Cost-Performance Planning of Municipal Wireless Access Networks

Paolo Giacomazzi and Alessandro Poli

Dipartimento di Elettronica e Informazione, Politecnico di Milano
via Ponzio 34/5, 20133 Milano, Italy
{giacomaz,poli}@elet.polimi.it

Abstract. Planning Municipal Wireless Access Networks is a challenging task, since many optimization choices must be taken into account in large metropolitan areas and, in turn, the number of free variables is very large, on the order of several millions, if the position of wireless access points is to be optimized, as well as their connection to the backbone network. This paper formalizes the problem of choosing appropriate access points' locations and to connect the so built wireless access network to the backbone network. We provide an optimization algorithm able to find a solution with a fast heuristic approach, and apply it to a real-world scenario about a 51 km^2 area of Milano (Italy). The proposed heuristic solution is compared with the result of a simulated-annealing-based optimization algorithm and the result is that the cost of the solutions of the heuristic algorithm is larger than that of the optimal solution by a few percents.

Keywords: access networks, wireless, municipal, planning, optimization, cost, performance.

1 Introduction

Municipal Wireless Access Networks (MWAN) planning is a hot topic nowadays. In the world, about 500 projects have been started [1], and the municipalities' interest in the field is rapidly increasing, due to the significant and important services for citizens enabled by these infrastructures. A municipal wireless access network is a viable solution for municipalities to provide the community a broadband network with a relatively low deployment cost, if compared to wired networks [2]. The success of MWANs greatly varies from city to city. Cities where a well-planned risk sharing model between the private and the public sectors has been implemented, typically obtained greater success. Most current literature focuses on economic or social issues of municipal wireless networks. In [3], the role of municipalities as wireless broadband access providers is examined, also presenting current business models. [4] provides an high-level framework, for guiding communities that seek to implement a MWAN, composed of three steps: identification of goals, planning and implementation. [5] analyzes the motivation driving MWAN deployment in three U.S.A. major metropolitan areas, considering standards and technology evolution. A

C. Wang (Ed.): AccessNets, LNICST 6, pp. 46–61, 2009.

variety of technologies and options for wireless access networks is explored in [6], in addition to actual deployment examples. [7] is an attempt to study the feasibility of a MWAN in a specific context, exploring political issues and technological options; a coarse-grained estimation of costs is also reported. The topic of Municipal Wireless Access Networks planning is not faced thoroughly by current literature, mainly due to the extreme difficulty of the problem. This complex task requires the identification of access points' sites and the connection of the wireless access network to a backbone infrastructure. This paper formalizes the problem and provides an optimization algorithm that minimizes the total cost of the infrastructure to be deployed. A real case scenario about a 51 km^2-area of the city center of Milan is studied.

The paper is organized as follows. Section 2 presents the optimization problem which is modelled in Section 3. Section 4 describes the algorithm proposed to solve the optimization problem. Section 5 illustrates the optimization scenario, and empirical results are discussed in Section 6. Section 6 also provides a comparison between the proposed heuristic algorithm and a simulated annealing optimization algorithm. Conclusions are drawn in Section 7.

2 The Optimization Problem

The objective of the methodology presented in this paper is the outdoor coverage of a given metropolitan area with a wireless network at minimum cost. All chosen access points must be connected to a backbone network. In this section, technology requirements and technology resources are presented.

2.1 Technology Requirements

This section describes the formalization of the requirements of the municipal wireless network.

Metropolitan Area Map. The reference area conformation is described by the *metropolitan area map matrix* m_{ij} , identifying, for each latitude-longitude pair, the height in meters of the highest construction in the point, from a fixed reference altitude rh . The actual distance among consecutive indexes (rows or columns) in the matrix, measured in meters, is set by the constant l. As a consequence, in the concrete, each ij-point identifies an $l \times l$ area. The geographical coordinates of the m_{11} position should also be known, so that absolute geographical positions can be mathematically translated into points of the map matrix, and *vice versa*.

Coverage Requirement Map, Coverage Radius Map and Cluster Size Map. For each ij-point in the metropolitan area map, the *coverage requirement map* c_{ij}^{req} identifies whether an outdoor point must be covered by the wireless network ($c_{ij}^{req} = 1$) or not ($c_{ij}^{req} = 0$). A *coverage ratio requirement* parameter cr^{req} defines the minimum acceptable ratio of points actually covered by the wireless network to the total number of points that are required to be covered (i.e. $c_{ij}^{req} = 1$). cr^{req} value is typically near to 1.

Since different zones in the metropolitan area map have different environmental peculiarities, different population densities and users with different behavior, the resulting wireless network should be designed accordingly. Different requirements are taken into account by defining, as further input data, the *coverage radius map* r_{ij}, defining, for each point, the maximum distance, measured in meters on the latitude-longitude plane, from any access point to the points it is allowed and requested to cover, and the *cluster size map* s_{ij}, defining, for each point, the maximum size of a cluster of access points that should share one connection to the backbone network; cluster size greatly affect the wireless network available bandwidth. A *backhaul radius* br, in meters, defines the maximum distance between a pair of access points that can reciprocally communicate, and thus possibly form a wireless mesh *cluster*. Access points in the same cluster share a single connection to the backbone network.

