

Routing and Network Design for HEAnet

Deepak Mehta, Barry O’Sullivan, and Luis Quesada
4C, University College Cork, Cork, Ireland

David Payne, Linda Doyle, and Marco Ruffini
CTVR, University of Dublin, Trinity College Dublin, Ireland

Abstract—HEAnet is Ireland’s national education and research network. In this paper we address the problem of deploying long-reach passive optical access points and core nodes for this network. We decompose the deployment problem into two phases. First we address the problem of locating a number of core nodes and local exchanges to cover the pool of HEAnet customers. Then, we compute the actual route that the cable needs to follow to reach every customer taking into account that there is an upper bound on the length of the path from the core node to the customer. The first phase of the problem is solved as a mixed integer program (MIP). We empirically study this phase in detail. The second phase is also solved as a mixed integer program. This phase is unlikely to be suitable for solving with MIP when the problem size increases as the formulation is quadratic in terms of the number of nodes involved in the network, so we also explore an alternative constraint programming mode, which is more compact.

I. INTRODUCTION

In this paper we study the challenge of developing a combinatorial optimisation approach to designing a long-reach passive optical network (LR-PON) [4] infrastructure for Ireland’s HEAnet network, a core infrastructure for the country’s education and research systems.¹ Our approach to solving this problem comprises two phases.

The first phase involves locating a number of metro nodes and local-exchanges to cover the pool of HEAnet customers. We solve this phase using mixed integer programming (MIP). Due to the size of the problem, it is further decomposed into two subproblems that are solved interleavedly until an optimal solution is found. The second phase involves computing the actual route that the cable needs to follow from a metro node to a customer taking into account that there is an upper bound on the length of the path from the core node to the customer. This phase ensures that each user connects with an optical fibre to a local exchange, which supports a number of other users following a tree topology. Each local-exchange connects to two metro nodes, one for primary service and one for backup purpose. Figure 1 shows the location of 3906 users, 1121 local-exchanges and 59 candidate locations for metro nodes. We solve this phase also using mixed integer programming. Solving this phase of the problem using MIP would be impractical for big networks because of the size of the formulation that would be required.

¹<http://www.heanet.ie>

Therefore, for the second phase we also explore a pure constraint programming approach using global constraints. The overall objective is to cover every user with a PON, using the least amount of fibre. The fibre length is calculated considering the Euclidean distance between two points and then multiplying by a routing factor of 1.4. In this paper we study both phases of the problem in detail, presenting detailed empirical result.

II. DEPLOYMENT OF METRO NODES AND EXCHANGE SITES

Table I presents a mixed integer programming model for the deployment of metro nodes and exchange sites. In the model we use the following notation: C is a set of customers. E is a set of exchange sites. M is a set of metro nodes. d^{ce} is the maximum allowed distance between a customer and its exchange site. d^{em} is the maximum allowed distance between an exchange site and a metro node (this will include routing factor). s^e is the maximum number of exchange sites that can be selected. s^m is the maximum number of metro nodes that can be selected.

We say that a customer is covered by an exchange site if the distance between them is less than d^{ce} . We say that an exchange site is covered by a metro node if the distance between them is less than d^{em} . The problem is to find a set of exchange sites, denoted by S_E , and a set of metro nodes, denoted by S_M , such that $|S_E| = s^e$, $|S_M| = s^m$, each customer $i \in C$ is covered by 1 exchange site in S_E , each exchange site $j \in S_E$ is covered by 2 metro nodes in S_M , and the sum of the lengths of optical fibres required to connect each customer to its exchange site and each selected exchange site to its metro nodes is minimised.

In the model we also use $E^c(i)$ to denote the set of exchange sites that are within distance d^{ce} from a customer i , $C^e(i)$ to denote the set of customers that are within distance d^{ce} from an exchange site i , $M^e(i)$ to denote the set of metro nodes that are within distance d^{em} from an exchange site i , $E^m(i)$ to denote the set of exchange sites that are within distance d^{em} from a metro node i , $M^c(i)$ to denote the set of metro nodes that cover a customer i (i.e., $M^c(i) = \bigcup_{j \in E^c(i)} M^e(j)$).

