IncrEase: A Tool for Incremental Planning of Rural Fixed Broadband Wireless Access Networks

Giacomo Bernardi and Mahesh K. Marina
School of Informatics
The University of Edinburgh, UK

Francesco Talamona and Dmitry Rykovanov
EOLO Lab
NGI SpA, Milan, Italy

***Abstract*—The coverage footprint of most Broadband Wireless Access (BWA) networks is not set in stone but varies over its lifetime, usually by expanding incrementally as new transmissions sites are rolled out or by adding capacity as demand changes and increases. We thus propose the *incremental planning* paradigm as an effective strategy for planning and deploying BWA networks, such as those operated by WISPs or community organisations. We believe incremental planning allows to anticipate return of investment and to overcome the limited network infrastructure (e.g., backhaul fibre links) in rural areas. Our `IncrEase` tool elaborates a varied set of operational metrics to guide the operator in identifying the regions that would benefit the most from a network upgrade, automatically suggesting the set of "moves" (e.g., new transmission sites to deploy) that leads to the best long term strategy. We benchmark the computation time of our tool and evaluate it using data from a large-scale ISP network with access to over 8,900 transmission towers.**

I. Introduction

Academia and industry offer a myriad of software tools for wireless network planning. Broadly speaking, their aim is to identify the best locations for transmission towers and/or to plan their interconnection to the rest of the network. However, these tools are often unavailable or unsuitable to communities and small Wireless Internet Service Provides (WISPs), which often resort to an ad-hoc approach towards network planning. Our focus is on the needs of WISPs operating in rural areas, which are faced with the unique challenge of extending their coverage on a tight investment budget in an environment of limited profitability. The key, for such organizations, is to identify the most cost-effective deployment strategy for planning their core network while taking user coverage into consideration.

The design of rural fixed Broadband Wireless Access (BWA) networks is significantly different from mobile broadband networks (e.g., 3G), and its planning process can benefit from the following two observations:

- *Only outdoor propagation is relevant*, as client devices (Customer Premises Equipment, or CPE) are typically installed on rooftops. Taking this notion into account helps in driving down costs (or increases coverage for the same cost) because of the lower path loss as signal does not have to penetrate walls. Several recent market studies, such as [1], concludes that the use of outdoor CPEs is the most cost effective choice to reach the "final third" of the population residing in rural areas.
- Users are fixed: there is rarely any need to support nomadic or mobile services. There is *no need to provide blanket coverage* or overlap in coverage between cells, since no handover is required. The planning process can concentrate only on residential locations where outdoor CPEs will be placed, thereby simplifying the coverage planning aspect of the problem. In comparison, mobile network planning software typically calculates coverage over a grid of equally spaced points, requiring each of them to be above a threshold signal level.

The network model of rural fixed BWA networks almost invariably follows a two-tier model: an access tier composed of point-to-multipoint (PMP) between transmission towers and customers; and a backhaul tier formed of point-to-point (PTP) links typically or exclusively wireless, as dedicated wired links are seldom available in rural areas.

While wireless network planning is traditionally a very active area of the research community, the focus of research literature is mainly on mobile broadband networks and wireless local area networks, such as [2]–[5]. More importantly, the network planning formulation is aimed at *"all-at-once"* deployment, largely based on mathematical optimization methods and meta-heuristics like [6]–[10]. As we argue below, this is an impractical approach to deployment for rural WISPs and community organizations. Limited research has so far focused on the rural domain, which also follows the all-at-once deployment approach. An example is [11], which proposes a four-step scheme to determine the network topology, tower heights, antenna types and radio transmit power levels. The paper provides several interesting design considerations to minimize costs, such as the reduction in tower heights as they often represent the largest investment for a rural ISP. [12] tackles a similar problem: constructing a topology for inter-village rural mesh networks and presents a solution based on greedy approximation algorithms.

