all the IAB-nodes of the network and react to changes.

It is extremely important to stress that our effort went in the direction of producing code that can be directly reusable in a real 5G scenario. For as much as possible we used open-source frameworks that already implement the standard interfaces and when needed we implemented messages that extended the standards. When the standard was lacking (recall, O-RAN and OAI are still work in progress) we used open source solutions that are re-usable as is. This of course required a larger effort in studying the standards and the available code compared to any custom-developed solution (which would have been much easier to implement) but fulfils the goal of the project of producing code that can have direct impact beyond the end of the project. The code-base is published and documented on github.

Co-funded by the Horizon 2020
Framework Programme of the European Union

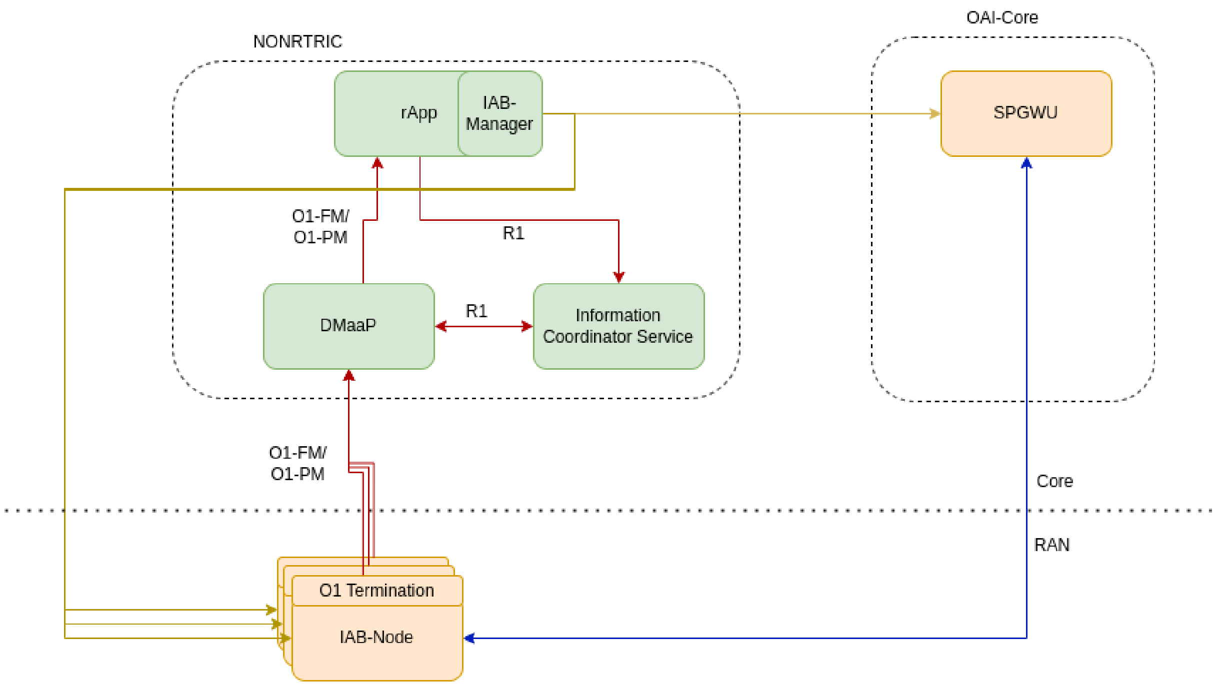

Figure 14: Architecture of the non-real time RIC and rApp.

RAN Intelligent Controller (RIC)

The O-RAN architecture foresees two different RICs. The first one is called near-real time RIC and should be deployed on the network's edge. It should manage a relatively small number of network nodes in a timescale varying from 100ms to 1s and runs the so-called xApps. The second one, which is called non-real time RIC, should be deployed in the network's core and manages a large number of network nodes in a timescale longer than 1s, by taking advantage of the so-called rApps. As explained in D1, the deployment and monitoring of IAB networks comprises the management of the whole backhaul topology, made of tens of gNBs and thus it must be made at the network core. Furthermore for the time being, it can not realistically follow real time dynamics, thus we decided to develop our IAB intelligent controller as an rApp.