2.2 Technology Resources

The technology requirements for the municipal wireless network can be satisfied by means of the infrastructural resources defined in the following.

Backbone Connection Technologies. A set of available *backbone connection technologies* is given as input. For each connection technology ct_k, the following parameter is specified:

1. Cost c of the connection, as a function of the link length. This cost also includes a fixed cost for the activation of the connection.

Backbone Interconnection Points. A set of all the available interconnection points b_j is considered. An interconnection point can be directly connected to the backbone network, or can require to be connected to another interconnection point. For each interconnection point, a set of properties is defined:

1. Latitude and longitude.
2. Interconnection point type bt_i.
3. Sets of backbone interconnection points that can be reached with the different connection technologies. For each interconnection point, the estimated link length is also known.

The backbone interconnection point type bt_i defines the following characteristics:

1. Maximum number of ingoing connections, for each connection technology ct_k.
2. Whether or not the interconnection point is directly connected to the backbone.
3. Whether or not the interconnection point can establish a connection to another interconnection point using the connection technology ct_k. This option is specified for all connection technologies.
4. Cost c of the interconnection point, as a function of the number of ingoing connections for each technology.
5. Additional cost for using each specific connection technology for the outgoing connection.

6. Estimated average cost of one ingoing connection tc to the interconnection point, for each technology. The heuristic methodology is driven by this cost, which, in turn, is strictly related to the actual cost function of the interconnection point.

Candidate Poles for Access Points. The planning methodology considers a set of all the poles p_j in the metropolitan area that can host a wireless access point. For each pole we know the geographical position and the reachable backbone interconnection points. Altogether, the following properties are available:

1. Latitude and longitude.
2. Sets of backbone interconnection points that the access point installed on the pole can reach with the different connection technologies. For each interconnection point, the estimated link length is also known.
3. Pole type pt_i.

The pole type pt_i defines the following characteristics:

1. Height from the fixed reference altitude rh.
2. Cost c for the use of the pole if it is chosen for hosting an access point. This cost also includes the AP equipment, arrangement, and installation costs.
3. Additional cost for using each specific connection technology.

3 The Optimization Model

3.1 Decision Variables

The following choices should be taken by the optimization methodology:

1. Whether or not a pole is used for hosting an access point. Since there is no need for discrimination in the model, a *pole* used for hosting an access point is also simply referred to as *access point*.
2. Every access point is assigned to a particular *cluster* of access points c_i.
3. For each cluster, one access point is chosen for being connected to the backbone network.
4. Whether or not a backbone interconnection point is used for receiving ingoing connections.
5. For each connected access point, one connection technology is used.
6. For each connected access point, one backbone interconnection point is reached with a *connection*.
7. For each used interconnection point not directly connected to the backbone network, one connection technology towards another interconnection point is used.
8. For each used interconnection point not directly connected to the backbone network, another backbone interconnection is reached with a connection.

Decision variables from 4 to 8 also implicitly define the number of *actually used ingoing connections* for each interconnection point, for all connection technologies.

3.2 Line of Sight and Coverage Evaluation

A three-dimensional point in the map is identified by its three coordinates (x,y,z), where z is the height in meters form the reference altitude rh, while x and y are the latitude and longitude coordinates. x and y can be easily translated into the ij-indexes of the metropolitan area map matrix m_{ij}.

In order to evaluate the mutual visibility of two 3D points, the 3D straight line between the points is projected on the (x,y) plane, and thus all crossed ij-positions in the area map matrix are identified. 3D points are said to be in *sight* when, for every ij-position, the height z of the straight line is above the area map m_{ij} quota, i.e. no obstacles are in between. To cope with rounding and approximation issues, this constraint is relaxed, allowing the existence of a limited overall length occupied by obstacles (obstacles are the points where m_{ij} is above the straight line quota).

For each ij-point in the metropolitan area map, the *coverage map* c_{ij} identifies how many access points are able to reach and cover the specific point. When no poles are used for hosting access points, c_{ij} is equal to zero for all coordinates. For each access point, the coverage area is identified, and then the corresponding c_{ij} values are set. An ij-point is said to be covered by an access point placed on a pole if and only if the 3D point at quota m_{ij} (i.e. it is on the ground) is in sight of the point placed on the top of the pole (its height can be computed by summing the height of the pole to the m_{hk} quota, where h and k are the indexes corresponding to the pole coordinates) and the distance between the two points in the (x,y) plane is no greater than the radius identified by the coverage radius map r_{hk}.