The model described in Table I does not scale for our problem instance. Therefore, we decompose the original

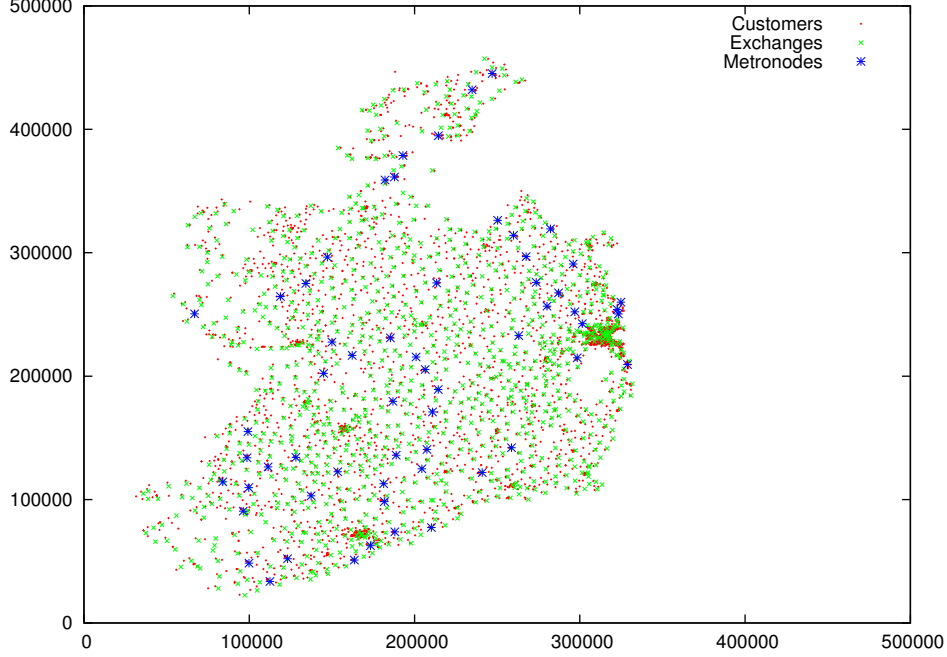


Figure 1. Locations of customers, exchanges and metro nodes.

Table I
MIP MODEL FOR THE DEPLOYMENT OF METRO NODES AND EXCHANGE SITES

Variables	
x_j^e	A Boolean variable that denotes whether exchange site j is selected or not.
x_j^m	A Boolean variable that denotes whether metro node j is selected or not.
y_{ij}^{ce}	A Boolean variable that denotes whether a customers i is connected to an exchange site j .
y_{ij}^{em}	A Boolean variable that denotes whether an exchange site i is connected a metro node j .
Constraints	
$\forall i \in C : \sum_{j \in E^e(i)} y_{ij}^{ce} \geq 1$	Each customer should be connected to one exchange site.
$\forall j \in E \forall i \in C^e(i) : x_j^e \geq y_{ij}^{ce}$	An exchange site is selected if a customer within its reach is connected to it. Another way to impose this constraint is $\forall j \in E : C^e(j) \times x_j^e \geq \sum_{i \in C^e(j)} y_{ij}^{ce}$. However, the LP bound would be weaker.
$\sum_{i \in E} x_j^e = s^e$	The number of selected exchange sites should be equal to s^e .
$\forall i \in E : \sum_{j \in M^e(i)} y_{ij}^{em} \geq 2 \times x_i^e$	Each selected exchange site should be covered by 2 metro nodes.
$\forall j \in M \forall i \in E^m(j) : x_j^m \geq y_{ij}^{em}$	A metro node j is selected if any exchange site i within its reach is connected to it. An alternative way that would result in a weaker LP bound would be $\forall j \in M : E^m(j) \times x_j^m \geq \sum_{i \in E^m(j)} y_{ij}^{em}$.
$\sum_{i \in M} x_j^m = s^m$	The number of metro nodes should be equal to s^m .
$\forall i \in C : \sum_{j \in E^e(i)} x_j^m \geq k^m$	A redundant constraint that forces each customer to be connected to k^m number of metro nodes can be enforced.
Objective	
$\min \sum_{i \in C, j \in E^e(i)} 1.4 \times d_{ij}^{ce} \times y_{ij}^{ce} + \sum_{i \in E, j \in M^e(i)} 1.4 \times d_{ij}^{em} \times y_{ij}^{em}$	Minimize the sum of the cost of connecting customers to exchange sites and exchange sites to metro nodes

problem into two subproblems. In the following we present two such decomposition approaches:

1) *Decomposition Approach I:* Let $cost$ denote the cost of a solution of the original problem. Let $cost_E$ be the cost of connecting customers to their closest exchange sites in S_E , respectively. Let $cost_M$ be the cost of connecting S_E exchange sites to their two closest metro nodes in S_M .

The first subproblem is a decision problem where the task is to find a feasible solution of the original problem such that the sum of the lengths of the optical fibre required to connect

customers to their exchange sites is within bounds of $cost$. The second subproblem is an optimisation problem where the task is to find a set of s^m metro nodes that connects a given set of exchange sites S_E (determined in the first step) such that the cost of connecting exchange sites in S_E to their closest and second closest metro nodes is minimum subject to the constraint: $cost_M < cost - cost_E$.

These two subproblems are solved one after another repeatedly until an optimal solution of the original problem is found. Let Sol_E be the set of sets of exchange sites

known by solving the first subproblem. While solving the first subproblem, we ensure that the same set of exchange sites are never found again. The search stops when the first subproblem becomes infeasible. If the second subproblem is not infeasible then the upper bound of $cost$ is set to $cost_M + cost_E$.

2) *Decomposition Approach II.*: In the first decomposition approach the decision problem is to find a feasible solution of the original problem such that the cost of connecting customers to exchange sites is less than a given bound, whereas here it is to find a feasible solution of the original problem such that the cost of connecting exchange sites to metro nodes is less than a given bound. The second subproblem is an optimisation problem where the task is to find a set of exchange sites such that the sum of the lengths of optical fibres required to connect a customer to its exchange site and an exchange site to its two metro nodes is minimised.

The sum of the lengths of optical fibres required to connect customers to their exchange sites is $cost_E$. The sum of the lengths of optical fibres required to connect exchange sites to their metro nodes is $cost_M$. If the second subproblem is not infeasible then the upper bound of $cost$ is updated to $cost_M + cost_E$. The two subproblems are solved one after another repeatedly until an optimal solution of the original problem is found. While solving the first subproblem, we ensure that a same set of metro nodes are never found again. The search stops when the first subproblem becomes infeasible.

III. RESULTS FOR THE DEPLOYMENT OF METRO NODES AND EXCHANGE SITES

In this section we present empirical results and analysis. We used CPLEX (version 12.2) to solve all the integer linear programming formulations. All algorithms were implemented using Java. All the experiments were run on Linux 2.6.25 x64 on Dual Quad Core Xeon CPU with overall 11.76 GB of RAM and processor speed of 2.66GHz. For each pair of combination of the number of metro nodes, $s^m \in \{12, 14, 16\}$, and the number of exchange sites, $s^e \in \{135, 145, 155\}$, we tried all the three approaches. The rationale is that the operator might want to minimize the number of local-exchanges and metro nodes providing the LR-PON services.

The approach based on the MIP model presented in Section II could not scale because of its space requirements. The lower bounds were computed by relaxing the integrality constraint of all the variables of the full sized MIP model except the Boolean variables associated with the selection of the metro nodes. The lower bounds are shown in the column labelled as LB in Table II. These bounds were then used for both decompositions approaches. For the two decomposition approaches the time-limit was set to 4 hours. We observed that finding a feasible solution for the first subproblem of the

first decomposition approach was harder for CPLEX than finding a feasible solution for the first subproblem of the second decomposition approach. While both decomposition approaches were able to find solution for all the instances, in Decomposition II, we were able to carry out more iterations thus leading to better results.