Currently, the only open-source platform to realise a non-real time RIC is NONRTRIC[7], developed by the O-RAN Software Community, which is already integrated in a Service Management & Orchestrator (SMO) framework called ONAP[8], developed by the Linux Foundation. By following the documentation available in ONAP, we designed a rApp that integrates with NONRTRIC. A shown in Fig. 14. the rApp can subscribe to different information services through the Information Coordinating Service (ICS) using the R1 O-RAN interface. The relevant messages are then filtered and pushed from the RAN to our rApp by dMaaP, which is the NONRTRIC message broker.

[7] https://docs.o-ran-sc.org/projects/o-ran-sc-nonrtric/en/latest/overview.html
[8] https://www.onap.org/

Co-funded by the Horizon 2020
Framework Programme of the European Union

This choice makes it possible that any network operator could, in principle, deploy our rApp in a production network where either NONRTRIC or ONAP are already deployed.

O1 Interface Implementation

To maintain a representation of the IAB topology in the rApp, the IAB nodes must be able to periodically transmit KPI indicators of the radio links to the RIC. In order to do this, we have defined new O1 Performance Management (O1-PM) messages, containing aggregated metrics from the Physical, MAC and RRC layers such as the used MCS, the Block Error Rate (BLER) and the SINR. These messages are transmitted periodically from the network nodes to dMaaP. Additionally, we have defined O1 Fault Management (O1-FM) messages, which are sent from the gNBs as soon as a link failure is detected so that the rApp can immediately react and reconfigure the backhaul. This required modifications to OAI in both gNB and softUE code base since both software are used in the IAB-like implementation. The periodic reporting has been implemented as a separate C thread in order to avoid interfering with the real time threads managing the radios.

Together with PM and FM interfaces, O1 provides a specific control interface based on the NETCONF protocol to deploy and modify configurations on the network node. This could be used to push topology modifications to the gNBs. Unfortunately, OAI has not implemented this interface and besides OAI, the number of O-RAN compatible gNBs implementing this interface is still scarce. Moreover, IAB has not yet been integrated with O-RAN and thus the standards to interact with IAB nodes are still undefined.

For this reason, instead of implementing a non-standardised interface, we have decided to rely on iab-manager, the tool we extended and used to orchestrate IAB experiments [Moro2022]. Iab-manager behaves similarly, by deploying text-based configurations using the SSH protocol and also allows management of the lifecycle of the different software components running in the network node. We use iab-manager as a library.

The link to the github repository containing the O1 interface implementation of OAI is available in Table 3 of section 11.

rApp implementation

The implementation of our rApp has been done by taking advantage of an open-source Python web framework called fastAPI[9], enabling us to easily integrate the application with the optimization models that have been implemented using Python libraries. The app exposes two different sets of REST interfaces. The first one is the O1 interface where the PM and FM messages are pushed by dMaaP. The second one is used by the user and allows them to initialise the app and deploy the IAB topology. In the current release the rApp is able to implement the single-tree topology generated by the single-tree optimization models.

[9] https://fastapi.tiangolo.com/

Co-funded by the Horizon 2020
Framework Programme of the European Union

Once the topology is deployed, the rApp listens for O1 messages and uses them to maintain an updated graph representation of the RAN, where each vertex represents either an MT or a DU and each edge a wired or wireless connection, with its associated metrics.

When some link in the network fails, the rApp is able to reboot the network and implement a different topology. This is possible, for instance, if we use the multi-tree optimization model in which one tree is a backup tree to be used when the first one fails. Our rApp can be trivially extended to support this healing strategy.

It is important to note that the optimal solutions for the optimization problems are found in a few seconds for low density and a few tens of seconds for higher densities while the flow-based model takes tens of minutes and a few hours for 45 and 60 gNB per square km. With enough computing power we could pre-generate topologies that exclude any link/edge or some combinations of them, and obtain tens/hundreds of valid topologies that can be enforced when a link or an edge fails, from which the rApp can simply choose the correct one when some node/edge fails. Note also that while the problem of jointly deciding donors and edges needs to be solved at design time, the problem of network healing is a simpler one, as the donors can not be changed once the network is deployed, and thus it requires less time. As a future work we will test the feasibility of an approach in which the optimal solution is found in real time.

Note however that we focused on non-real time network healing because the implementation of all Handover procedures are missing from both the implementation of the UE and the implementation of the gNB in OAI. Without the handover, the IAB-node cannot switch its upstream connectivity without losing the connection to the core, thus it is impossible to seamlessly reconfigure the network. When the handover capability will be present this will open more exciting research directions. For instance using the multi-tree model, we may implement an xApp working at real time that simply switches the single gNB connectivity to the backup tree, locally healing the network without the need to reconfigure the whole backhaul.