We instead advocate the importance of *incremental planning*, a design methodology that guides WISPs – especially those operating in rural scenarios – in planning their growth by extending their coverage. Our approach is based on the following observations:

- Rural deployments are typically *coverage-driven*, rather than capacity-driven[1]. A reason is that low population density of rural areas plays a positive factor in keeping the required capacity of a cell low. Another reason is that ISPs often operate on a tight budget, so they need to anticipate return on investment from the early stage of roll-out. Even more, in an environment of limited profitability such as rural regions, their priority is to reach the largest number of potential customers as early as possible in the deployment phase by focusing on areas where users are clustered (e.g., larger towns and villages lacking broadband access). A similar reasoning is also applicable to community or council networks, where part or all the population of a region needs to be covered.

- Network infrastructure (e.g., fibre links for backhauling) is often limited or unavailable in rural areas, which means that the operator has to roll-out its own (wireless) back-haul tier as it expands, entailing an additional deployment cost.
- Limiting the geographical extent in which the ISP initially operates is effectively a way to keep ancillary operational costs (e.g., customer support) low.

Beyond the initial deployment stage, the network operator can take two actions to extend its business: to increase the network coverage, or to improve it in areas already covered. In either cases, its moves are likely limited by budget and only a small subset of the potential actions can be addressed.

Our aim in this paper is to systematically identify, and suggest to the network operator, a sequence of actions that results in the best long term deployment strategy. Towards this end, we develop an open-source tool called `IncrEase` that enables the incremental planning paradigm in practice. `IncrEase` supports two operational modes, presented in Section II: (1) Targeted Increase, where the operator selects a specific region to be covered as part of network expansion; and (2) Strategy Search, where the tool guides the operator in deciding the deployment order of transmission sites in the near to long-term horizon based on expected profitability. In Section III, we evaluate the tool on a real-world scenario of over 8,000 available towers and benchmark its computation time. To validate the quality of its output, we posed a sample set of planning scenarios to experienced wireless engineers at a large WISP. We then compared their answers with the deployment strategies proposed by `IncrEase`, obtaining similar or better results in most cases.

Even though the term "incremental" has been used elsewhere in the network planning literature (e.g., [13], in the context of capacity planning for fixed lines) this is, to the best of our knowledge, the first paper considering the incremental planning approach for rural fixed BWA networks.

II. The IncrEase Tool

`IncrEase` is an open-source software, implemented as a cross-platform desktop application in Java. It is based on the NASA World Wind Java[2] open-source GIS library and the Neo4J[3] graph database. It allows the network owner to import, model and elaborate customer-related statistics from the BWA network in order to systematically identify the strategy that best improves and extends the network coverage.

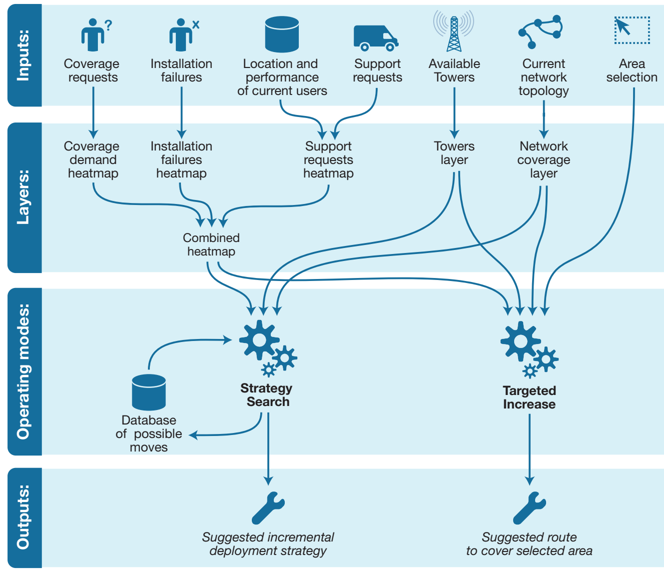

Fig. 1: Example of information flow in the `IncrEase` tool.