Table II
COST RESULTS IN KMS.

s^e	s^m	LB	Decomposition I	Decomposition II
155	16	46125	60574	50421
145	16	46837	61147	50973
135	16	47513	62720	51788
155	14	47618	60370	51355
145	14	48437	63489	51842
135	14	49070	61549	52626
155	12	48969	63109	52449
145	12	49808	63941	53063
135	12	50216	65069	53652

Figure 2(a) shows the distribution of distances between customers and their closest local exchange. It can be seen that most customers are within 10Km of their local-exchanges. Figure 2(b) shows the fibre length required for each customer to reach their local-exchanges when the maximum number of local exchanges is 145. This also implied that the distances are longer than those in Figure 2(a), where each customer was connected to its closest local exchange. Figure 2(c) shows the fibre length required for each local-exchange to connect to a metro node when the number of metro nodes is restricted to 14. We observed that the lower the number of metro nodes, the longer the overall distance, so there is a trade-off between cost of fibre distance versus number of metro nodes (thus overall cost of fitting equipment in metro nodes). Figure 2(d) shows the fibre length required for each local-exchange to connect to a secondary metro node for protection purpose.

Finally we show the distribution of users per local-exchange, and per metro node. Figure 2(e) shows the number of users for each local-exchange when the number of local-exchanges is set to 145 and the number of metro nodes is fixed to 14. There are 6 local exchanges that have less than 10 customers each, most exchanges have between 20 and 40 customers, and only 3 have more than 100 customers.

Similarly Figure 2(f) shows the number of users per metro node, respectively for primary service and backup service, both calculated for 145 local exchanges and 14 core nodes. The minimum number of users per metro node is close to 100, most are between 150 and 400, and one has almost 900 customers. Although most of the customers are within 100KMs from their metro node, 140 (3.5 %) require an overall primary link longer than 100Km. In addition 931 customers (23.8 %) require protection paths that are longer than 100Km.