Given an initial state of the coverage map $c_{ij}^{'}$, the use of the additional pole p_c, placed in ij-position, for hosting an access point, causes the coverage map to change to $c_{ij}^{''}$. The *relative coverage gain* cg for the access point p_c is the ratio of the number of hk-points required to be covered ($c_{hk}^{req} = 1$) and covered for the first time by p_c (such that $c_{hk}^{'} = 0$ and $c_{hk}^{''} = 1$), to the number of all mn-points required to be covered ($c_{mn}^{req} = 1$) inside the coverage circle with radius r_{ij}.

3.3 Additional Constraints

In addition to the constraints implicitly derived from the decision variables specification, the following constraints must be enforced:

1. *Coverage ratio* cr should be above the required value cr^{req}. Coverage ratio is the ratio of points required to be covered ($c_{ij}^{req} = 1$) and actually covered by at least one access point (for which $c_{ij} \geq 1$), to the total number of points that are required to be covered (i.e. $c_{ij}^{req} = 1$).

2. All used backbone interconnection points must be either directly connected to the backbone network, or connected to a backbone interconnection point that is connected to the backbone (directly, or through a chain of interconnection points).
3. For each connection going from an access point (or an interconnection point) to an interconnection point, the use of the same connection technology on both sides is required.
4. For each used interconnection point, the sum of the ingoing connections, for each connection technology, must not exceed the maximum number defined by the interconnection point type.
5. For each cluster of access points c_a, the number of access points in the cluster must not exceed the requirement of maximum cluster size. The maximum cluster size for cluster c_a is the minimum of the s_{ij} values in all the *ij*-positions where an access point's pole in cluster c_a is located.
6. Each access point not connected to a backbone interconnection point in a cluster c_a must be in sight, and no farther than the backhaul radius br in the (x,y) plane, of an access point, in the same cluster c_a, connected to a backbone interconnection point (directly, or through a chain of access points in sight of the same cluster). A pair of access points is said to be in sight when the two 3D points placed on the top of the poles are in sight (see Section 3.2) and no farther than the backhaul radius br in the (x,y) plane.

3.4 Objective Function

The objective function to be minimized is the total cost of the municipal wireless network. The total cost can be computed by adding the following costs:

1. Sum over each used interconnection point of its cost (a function of the number of ingoing connections for each technology plus the additional cost for using each specific connection technology for the outgoing connection).
2. Sum over each used interconnection point not directly connected to the backbone of the connection cost.
3. Sum over each pole used for hosting an access point of its cost.
4. Sum over each pole used for hosting an access point connected to the backbone of the connection cost.

4 The Optimization Algorithm

Our heuristic methodology for the planning of a municipal wireless network comprises two phases: (a) wireless *coverage* of the designated area, (b) *connection* of access points *to the backbone* network. Both phases are independent from the other; nevertheless, the first phase, while exploring its solution space, also takes into account the impact of its choices over the objective function because of the second phase.

4.1 Wireless Coverage

The wireless coverage algorithm aims at finding a set of poles where to install access points, such that coverage requirements are satisfied. Access points must be grouped in clusters where the wireless mesh network communication is feasible. Even though the minimization of overlaps among the wireless coverage of different access points is not a requirement, the proposed algorithm also implicitly tries to minimize such zones (i.e. the points where $c_{ij} > 1$). The algorithm uses the following additional parameters:

1. A *pole preference ratio* value pr, for each pole type pt_i. This value is set to 1 for standard poles, while it is smaller than 1 for poles that should be preferred during the wireless coverage phase and greater than 1 for poles to be penalized. pr values should be empirically determined in such a way to minimize the total cost of the infrastructure, privileging the choice of poles whose use is cost-beneficial.
2. *Minimum coverage gain* mcg for an additional access point. This parameter forces the coverage procedure to avoid the use of access points whose additional contribution to the global coverage is slight. This value should be carefully chosen, since with high mcg values the coverage ratio constraint could not be satisfied.

The algorithm also uses the following additional variable:

1. *Used poles stack* ps, a stack used to keep track of previously chosen poles.

The algorithm follows these steps:

1. A pole p_c is randomly chosen among those placed in a point required to be covered (such that $c_{ij}^{req} = 1$) and not yet covered ($c_{ij} = 0$). Current cluster is chosen among already existing *not full* clusters; the cluster containing the nearest pole to p_c, and in sight with p_c, within a maximum distance $\left(2 \cdot r_{ij}\right) - \varepsilon$, is selected. If no clusters satisfy all requirements and if current cluster is not empty (i.e. at least one access point has been assigned to it), then create a new empty cluster. A cluster is said to be *full* if the number of access points in it is not smaller than the minimum s_{ij} value, evaluated among all the *ij*-positions of the access points in the cluster and the currently chosen access point p_c.
2. The chosen current pole p_c is used for hosting an access point. As a consequence, the coverage map c_{ij} is updated, according to the procedure described in Section 3.2.
3. A new cluster c_h is created (h is incremented by one) if current cluster is *full.*
4. Chosen pole p_c is assigned to current cluster c_h and pushed into the ps stack.
5. Being *i* and *j* the positions in the metropolitan area map of current pole p_c, next pole is chosen among those which are in sight with current pole p_c, not placed in