An example of such information flow is presented in Fig. 1. A set of XML files containing network statistics is read and parsed. In the current implementation, we consider three sources of data. The first is *coverage demand*: the list of requests for coverage received (e.g., on the ISP website) by potential users living in unserved areas. The second is the set of details regarding those new users that *failed installation* stage due to insufficient coverage. Finally, we also import the *log of support calls* to the WISP helpdesk and the location of existing users. Extra data sets can be imported to capture other influencing factors (e.g., availability of DSL, 3G coverage, demographics, etc.). `IncrEase` elaborates each source of data to form a bi-dimensional array covering the geographical region of interest, with each cell value representing how many "items" (e.g., current users) are contained in the cells region. Cell values are then normalized as a fraction of the cell containing the most items. `IncrEase` visually presents the three 2D arrays on the map as *heatmaps*, and combines them into one as a weighted average, where weights are configured according to each metric's relative importance to the operator. Such combined heatmap provides a glance view of the areas that would benefit the most by a network upgrade. In this sense, *heat* (i.e. high cells values in the 2D array) is an indication of inadequate wireless coverage, which can be removed by a new transmission tower or directive sector. Heatmaps are stored in memory, and can be zoomed, shown or hidden by selecting the appropriate GUI elements.

`IncrEase` can import a a list of additional towers available

[1]The number of base stations is determined either by the size of the region to be covered ("coverage-driven") or by the total traffic that needs to be carried ("capacity-driven")

[2]http://worldwind.arc.nasa.gov/java/

[3]http://neo4j.org/

to be deployed. Such an inventory could include towers that exist already (e.g., available for rent from a tower operator) or suitable locations where new ones could be built. An XML description of the current network topology, including the location and height of each tower and configuration and number of the sector antennas can also be imported into `IncrEase`. Such information is used to generate a "network coverage" layer by performing, with a configurable granularity, line-of-sight calculation from each existing tower and considering azimuth and tilt of existing sectors and a "maximum distance" parameter specifying the admissible range for access-tier wireless links.

We define the following notation for the remainder of the paper:

- T: is the set of all the towers (deployed and available)
- N: a subset of T containing only the towers currently in use in the network topology.
- $h(t)$: is the "total amount of heat" for tower $t \in T$, defined as the sum of the heatmap cell values covered by the tower.
- $c(t)$: is the deployment cost of the tower t.

At launch, the software performs a pre-calculation step to store the memory structures needed at later stages. The most important is G, an "inter-visibility" undirected graph, in which vertices are elements of T. Two towers $t_1, t_2 \in T$ are connected by an edge if they are in line of sight and not farther apart than the maximum allowed distance for point-to-point links. As G is far from a complete graph, we save G on an internal graph database which allows for efficient storage of sparse graphs, and because the LOS calculation is computationally intensive, the graph database is persistently stored on disk, saving time at following launches.

Based on the data infrastructure described, the tool offers two operational modes: (A) Targeted Increase; and (B) Strategy Search. We describe these two modes in the following subsections.

A. Targeted Increase

Heatmaps are a visual aid for the network operator to see the areas that would benefit the most by an improvement in coverage. The *targeted increase* operation mode provides the lightest level of automation available in `IncrEase`, keeping the "human in the loop" by asking the operator to visually select on a map the geographical region where coverage should be improved. `IncrEase` then automatically identifies which is the hottest cell in the region, defined as the one with the highest value in the combined heatmap layer. The application determines the set of closest towers (i.e., 20 in the current default settings) from set $T - N$ that are in line-of-sight of the hottest cell to form the set of candidate locations that will cover the hotspot. Considering multiple source towers allows selection from among a larger number of potential backhaul paths, compared to focusing only on the single tower closest to the hotspot.

The software finds the best way to connect each of those towers to the existing network topology (i.e., the set N) by traversing links in the G graph. The "best" solution is the route that provides the lowest value for the $c(t) - h(t)$ difference for each tower t traversed. In this calculation, we carefully avoid accounting multiple times the "heat" associated with a cell that may be in line-of-sight with different towers, as it would bias the results. We consider heat for such cells only once.