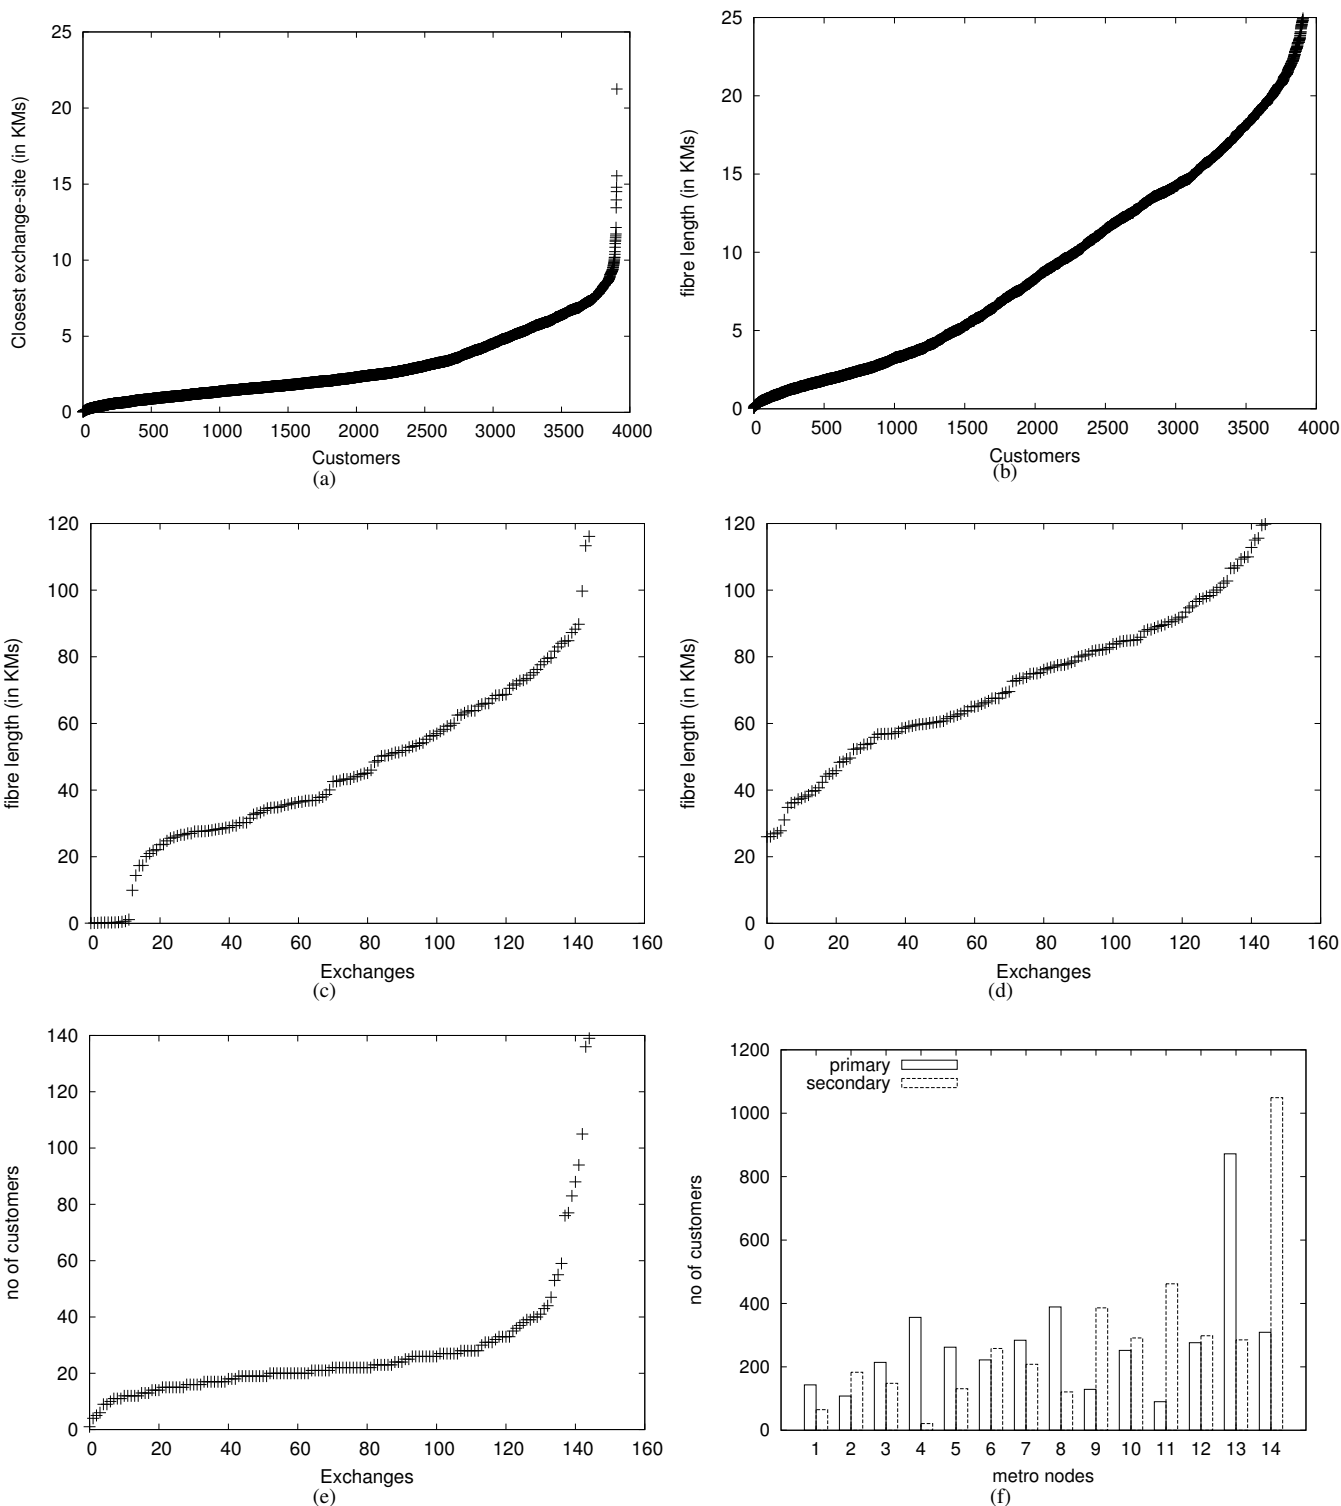


Figure 2. (a) Closest exchange site for each customer. (b) Distribution of fibre lengths required to connect customers to their closest exchange sites when the number of metro nodes is fixed to 14 and the number of exchange sites is fixed to 145. (c) Distribution of fibre lengths required to connect exchange sites to their primary metro nodes when the number of exchange sites is fixed to 145. (d) Distribution of fibre lengths required to connect exchange sites to their secondary metro nodes when the number of exchange sites is fixed to 145. (e) Distribution of number of customers over number of exchange sites when the number of exchange sites is 145 and the number of metro nodes is 14. (f) Distribution of number of customers over primary and secondary metro nodes when the number of exchange sites is 145 and the number of metro nodes is 14.

IV. ROUTING A BRANCH OF THE LR-PON

In this section we consider the task of routing and branching the cables in the context of an already existing association of exchange sites with customers and metro nodes with exchange sites. That is, for every customer there is already an exchange site associated with it. Similarly, every exchange site has been already assigned its two metro nodes. The remaining task is then to find the route for the cables.

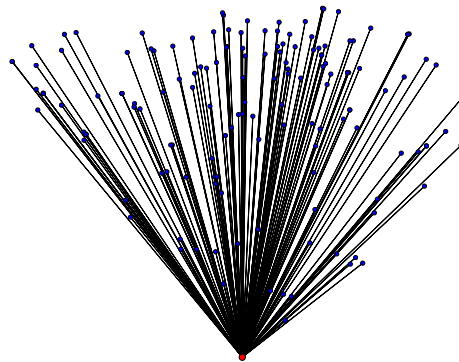
The problem of routing cable from the metro nodes to its exchange sites is different from the problem of routing cable from the local exchanges to the customer (i.e. the optical distribution network (ODN) side). In the ODN in fact the fibres are distributed to the users through a tree distribution network based on optical power splitters. At each split the optical power is reduced by $10\log N + \alpha$ dB, where N is the number of splits at each splitter and α is a constant that considers additional loss due to the non-ideality of the splitter. Thus the ODN problem also requires to introduce physical layers constraints due to the optical power loss.

In this paper we focus on the problem of routing the cable from the metro nodes to the local exchanges only. As an association between a metro node and its exchange sites is already given, we can treat each one-to-many relation (i.e., tree) independently. That is, we can focus on one single metro node and find the paths from that metro node to all the local exchanges that it covers.