an already covered area, and within a maximum distance $\left(2\cdot r_{ij}\right)-\varepsilon$ in the (x,y) plane, where r_{ij} is the coverage radius requested for the specific area, and ε is a short distance used to force a small overlap of the coverage circles, such that their union does not leave small uncovered pieces of land on road sides. The methodology supposes that, in all areas, the backhaul radius br is greater than $\left(2\cdot r_{ij}\right)-\varepsilon$. Poles not reaching the minimum coverage gain (such that $cg < mcg$) are discarded and not taken into account. Identified candidate next poles are ranked by their decreasing distance from p_c divided by their preference ratio pr. The first pole is then assumed to be the following current pole p_c.

6. If not any pole has been identified by the previous step:
   a. If ps stack is not empty, next current pole p_c is identified by popping a pole from the ps stack. A new empty cluster is created. Then go to step 5.
   b. If ps stack is empty create a new empty cluster. Then go to step 1.
7. If coverage ratio is above the required threshold ($cr \geq cr^{req}$) then stop. Else go to step 2.

4.2 Connection to the Backbone Network

This phase goal is the identification, among the access points in a cluster, of the access point designated to be connected to an interconnection point, and the identification of the interconnection points to be used.

The algorithm uses the following additional parameters:

1. *Security fulfillment ratio threshold* ft for each connection technology ct_i. This ratio threshold (spanning from 0 to 1) is used to avoid, in the heuristic algorithm, states when a pole or an interconnection point can not reach any interconnection point with available ingoing connections. Its use will be better explained in the following.
2. *Average number of ingoing connections* $bt_i.acn\left(ct_j\right)$ for each backbone interconnection point type bt_i, for all its available connection technologies ct_j.

These parameters guide the cost estimation of the heuristic, and should be set with values reflecting the actual average number of ingoing connection for each interconnection point type. The estimation of these parameters can be performed by evaluating the solutions found by the whole heuristic algorithm.

The algorithm also uses the following additional variable:

1. *Average overall cost per termination* otc, for each backbone interconnection point, and for each connection technology ct_j, referring to the estimated overall cost for one ingoing connection. This cost includes all the costs for linking one ingoing connection till the backbone, and thus accounts for the entire path to the backbone network. In the following, we will refer to the average overall cost per

termination with technology ct_j for the backbone interconnection point b_i as $b_i.ct_j.otc$.

The algorithm operates as follows:

1. For all the interconnection points b_i directly connected to the backbone, for each connection technology ct_j, the average overall cost per termination otc is set as the average cost tc of one ingoing connection, defined by its type bt_k. At this stage, all backbone interconnection points are set as *not used* for receiving ingoing connections.
2. For all the interconnection points b_i not directly connected to the backbone and not already set as *used*, for each connection technology ct_j, all the available interconnection points that can be reached, that are either directly connected to the backbone or that have an outgoing connection already set, and such that the number of *actually used ingoing connections* is not above the maximum number of available ingoing connections multiplied by the *security fulfillment ratio threshold* ft, are explored. The minimum cost option is chosen and, consequently, destination interconnection point and related connection technology ct_j are set for current interconnection point. If not any option is found, the security fulfillment requirement is relaxed, and the evaluation is performed again. Costs are evaluated as the average overall cost per termination otc of one ingoing connection of the destination interconnection points, plus the cost $ct_j.c$ of the connections as function of link lengths. The cost of the chosen option is then divided by the average number of ingoing connections $bt_k.ct_j.acn$ and added to the average costs $ct_j.tc$, for all incoming connection technologies; the resulting value is stored as the average overall costs per termination $ct_j.otc$.
3. In the previous step, destination interconnection points and related connection technologies ct_j have been set for a part of all backbone interconnection points, and average overall costs per termination otc, for all connection technologies, have been computed. Now, for each backbone interconnection point, a larger set of reachable interconnection points, that are either directly connected to the backbone or that have an outgoing connection already set, is available. Step 2 is consequently iteratively executed until not any average overall cost per termination otc does change, i.e. the minimum cost path from any backbone interconnection point to the backbone has been identified.
4. Randomly choose a *cluster* of access points c_i not already connected to the backbone network. All the available interconnection points that can be reached with all the connection technologies available from cluster's access points, and that are either directly connected to the backbone or have an outgoing connection already set, and such that the number of *actually used ingoing connections* is not above the maximum number of available ingoing connections multiplied by the

security fulfillment ratio threshold ft, are explored. The minimum cost option is chosen and, consequently, destination interconnection point and related connection technology ct_j are set for corresponding access point and cluster; the access point and the cluster are then marked as connected to the backbone. If not any option is found, the security fulfillment requirement is relaxed, and the evaluation is performed again. Costs are evaluated as the average overall cost per termination otc of one ingoing connection of the destination interconnection points, plus the cost $ct_j.c$ of the connections as function of link lengths. Access point's destination interconnection point and, iteratively, all the interconnection points met along the path to the backbone network, are consequently set as *used*; concurrently, the number of *actually used ingoing connections* for all met interconnection points is updated; if the assignment caused a *security fulfillment threshold* to be passed for any backbone interconnection point, a smaller set of reachable interconnection points is now available. Step 2 (only) is consequently iteratively executed until not any average overall cost per termination otc does change.