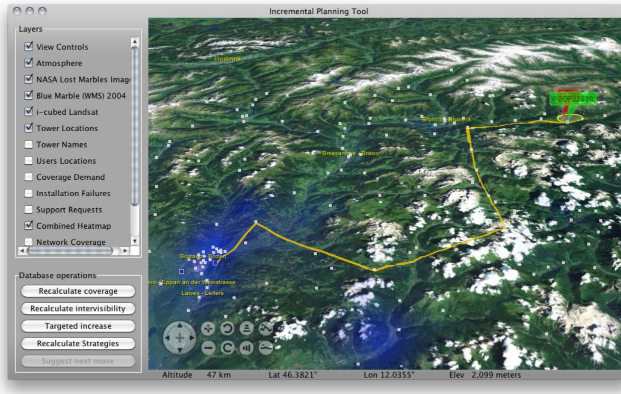

Fig. 2: A screenshot of the `IncrEase` application showing the results of a *Targeted Increase* query.

For pathfinding over the G graph, `IncrEase` uses the A* (A-star) algorithm. A* uses a best-first search, based on a distance-plus-cost heuristic function, to find the least-cost path from an initial node to a goal node. Our implementation has two slight changes with the original A* algorithm. First, it takes as input a set of start nodes (closest towers to the hottest cell in the selected region) and a set of goal nodes instead of a single start/end nodes, as the backhaul path could start from any of the candidate locations and terminate at any of the existing towers. Second, in the G graph, costs are associated with the vertices (i.e., towers) rather than edges, so we consider the cost of an edge (i, j) to be that of the departing node i [14].

A* requires an heuristic function that is the minimal lower bound of possible path cost (e.g., for traveling between two cities, it is distance by straight line), so in our case we need to design an estimate of the best $c(t) - h(t)$ achievable for the rest of path from a given tower to the goal towers. We adopt $(l/d) \cdot c_{min}$ as such heuristic, where l is the straight-line distance between the current tower being analyzed and any one of goal towers, d the maximum distance allowed for point-to-point links (both in km), and c_{min} the minimal $c(t) - h(t)$ of all towers.

Finally, we introduced two modifications to the cost functions presented above. As A* requires non-negative edge costs, we sum an arbitrary constant large positive value to all $c(t) - h(t)$ values. Lastly, to let the user balance the importance of saving money and extending coverage, we allow two variable coefficients c_0 and h_0 and define cost as $c_0 \cdot c(t) - h_0 \cdot h(t)$.

The search result for the best path is presented as a path on the map together with a text indication of which towers to be deployed and their order, as shown in Fig. 2.

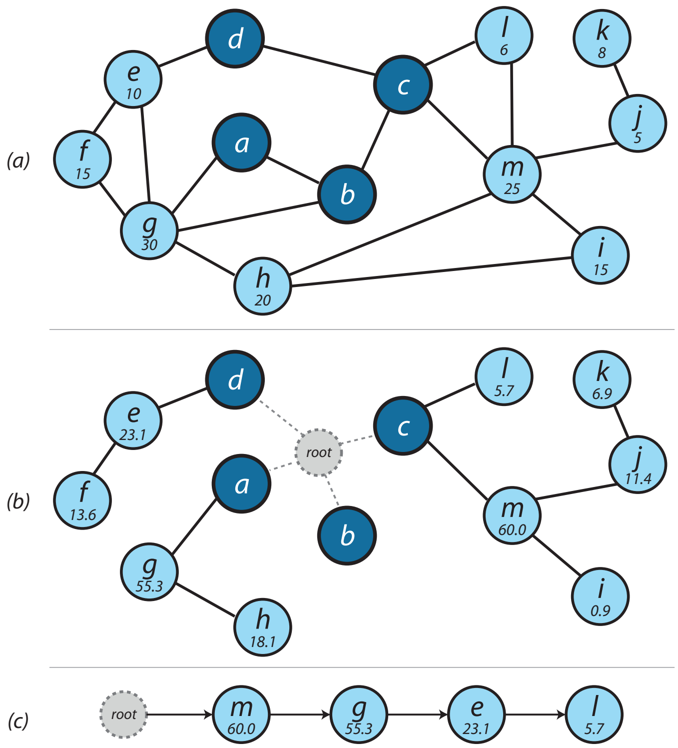

Fig. 3: Data structures used in the *Strategy Search* operation mode. (a) G graph. (b) R graph. (c) L list.