In Figure 3 we show two ways of connecting a given set of local exchanges to a metro node. In the first case (Figure 3(a)), we are simply connecting each local exchange directly. In the second case (Figure 3(b)), we are computing a minimum spanning tree rooted at the metro node. Certainly the option of connecting each local exchange directly to the metro node leads to shorter connection paths (as shown in Figure 4). However, the drawback of connecting each local exchange directly is the total amount of cable used (as shown in Figure 5, where a point (x, y) means that y units of cable are used to cover x local exchanges)².

We are interested in both restricting the length of the paths and the total amount of cable used. Keeping both requirements is known to be a hard problem [3]. There has been a lot of work on this bounded version of the spanning tree problem (see [5] for a short summary of the most relevant approaches). It has also been suggested CP for tackling this problem ([5], [2]). We assume that a cost function is given. This cost function associates each arc and capacity with a cost. The objective is then to minimize the cost of the used arcs taking into account the paths going

²In the path length comparison, local exchanges are sorted by the length of the path from the local exchange to the metro node. In the coverage comparison, the links of the tree are sorted with respect to their length, where a point (x, y) represents that there are x local exchanges connected by the k shortest links and the cumulative length of those k links is y .



(a) **Trivial tree.** Each local exchange is directly connected to the metro node assuming that the branching capacity of the metro node is not limited.



(b) **Minimum spanning tree.** Local exchanges are connected to the metro node through a spanning tree rooted at the metro node

Figure 3. Two ways of connecting the local exchanges to a given metro node

through them³. What we are trying to model is a cable layout where no trenching cost is required, as the network reuses existing ducts. The costs to be considered are: those resulting from laying new fibre on an existing duct and those resulting from the capacity of the cable (which are a step-wise function of the number of fibres in the duct). These two costs generate a trade-off, as while reducing the overall amount of fibre reduces the cost of laying fibre in the duct, it also increase the size of cables on the used ducts.