5. If still exist at least one *cluster* of access points c_i not already connected to the backbone network then go to step 4.
 Else stop.

5 The Optimization Scenario

The algorithm described in Section 4 has been applied on a large area of the city of Milan. The area of interest is a rectangular 51 km^2 zone of the city center. Since the information about the height of the buildings is not extensively available, the *metropolitan area map matrix* only discriminates among points where a street or a building is there. The used *coverage requirement map* is such that all the streets and all the parks in the area are required to be covered by the wireless network. The map is shown in Fig. 1; the white points are the areas to be covered ($c_{ij}^{req} = 1$). A *coverage ratio requirement* parameter of 0.97 has been used in the performed optimizations.

The *cluster size map*, defining, for each point, the maximum size of a cluster of access points, has been set to 3 for all points. The *backhaul radius* has been safely considered equal to 200 m.

The available *backbone connection technologies* are two classes of *Optical Fiber*, respectively provided by a *Third Party* company or by the *Municipality of Milan*, *Power Line Communication* (PLC), and *Direct Connection* with a local area network link. A direct connection is a connection that can be reached by an access point with a local area link, for example when the access point is located on the same building as the interconnection point. The cost of the Third Party Optical Fiber is 7 €/m, while the Optical Fiber of the Municipality of Milan is provided for free. A Power Line Communication link has a fixed cost of 95 €, accounting for the PLC client installation, and a cost of 480 € every 150 m, if the total length of the link is above 150 m, and a signal repeater is required. A Direct Connection has a fixed cost of 500 €.

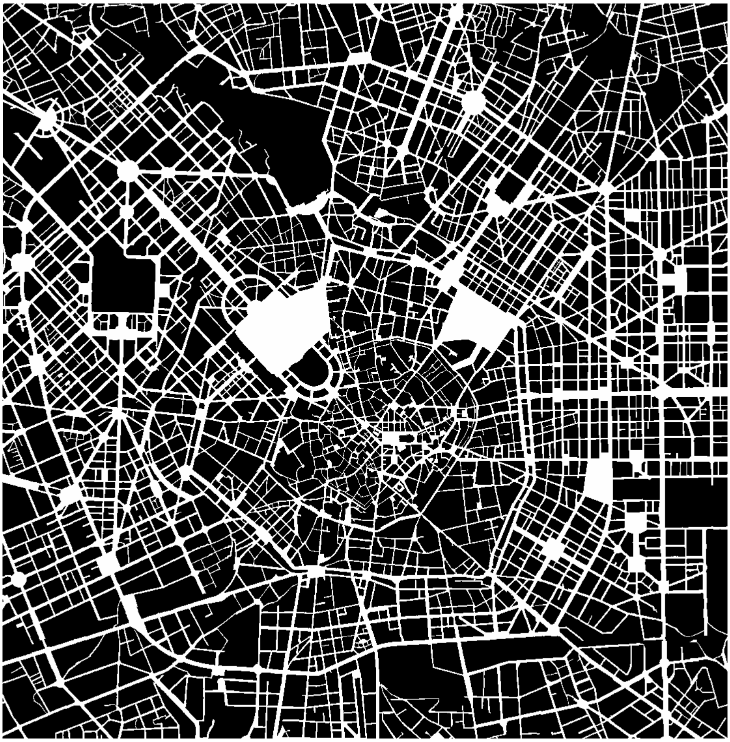

Fig. 1. Coverage requirement map of the considered zone, representing a 51 km^2 area of the city center of Milan

There are 4 types of backbone interconnection points:

1. *Third Party Cabinets* owned by a third party company, 8 in the considered area, directly connected to the backbone, not any reachable interconnection point, can terminate an unlimited number of Third Party Optical Fibers, and 3 PLC connections. The cost is given by the Router Chassis (8,864 € each) capable of hosting up to 6 2xEth 100base FX adapter modules (6,332 € each) able to terminate up 2 optical fibers, and by the PLC Head End (465 €), if required.
2. *CampusII Buildings* owned by the Municipality of Milan, 319 in the considered area, directly connected to the backbone, not any reachable interconnection point, can terminate an unlimited number of Third Party Optical Fibers, 3 PLC connections, and one Direct Connection. The cost is given by the Router Chassis (8,864 € each) capable of hosting up to 6 2xEth 100base FX adapter modules (6,332 € each) able to terminate up 2 optical fibers, and by the PLC Head End (465 €), if required.
3. *Municipal Traffic Control Facilities* sites owned by the Municipality of Milan, 582 in the considered area, directly connected to the backbone, not any reachable interconnection point, can terminate an unlimited number of Municipal Optical Fibers, and one Direct Connection. The cost is given by the Router Chassis (8,864

€ each) capable of hosting up to 6 2xEth 100base FX adapter modules (6,332 € each) able to terminate up 2 optical fibers.