B. Strategy Search

While targeted increase is a semi-automatic mode that requires the operator to select a region, the *strategy search* operational mode identifies and suggests the best network expansion strategy. We assume the network topology to evolve over arbitrary discrete intervals of time (e.g., weeks, months) and the capital investment budget of the WISP to be limited by a discrete parameter that determines how many moves (i.e., tower installations) can be performed in each time interval. The overall aim of the strategy search is then to suggest the "best" actions that the WISP should take during the next interval of time. An obvious practical limitation is the so-called *horizon effect*: as in many artificial-intelligence games, the number of possible states is so large that it is only feasible to search a small portion of all the potential moves within the time horizon. The search algorithm needs to be able to cut down the number of possible strategies to analyse, while limiting the risk of excluding potentially good ones.

Below we describe the strategy search algorithm, which is triggered via the "recalculate strategies" button in the IncrEase user interface.

Step 1. A "multiple-source lowest-cost" path search algorithm is run on the inter-visibility graph G to identify the lowest-cost paths, with costs being $c(t) - h(t)$ as before, from each of the nodes in the set $T - N$ (e.g., the available towers) to any of the nodes in N. The output is a tree R, which intuitively provides the best path from the existing network to each available tower. To generate R, `IncrEase` adds a fictional zero-cost 'root' tower connected to every tower in N and runs Dijkstra's algorithm from such root to each node of the graph G. An example is provided in Fig. 3, where (a) shows the inter-visibility graph G with shaded nodes being those already deployed and (b) the resulting R paths after Dijkstra is run.

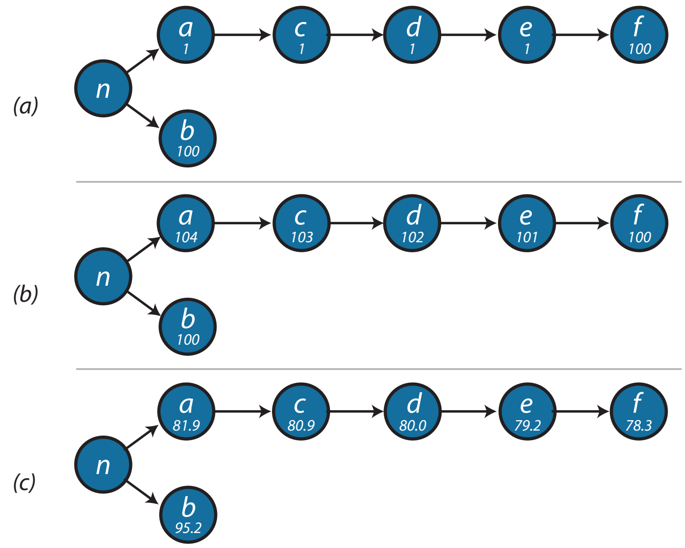

Fig. 4: Net Present Value based score adjustment on an example network (a) in Strategy Search mode: increasing the discount value, C, from 0% (b) to 5% (c) changes the greediness of the selection.

Step 2. Graph R is traversed depth-first starting from the 'root' node and, while doing so, towers are tagged with the score:

$$\frac{h(r) - c(r)}{(1 + C)^{distance(r)}}$$

where $distance(r)$ is the count of newly-deployed towers that have to be traversed to go reach r from the 'root' on R (e.g., towers that can be directly connected to the network have $distance = 0$). As not all towers can be immediately connected to the network, to gain the coverage benefits associated any one of them, others may have to be deployed first to serve as back-hauling relays. To weight future coverage benefits to the present day we bring in the financial concept of Net Present Value (NPV), which applies a constant discount rate C (e.g., 5% being 0.05) to revenues that will happen in the future. An example is provided in Fig. 4(a). Here, two towers *a* and *b* could be installed and directly connected to the existing mast *n*. Nodes *b* and *f* each bring a $h(t) - c(t)$ benefit of 100, while all other towers provide significantly lower benefits. Parameter C is a measure of the greediness of the selection: if we were to pick between *a* and *b* based on the total benefit that they and their descendants could bring, we would decide to install tower *a* as in Fig. 4(b). However, if we increase the value of C to 5%, *b* becomes more attractive (Fig. 4(c)). NPV controls how far it is worth going for installing profitable towers, allowing the network owner to tune the duration of the benefit delay.