5 Present and Foreseen TRL

The various components realised and improved during the project have the following TRL levels:

- The optimization algorithms have been developed and executed in a standard PC using as input the topologies obtained by the analysis of the 3D open-data. They have shown to be fast and thus usable in a lab scenario: TRL 4
- The implementation of the rApp that enforces the given topology in the Colosseum testbed, the OAI extension to the O1 interface (a component of the standard O-RAN architecture), the iab-manager extensions were validated in a lab scenario: TRL 4

Co-funded by the Horizon 2020
Framework Programme of the European Union

6 Exploitation, Dissemination and Communication Status

As detailed in the previous section of the deliverable, the production of open source code and open data were two of the main goals of this project, together with the realisation of scientific publications. We can summarise the actions takes as:

- The open source code to replicate the results of the project has been published. The URLs of all the repositories are reported in section 11.3
- The open data generated during the project, following the steps described in Fig. 6 have been published. The URLs of all the repositories are reported in section 11.2
- The Colosseum platform has been enriched with the scenarios generated during the project and will be used for future experiments
- One scientific paper has been accepted for publication in an IEEE/IFIP conference, using the data generated in the project and a new methodology proposed to estimate node demand.
- Two scientific publications have been drafted. They will be submitted to two conferences as described in deliverable 2. Since these conferences use a double-blind review process we don't provide precise references in this public deliverable. The paper drafts have been attached to the material submitted for project review.
- The project has been presented in a joint meeting at WIoT in the visit organised to NEU in the presence of local researchers and remote researchers, including researchers from the Italian University of Padua and Polytechnic University of Milan.

Three other dissemination activities are being prepared for 2023:

- The inclusion of the themes researched in the project in a specific seminar to be held to Master students of Computer Science in Venice. Leonardo Maccari is the professor of the Network Security Master course, which includes also network robustness metrics[10]. The material produced in the project (network topologies, with redundancy) can be the subject of student experimentation on robustness metrics (end of April/beginning of May)
- In May 2023 Leonardo Maccari will be visiting The University of Sevilla for a series of seminars for students regarding the latest advances in network design. This will also include 5G and the backhaul design.
- A public presentation on the project themes at the "Science Cafè", in which Leonardo Maccari has been invited to participate, in the 2023 programme.

[10] https://www.unive.it/data/course/398301/programma

Co-funded by the Horizon 2020 Framework Programme of the European Union

7 Impacts

Impact 1: Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation.

The project has financed the mobility of Gabriele Gemmi, a PHD student of UniVe that moved to NEU at the beginning of September and spent three months working on the Colosseum platform. The experimental results and the publications are based on the joint work of personnel from UniVe and NEU. Gemmi is confirmed to remain in NEU after the end of the project (Dec/22 - Feb/23) with funding offered by NEU and UniVe Ph.D funding. The visit of Prof. Maccari and Marin to NEU gave the possibility of working side-by-side with NEU personnel, preparing the publications and opening new research directions that will be carried out after the end of OR2AN2G. Maccari and Marin presented the OR2AN2G project results to the whole WiOT personnel, and external guests that were present in the institute. A research plan was prepared based on the publication drafts.

Impact 2: Reinforced collaboration and increased synergies between the Next Generation Internet and the US Internet programmes.

Colosseum is fully accessible from remote, but has a steep learning curve. In the initial phases of its usage, support from NEU personnel helps to speed up the set-up of experiments. Having gone through this process, we are opening the possibility for future students and researchers from UniVe to work on Colosseum and reinforce the scientific cooperation between US and NGI partners. In particular, all the experience gathered so far is key to make the interchange smoother than it has been in this first occasion, also at the administrative level.

The new scenarios we produced become part of the assets of Colosseum and can be re-used by future researchers.

Impact 3: Developing interoperable solutions and joint demonstrators, contributions to standards.

One of the key features of 5G is the use of an open design, as opposed to the monolithic approach of previous generations of mobile networks. The open interfaces are currently under development, the O-RAN specifications and the OAI implementation are both two moving targets. This project produces results using real and reusable code, which has a key impact: It serves as a demonstrator for the validity of the specifications and the implementation. Our experiments outlined some of the limitations of OAI, which we reported and discussed with the international team led by Eurecom, in order to be solved.