4. *Medium to Low-Voltage Electric Cabins* owned by the Local Electric Grid Company of Milan, one per electric cell, 778 in the considered area, not directly connected to the backbone, can reach any Third Party Cabinets and any CampusII building with Third Party Optical Fibers, can terminate 3 PLC connections. The cost is given by the Power Line Communication Head End (465 €), used to terminate no more than 3 connections. The number and the position of the Medium to Low-Voltage Electric Cabins, at current stage, are not exactly known, and have been estimated.

We consider 3 types of poles:

1. *Streetlamp poles* owned by the Local Electric Grid Company of Milan, 85,285 in the considered area, can reach any Third Party Cabinet and any CampusII building with Third Party Optical Fibers, or the nearest Medium to Low-Voltage Electric Cabin with PLC technology. The cost is given by the access point device and its installation (2,900 €) plus an additional pole adaptation cost (3,200 €) if using Third Party Optical Fibers as connection technology. The number and the position of streetlamp poles, at current stage, are not exactly known, and have been estimated.
2. *CampusII poles* placed on CampusII buildings, 319 in the considered area, owned by the Municipality of Milan, can reach the corresponding CampusII buildings with a Direct Connection. The cost is given by the access point device (1,900 €).
3. *Traffic Lights and CCTV poles* owned by the Municipality of Milan, 909 in the considered area, can reach any Municipal Traffic Control Facility with Municipal Optical Fibers, and the nearest Medium to Low-Voltage Electric Cabin with PLC technology; every Traffic Lights pole can also reach one specific Municipal Traffic Control Facility with a Direct Connection. The cost is given by the access point device and its installation (2,400 €).

The above described available technology resources are such that can be mixed together only in ways that do not cause bottlenecks to appear, since the bandwidth availability sufficiently increases along the paths from the Access Points to the backbone network. This is an implicit requirement of the optimization problem that must be ensured by appropriate input data.

6 Optimization Results

Optimizations have been performed by an ad-hoc software tool, called MUWI (MUnicipal WIreless optimizer), implementing the algorithm presented in this paper. Results are about the scenario described in Section 5, considering a set of different values for the *coverage radius maps*, defining, for each point, the maximum distance reached by an access point. The considered radiuses, the same for all the points in the same map, range from 40 m to 90 m. Optimizations have been carried out in 5 different variations of the considered scenario:

1. Scenario **THIRD PARTY**: availability of *Third Party Cabinets* interconnection points, and *Streetlamp poles* only.

2. Scenario **CAMP+TC**: availability of *Third Party Cabinets*, *CampusII Buildings*, *Municipal Traffic Control Facilities* interconnection points, and *Streetlamp poles*, *CampusII poles* and *Traffic Lights and CCTV poles* only.
3. Scenario **TC+PLC**: availability of *Third Party Cabinets*, *Municipal Traffic Control Facilities*, *Medium to Low-Voltage Electric Cabins* interconnection points, and *Streetlamp poles* and *Traffic Lights and CCTV poles* only.
4. Scenario **CAMP+PLC**: availability of *Third Party Cabinets*, *CampusII Buildings*, *Medium to Low-Voltage Electric Cabins* interconnection points, and *Streetlamp poles* and *CampusII poles* only.
5. Scenario **CAMP+TC+PLC**: availability of *Third Party Cabinets*, *CampusII Buildings*, *Municipal Traffic Control Facilities*, *Medium to Low-Voltage Electric Cabins* interconnection points, and *Streetlamp poles*, *CampusII poles* and *Traffic Lights and CCTV poles* (the full scenario).

The execution of the algorithm for one optimization took about 3 hours on one CPU/core of a machine equipped with two Intel Xeon E5440 quad-core 2,83 GHz. The total costs of the municipal wireless infrastructures generated by the optimization algorithm are shown in Fig. 2.