Step 3. R is traversed again, this time from leaves to the root. While doing so, we update the score of each node r to the sum of its own score and that of its descendants. Fig. 4(b) and 4(c) actually show the scores obtained after step 3.

Finally, at each click on the "suggest next moves" button of the UI, `IncrEase` asks for the number of moves (towers to be deployed). It then generates a sorted list L (Fig. 3(c)) which includes the towers that could be immediately deployed ordered by decreasing benefit score as calculated at the end of step 3; extracts the top nodes from L; and presents them on the map as results.

III. Evaluation

We have implemented a prototype of `IncrEase` that implements both operation modes described in the previous section. It is being used by NGI SpA, a large Italian ISP operating a fixed 802.11 and 802.16e wireless access network. The service covers Northern Italy, including both metropolitan cities, towns and small rural villages. As ISPs commonly do, NGI has agreements with mobile operators and TV/radio broadcasters in order to acquire space on existing towers, resulting in over 8,900 towers available for immediate installation. The existing network spans over 513 transmission sites, mostly connected over wireless point-to-point (PTP) links of up to 7km in length. For determining coverage, NGI uses a simple line-of-sight (LOS) criteria with a 20km maximum allowed distance, since the access tier operates on the 5.4-5.7Ghz spectrum with outdoor CPEs typically on customers' rooftops. NGI's network is a typical BWA network, so we use it as a representative scenario to benchmark `IncrEase` and demonstrate its usefulness.

A large set of performance metrics has been obtained including coverage requests from prospective clients, details of customers that could not be connected to the network because of insufficient coverage, the log of support request received at NGI's helpdesk, and the geographic location of all current users. We imported this data set in `IncrEase` to generate heatmaps and drive the incremental planning process, as outlined in Fig. 1.

The first step to visually assess areas that would benefit the most by an improvement in coverage is to calculate the current coverage extent. Fig. 5 shows the time taken to calculate and display on a map the area in LOS with any of the towers in the existing network over an area of 273,000 square kilometres. The coverage calculation considers the type, orientation and number of sector antennas installed on each tower. In NGI's network, there are normally 3 sectors per tower, except small towers in mountain areas composed by a single omnidirectional cell, and few critical sites where up to 20 sectors have been installed to add capacity. Benchmarking results, taken on a 2.7Ghz dual-core CPU, validate our implementation by showing strictly quadratic time complexity on the map resolution (in points per latitude/longitude degree), which is expected as the number of points in the region has a power of two relationship with the resolution.

Computation time for building the inter-visibility graph G (Fig. 6) shows linear complexity with the number of available towers in the network, taking a minute to compute G with 8,900 available towers: we get this result by comparing the LOS path of only those tower pairs that are within the

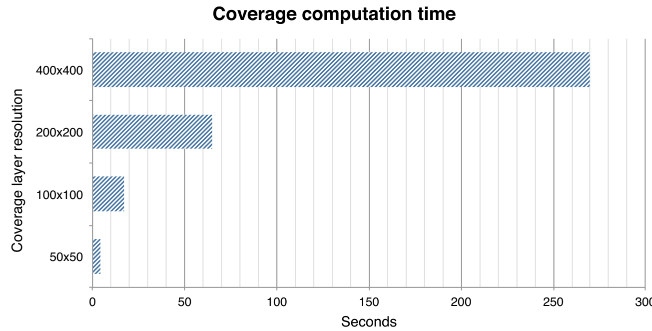

Fig. 5: Benchmarking the computation time of the coverage layer on a network of 513 towers on an area covering the whole of Northern Italy.

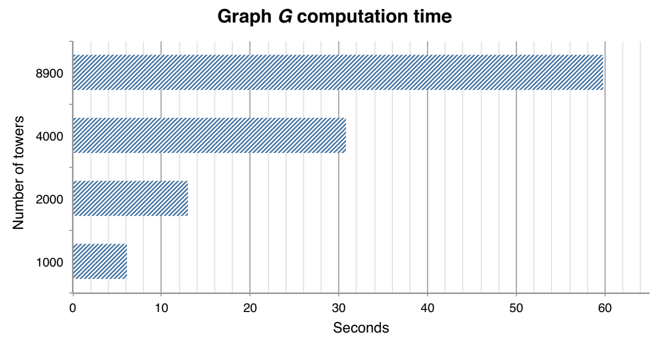

Fig. 6: Time taken to calculate the inter-visibility graph G, depending on the number of available towers.

maximum allowed distance configured for PTP links (7km in our experiments). It is important to note that both the coverage layer and G are persistently saved in the internal graph database and are recomputed only if needed.