Co-funded by the Horizon 2020
Framework Programme of the European Union

We extended OAI realising the code for the O1 interface, needed to monitor the network state and trigger the network reconfiguration in the rApp. We also validated the design of the IAB-like solution we used, that is compatible with the current state of OAI.

The code we realised is interoperable with any other research facility respecting O-RAN. As a confirmation of this, Gabriele Gemmi has applied to the POWDER winter school with the goal of repeating experiments on the POWDER test-bed, which is another relevant 5G testbed in USA[11]. The final notification of its accepted participation will be on Dec. 19th.

Impact 4: An EU - US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the Internet

The code and the scenarios we realised for Colosseum is directly reusable in real networks. Colosseum is essentially an emulated real network, so all the components we develop could potentially be re-used in a real setting. Our work then contributes to the available resources that can foster the market of 5G solutions. 5G is intended to have a multiplayer approach, the open source controller we will develop could be used by a real operator as a base for his on-field work. This is exactly the inter-play that is desirable between research institutions and start-ups, in which the former proposes innovative solutions based on high-level research that are used by the latter to foster a market in the area. The open source code for the optimization of the backhaul, and the new scenarios we make available to the community are examples of this kind of impact.

It is important to stress that the whole toolchain we produced and published, from data collection to optimization and test can be reused by anyone (including operators) to design their networks, optimise them and even test them on Colosseum before deployment.

8 Conclusion and Future Work

Referring to the objectives expressed in D1, and the KPI mentioned in the project submission, we consider the project successful, as we were able to deliver:

- source code to replicate and extend the experiments
- Optimization algorithms
- The rApp
- The open data and scenarios to replicate the experiments
- Scalability results on the optimization algorithms, which confirm their applicability
- We confirmed and tested the possibility to set-up network topologies using real-world code with the rApp, which is the first requirement for self-healing.

Even if some of the experiments could not be completed, we extended the state of the art with a new demand model, we started a line of research on stochastic modelling based on our data: under a scientific point of view the project has produced code, data, one published short paper and two drafts publications. SUB1 will be completed in the forthcoming weeks,

[11] See https://www.powderwireless.net/mww2023

Co-funded by the Horizon 2020
Framework Programme of the European Union

detailing the platform, the interfaces, the application, and will be submitted in January 2023. We are confident that in 2023 we will be able complete our experiments and submit SUB2 in the summer.

There are many future directions opened by the results. Among them we mention the development of cost metrics that can be used to effectively assess the economical gain in using IAB, the joint optimization of gNB positioning and backhaul network in order to produce backhaul networks that have a higher density than the ones we used, and allow more complex topologies, the use of graphs and not trees, so that network reconfiguration can lead to repair a broken link without necessarily switching to a new distribution tree, the follow up on the modelling effort with stochastic geometry. Under a technological point of view we expect OAI to mature and provide some of the components that are still missing, for instance the possibility of fast-handover that is required to have a seamless network reconfiguration upon failures, triggered and guided by the rApp.

Under the point of view of the international cooperation, the project opened a link between UniVe and NEU, which we will strengthen in the next months. Gabriele Gemmi will remain at NEU up to early 2023, one of the target conferences chosen for the publications will be held in Boston and organised also by personnel of NEU, and it will be the next step for a meeting to continue the cooperation between the two parties.

Co-funded by the Horizon 2020
Framework Programme of the European Union

9 References

[Gemmi2023] Gabriele Gemmi, Michele Segata, Leonardo Maccari, "Vehicles or Pedestrians: On the gNB Placement in Ultradense Urban Areas", accepted for publication to IEEE/IFIP WONS 2023.

[Gemmi2022] Gabriele Gemmi, Renato Lo Cigno, Leonardo Maccari "On Cost-effective, Reliable Coverage for LoS Communications in Urban Areas", IEEE Transactions on Network and Service Management, Sept. 2022

[TR38913] Study on Scenarios and Requirements for Next Generation Access Technologies (3GPP TR 38.913 version 14.2.0 Release 14)

[Polyakov2022] Nikita Polyakov, Anna Platonova, "Assessing Latency of Packet Delivery in the 5G 3GPP Integrated Access and Backhaul Architecture with Half-Duplex Constraints" Future Internet 2022, 14(11)

[Moro2022] Eugenio Moro, Michele Polese et Al. "IABEST: an integrated access and backhaul 5G testbed for large-scale experimentation", Proceedings of the 28th Annual International Conference on Mobile Computing And Networking (MobiCom 2022)