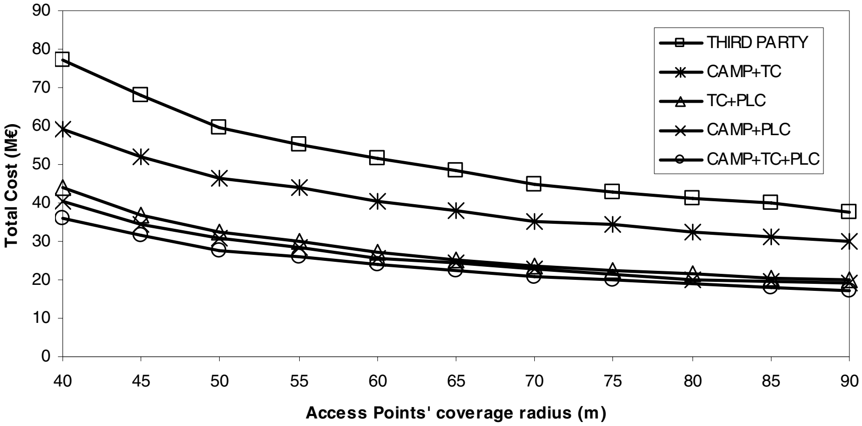

Fig. 2. Total cost of the municipal wireless network generated by the optimization algorithm in different scenario variations

Total cost doubles when coverage radius decreases from 90 m to 40 m, in all scenario variations. However, smaller radiuses allow offering higher bandwidth capacity per unit of area, depending on the specific wireless technology. The use of the third party company's infrastructures only, brings to the highest cost network, equal to 59.4 M€ with a coverage radius of 50 m. When preexisting municipality-owned infrastructures are taken into account (CAMP+TC scenario), a cost reduction of about 20% can be achieved in all coverage radius options. The Power Line Communication technology has the highest impact over the infrastructure cost; an additional 40% (CAMP+TC+PLC scenario compared to CAMP+TC scenario) can be gained by installing PLC Head End on designated Medium to Low-Voltage Electric Cabins and PLC Clients on well-chosen city poles.

The cost shares of the devices in three different scenario variations, using a coverage radius of 50 m, are reported in Fig. 3. The use of municipality-owned infrastructures (CAMP+TC and CAMP+TC+PLC scenarios) allows a cost reduction because of a lower use of Third Party Optical Fibers, from a 32% share to 7% and 5%. Access points cost is similar, in absolute value, in the three scenario variations. The use of the PLC technology, having an impact of 2% only on costs, allows an additional cost reduction mainly affecting APs arrangement, APs installation, router chassis, and router modules, because of a smaller amount of optical fibers to be connected to access points and to be terminated in routers.

The solution found for the full scenario (CAMP+TC+PLC), using a coverage radius of 50 m, has a total cost of 27.7M€ and uses 6949 poles/access points, 6261 of which are *Streetlamp poles*, 319 *CampusII poles*, and 369 *Traffic Lights and CCTV poles*. A total number of 2558 poles are infrastructured by means of the PLC technology.

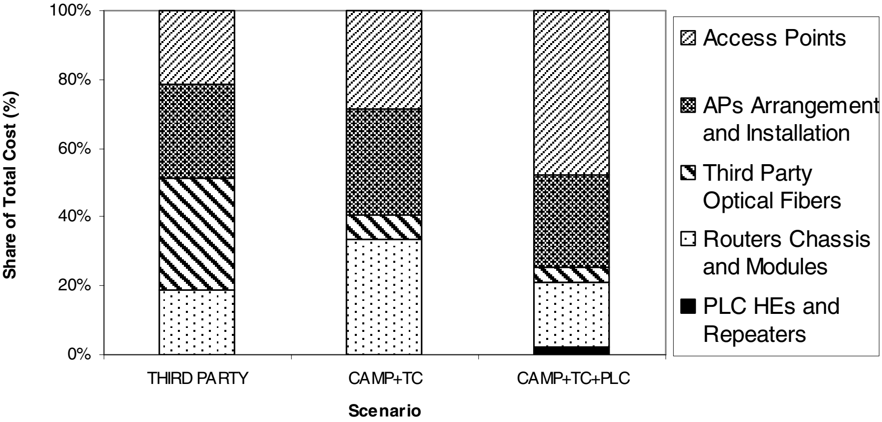

Fig. 3. Cost shares of the solutions obtained with a coverage radius of 50 m in three different scenario variations

6.1 Assessment through Simulated Annealing Optimization

The assessment of the heuristic algorithm proposed in Section 4 has been performed by means of comparisons with solutions found by a simulated annealing procedure. Simulated Annealing [8] is a probabilistic meta-algorithm for optimization problems, able to find good approximations of the global optimum in large search spaces. Simulated Annealing executes a random sequence of moves; every move is evaluated and is executed with a transition probability $P(\delta E, T)$, where δE is the variation of the objective function value caused by the move and T is the temperature of the system at a given moment.

In our implementation, if the move brings to a better solution, the move is always accepted (transition probability equal to 1), while it is $P(\delta E, T) = e^{\frac{-\delta E}{T}}$ if the move worsens the solution, according to the original formulation [8]. T decreases at every step until it reaches zero; we use exponential decreases with slow cooling ratios (0.99 or higher), and stop the Simulated Annealing algorithm when 1 million consecutive moves have not been accepted. The Simulated Annealing algorithm randomly chooses among a set of simple moves. A move transition probability is evaluated only if it brings to a feasible solution, i.e. no constraint is violated. Two groups of moves have been implemented: moves affecting the wireless coverage and moves affecting the connection to the backbone. Moves are described in the following.