A. Targeted Increase

Targeted Increase assists the network operator in finding the "best" strategy to cover a given geographical region and to build a set of suitable back-hauling connections to link it back to the existing network. `IncrEase` provides a suggestion in real time, by exploiting A* heuristics to navigate the inter-visibility graph G. We assess the quality of the proposed solutions with the following comparative study. The locations of available and existing towers were given to five different wireless engineers at NGI, with the request to cover five given small geographical areas and to propose valid backhauling strategies. Engineers could use any tool or technique to work their solution. We then ran the same scenarios on `IncrEase`, and compared the results in terms of number of new towers deployed and total heat removed from the map. Results are presented in Fig. 7: (a) in 4 cases out of 5, the strategy suggested by `IncrEase` requires fewer or equal number of towers to be deployed than solutions provided by wireless engineers; and (b) in all cases better in terms of "quality" in terms of total heat collected. Results demonstrate that Targeted Increase is able to find routes between towers that engineers may overlook. Also, in Scenario 3 `IncrEase` suggests to

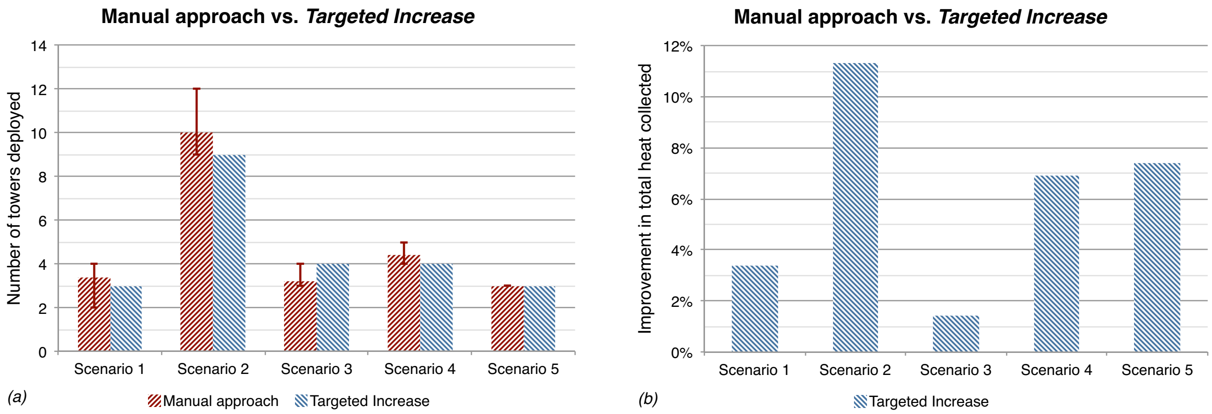

Fig. 7: Comparison, in terms of *(a)* average number of new towers deployed (bars shows best and worst case) and *(b)* total "heat" collected, between manually planning five deployment scenarios and by using the Targeted Increase mode in `IncrEase`.

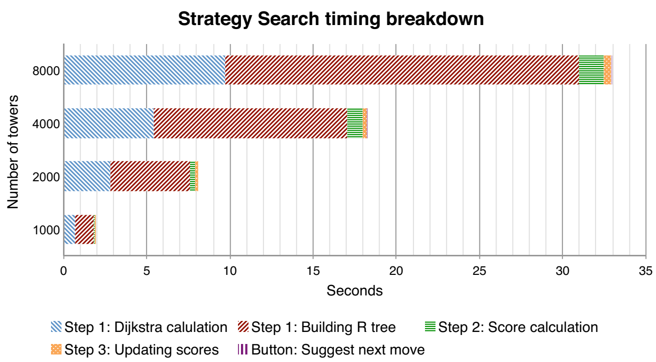

Fig. 8: Profiling of the Strategy Search operation mode, showing the computation time for different steps of the algorithm.

deploy a tower more than what engineers proposed in order to collect more heat. Based on this data, we believe that without the aid of a tool like ours and heatmaps, planning the best backhaul route may be non-obvious, even for a skilled technician, especially when a large number of towers are available.