Co-funded by the Horizon 2020
Framework Programme of the European Union

10 Glossary

5G	Fifth Generation (mobile/cellular networks)
IAB	Integrated Access and Backhaul
MT	Mobile Terminal
NGI	Next Generation Internet
NGWB	Next Generation Wireless Backhaul
OSM	Open Street Map
R&D	Research and Development
SDN	Software Defined Network
SDR	Software Defined Radio
SETU	South East Technological University (Coordinating Partner)
SRN	Standard Radio Node
UT	User Terminal
VNF	Virtual Network Function
5G	Fifth Generation (mobile/cellular networks)
IAB	Integrated Access and Backhaul

Co-funded by the Horizon 2020
Framework Programme of the European Union

11 Appendix

Here we detail the release of the open data and the source code. All the code has been released on github, to give it maximum visibility and re-usability. The code uses open source licences.

Since the open data have not been completed, as the link bit-rates are still missing, we temporarily publish the data in the project website[12].

As soon as the experimental part is over, we will update the open data and upload them in an open access platform such as Zenodo.

11.1 Source Geodata

This is the list of the sources of data we used (block 0 in Fig. 1). These are data that we don't republish as they are already open data from public sources.

Step	Source	Content	Area	License
0	link	Pointcloud	Barcelona	CC-BY 4.0
0	link	Raw DSM/DTM	Florence	CC-BY 4.0
0	link	Raw DSM/DTM	Luxemburg	CC0
1	link	Buildings/Roads	Europe	ODbL 1.0

[12]See https://or2an2g.dais.unive.it/software.html

Co-funded by the Horizon 2020 Framework Programme of the European Union

11.2 Released OpenData

The following table details the files names we used in the open data, with reference to the various steps depicted in Fig. 1 and their license.

Block in the flowchart (Fig 1)	File	Content	License
1	{area}.tif	Polished DSM	ODbL 1.0
1	{area}_dtm.tif	DTM	CC-BY 4.0
1	{area}_buildings_mask.tif	Rasterized Buildings	ODbL 1.0
1	{area}_roads_mask.tif	Rasterized Roads	ODbL 1.0
2	cars_raster.tif	Vehicular Demand Model	ODbL 1.0
3	index.csv	gNB placement	CC-BY 4.0
3	viewshed.tif	Area coverage	CC-BY 4.0
3	viewsheds.npy	Individual coverages from each BS	CC-BY 4.0
3	visibility.graphml.gz	visibility graph between BS, demand for each bs	CC-BY 4.0
4	interference.graphml	Interference graph	CC-BY 4.0
4	pathlossMatrix.csv	Colosseum scenarios	CC-BY 4.0
4	delayMatrix.csv	Colosseum scenarios	CC-BY 4.0
6	singletree.graphml	single backhaul graph	CC-BY 4.0
6	multitree.graphml	multi backhaul graph	CC-BY 4.0
7	rapp.graphml	rapp topology	CC-BY 4.0

Co-funded by the Horizon 2020 Framework Programme of the European Union

11.3 Code Repositories

Finally, the pointer to the released code. Every table includes the block in the flowchart of Fig. 1 (if relevant). The first table reports code that we realised in the past, that was updated during the project

Software	Block in the flowchart (Fig 1)	Repository
gNB placement heuristic, now supporting weights.	a	https://github.com/UniVe-NeDS-Lab/TrueBS/

This table reports code that was developed from scratch for the project.

Software	Block in the flowchart (Fig 1)	Repository
Colosseum RF scenario generation	b	https://github.com/UniVe-NeDS-Lab/ColosseumScenarioGenerator
Topology optimization algorithms	c	https://github.com/UniVe-NeDS-Lab/5G-backhaul-OR2AN2G
rApp and non real time RIC		https://github.com/UniVe-NeDS-Lab/IABTopologyOptimizer
Utilities to run OAI on Colosseum		https://github.com/wineslab/OAI-Colosseum

Finally, the code that is part of Colosseum, or OAI we largely contributed to during the project.

Software	Block in the flowchart (Fig 1)	Repository
Extension of the Python library and tool to orchestrate the IAB network		https://github.com/wineslab/iab-manager
OpenAirInterface extension with the O1 interface		https://github.com/UniVe-NeDS-Lab/openairinterface5g/tree/o1_reporting

Co-funded by the Horizon 2020 Framework Programme of the European Union

Co-funded by the Horizon 2020
Framework Programme of the European Union