Wireless Coverage Moves. This group of moves affects the selection of poles and, consequently, the wireless coverage.

1. *Remove AP*: remove an access point if it is not required for the communication among the access points in a cluster.
2. *Add Mesh AP*: add an access point to a pole and then add it to an already existing not full cluster in sight and no farther than the backhaul radius br .
3. *Add Infrastructured AP*: add an access point to a pole, add it to a new cluster, and then connect it to a random reachable interconnection point with free ingoing connections.

Backbone Connection Moves. This group of moves affects the connection to the backbone network.

1. *AP from IP to IP*: disconnect an access point connected to a backbone interconnection point and then connect it to a random reachable interconnection point with free ingoing connections. Update the set of used backbone interconnection points, if needed.
2. *AP from Mesh to IP*: remove an access point from a cluster, add it to a new cluster, and then connect it to a random reachable interconnection point with free ingoing connections. Update the set of used backbone interconnection points, if needed.
3. *AP from IP to Mesh*: disconnect an access point connected to a backbone interconnection point, and then merge its cluster's access points to an already existing not full cluster in sight and no farther than the backhaul radius br . Update the set of used backbone interconnection points, if needed.
4. *AP from Mesh to Mesh*: remove an access point from a cluster, and then add the access point to an already existing not full cluster in sight and no farther than the backhaul radius br .
5. *IP from IP to IP*: disconnect a backbone interconnection point connected to another backbone interconnection point and then connect it to a random reachable interconnection point with free ingoing connections. Update the set of used backbone interconnection points, if needed.

Simulated Annealing Results. Since the solution space of our optimization scenario is huge, the simulated annealing optimization would have required very slow cooling ratios and long time in order to converge towards a good approximation of the global optimum. As a consequence, our comparison has been performed on a smaller city area of 0.9 km^2 with reduced solution space. The reduction of the solution space required a decrease of the number of poles available for hosting access points (257 in the considered area).

The solution found by the optimization algorithm described in Section 4 has been used as starting solution for the Simulated Annealing algorithm. We tried a large set of different initial temperatures T and different cooling ratios. The minimum cost solution has been found with an initial temperature equal to 1,000 and a cooling ratio equal to 0.9999999. The cost of the best identified solution is 102,175 €, and execution took 275 minutes on our reference machine. This solution must be compared to the ones identified by the heuristic algorithm proposed in Section 4; the execution of 100 iterations of our heuristic algorithm with different random seeds required a total time of 20 minutes on our reference machine, and allowed to identify a solution with a cost equal to 105,625 €, i.e. about 3% more expensive than the solution found by the best parameterized Simulated Annealing algorithm.

7 Conclusions

We have proposed a new planning algorithm for Municipal Wireless Access Networks. The algorithm is based on a heuristic optimization procedure aiming at minimizing solutions cost. A heuristic approach has been followed because of the great complexity of the problem, which involves millions of decision variables when a large metropolitan area has to be taken into account all together. The algorithm has been implemented in a software tool and has been applied on a scenario about a 51 km^2 area of the city of Milan. Different scenario variations have been considered; results show that the reuse of preexisting municipality-owned infrastructures allows cost reduction in the order of 20%. An additional 40% cost reduction can be gained by the extensive use of the Power Line Communication technology.

The proposed heuristic algorithm has been assessed through Simulated Annealing optimization, and results show that our algorithm is able to provide solutions that are only about 3% far away from the approximation of the global optimum reached by Simulated Annealing algorithms.

References

1. MuniWireless, `http://www.muniwireless.com`
2. Gibbons, J., Ruth, S.: Municipal Wi-Fi: big wave or wipeout? IEEE Internet Computing 10(3), 66–71 (2006)
3. Tapia, A., Maitland, C., Stone, M.: Making IT work for municipalities: Building municipal wireless networks. Government Information Quarterly, Issues in Wireless Broadband 23(3-4), 359–380 (2006)
4. Mandviwalla, M., et al.: Municipal Broadband Wireless Networks. Communications of the ACM 51(2), 72–80 (2008)
5. Reinwand, C.C.: Municipal Broadband - The Evolution of Next Generation Wireless Networks. In: IEEE Radio and Wireless Symposium, pp. 273–276 (2007)
6. Sirbu, M., Lehr, W., Gillett, S.: Evolving wireless access technologies for municipal broadband. Government Information Quarterly, Issues in Wireless Broadband 23(3-4), 480–502 (2006)
7. Abelém, A.J.G., Stanton, M.A.: Alternatives for community metropolitan networks for the major cities of the Amazon region of Brazil: the case of Belem. In: 2nd International Conference on Broadband Networks, vol. 2, pp. 1492–1498 (2005)
8. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by Simulated Annealing. Science 220(4598), 671–680 (1983)