B. Strategy Search

The Strategy Search operation mode analyses the current network topology and the set of available towers to identify the long-term strategy that leads to the best long-term results. When the "Recalculate Strategies" button is pressed, the three steps presented in Section II-B are executed. Fig. 8 shows the time taken for their completion, which is clearly linear with the number of available towers. Note that timing of Step 1 is split between our modified Dijkstra calculation and the construction of R tree in memory: the latter is relatively more expensive operation as it requires the creation of a new Neo4J graph structure, which is then populated using Dijkstra's output. Finally, at each request to get the next n best actions, the algorithm performs a $O(1)$ operation by selecting the top n elements from ordered list L.

IV. Conclusions

In this paper we have presented `IncrEase`, an open-source tool for *incremental planning* of rural BWA networks. `IncrEase` prototype implementation efficiently handles large data sets and is currently on trial at a BWA operator, which provided the real-world data we used for our evaluation, including a database of over 8,900 real towers. In the future, we intend to improve our tool by incorporating redundancy in core network topology design, time-varying traffic demand patterns and realistic propagation modelling. Finally, as pre-existing transmission towers are often limited in rural regions, `IncrEase` could be enhanced to identify suitable locations (e.g., near roads, electricity available, etc.) for new transmission towers and to incorporate extra attributes in our mast cost model (e.g., power, tower height).

References

[1] "The costs and capabilities of wireless and satellite technologies 2016 snapshot," Analysys Mason, Tech. Rep., October 2010.

[2] E. Amaldi, A. Capone, and F. Malucelli, "Planning UMTS base station location: optimization models with power control and algorithms," *IEEE Transactions on Wireless Communications*, vol. 2, no. 5, Sept. 2003.

[3] S. Bosio, A. Capone, and M. Cesana, "Radio planning of wireless local area networks," *IEEE/ACM Trans. Netw.*, vol. 15, December 2007.

[4] E. Amaldi, A. Capone, and F. Malucelli, "Radio planning and coverage optimization of 3G cellular networks," *Wirel. Netw.*, vol. 14, 2008.

[5] M. St-Hilaire, "Topological planning and design of UMTS mobile networks: a survey," *Wirel. Commun. Mob. Comput.*, vol. 9, July 2009.

[6] R. Whitaker and S. Hurley, "Evolution of planning for wireless communication systems," in *Proc. Hawaii International Conference on System Sciences*, January 2003.

[7] T. Bu, M. C. Chan, and R. Ramjee, "Designing wireless radio access networks for third generation cellular networks," in *Proc. IEEE INFOCOM*, March 2005.

[8] L. Raisanen and R. M. Whitaker, "Comparison and evaluation of multiple objective genetic algorithms for the antenna placement problem," *Mob. Netw. Appl.*, vol. 10, February 2005.

[9] F. Gordejuela-Sanchez, A. Juttner, and J. Zhang, "A multiobjective optimization framework for IEEE 802.16e network design and performance analysis," *IEEE Journal on Selected Areas in Communications*, vol. 27, no. 2, February 2009.

[10] S. Hurley, S. Allen, D. Ryan, and R. Taplin, "Modelling and planning fixed wireless networks," *Wirel. Netw.*, vol. 16, April 2010.

[11] S. Sen and B. Raman, "Long distance wireless mesh network planning: problem formulation and solution," in *Proc. WWW*, 2007.

[12] D. Panigrahi, P. Dutta, S. Jaiswal, K. Naidu, and R. Rastogi, "Minimum cost topology construction for rural wireless mesh networks," in *Proc. IEEE INFOCOM*, April 2008.

[13] J. Charzinski and U. Walter, "Optimized incremental network planning," *Proc. IEEE MMB Conference*, March 2006.

[14] S. S. Skiena, *The Algorithm Design Manual*. Springer-Verlag, 1998.