A. A Mixed Integer Programming Model

Let x_{ij} be a Boolean variable that denotes whether there is a cable-link between nodes i and j . Let y_{ij} be an integer variable that denotes the number of optical fibres passing

³In the experiment that we have carried out we have ignored the dependency of the cost with the capacity, i.e., the cost of a link only depends on the cable distance.

Table III
MODEL USING ELEMENT AND ELEMENT2

Variables	
$s_i \in N$	an integer variable denoting the successor of the i -th node.
$l_i \in \mathbb{N}$	an integer variable denoting the length of the path associated with the i -th node.
$y_a \in \mathbb{N}$	an integer variable denoting the capacity associated with arc $a \in A$.
$q_{ij} \in N$	an integer variable denoting the successor at level j of the i -th node.
Constraints	
$s_m = m$	e is the only root
$\forall i \in C : i \notin \text{dom}(s_i)$	All customers need to be covered
$s_i = q_{i1}$	The successor at level 1 of node i is its immediate successor
$\text{ELEMENT}(q_{ij}, s, q_{i(j+1)})$	The successor at level j of node i is the successor of its successor at level $j - 1$
$q_{i(n-1)} = q_{in}$	As node i has at most $n - 1$ successors, the successor of its $n - 1$ -th successor is bound to be equal to the latter
$l_i = \text{ELEMENT2}(i, s_i, F) + \sum_{1 \leq i \leq n-2} \text{ELEMENT2}(q_{ij}, q_{i(j+1)}, F)$	The length of the path of node i is the sum of the costs of the arcs linking its successors
$l_i \leq \lambda$	The length of the path of node i is bounded by λ
Objective	
$\min \sum_{i \in N} \text{ELEMENT2}(i, s_i, F)$	To minimize the sum of the cost of the arcs of the tree

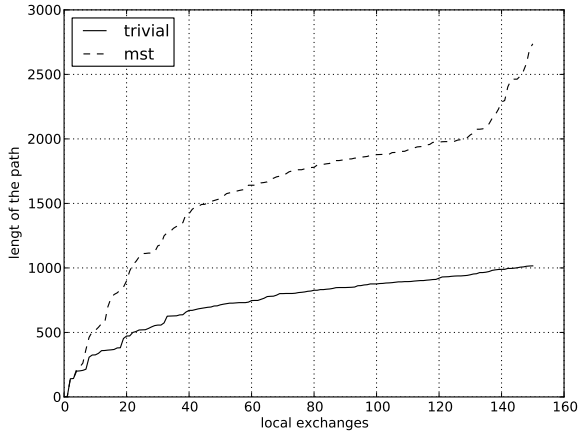


Figure 4. Comparing path length distributions




Figure 5. Comparing coverage distributions

through the cable-link between nodes i and j , where $i \in \{0, \dots, N\}$ and $j \in \{1, \dots, N\}$. We assume that Node 0 is a metro node. Let l_i be the upper bound on the length of the path of the optical fibre from the metro node to exchange site i . The number of fibres emanating from the metro node is equal to N :

$$\sum_{i=1}^N y_{0i} = N$$

Each exchange site $i \in \{1, \dots, N\}$ has only one incoming cable: $\sum_{j=0}^N x_{ji} = 1$. There exists an optical fibre that (in)directly connects each exchange site $i \in \{1, \dots, n\}$ to the metro node:

$$\sum_{j=0}^N y_{ji} - \sum_{j=1}^N y_{ij} = 1$$




Figure 6. Unbounded MST.




Figure 7. Bounded MST.




Figure 8. Performance in the computation of the optimal solution of the case in Figure 7

If there is no cable-link between nodes i and j then there is no optical fibre that connects nodes i and j directly:

$$x_{ij} = 0 \Rightarrow y_{ij} = 0$$

If there is a cable-link between exchange sites i and j then the length of the optical fibre required to connect j to the metro node is greater than or equal to the sum of the length of the optical fibres required to connect exchange site i to its metro node plus and exchange site i to exchange site j :

$$x_{ij} = 1 \Rightarrow l_j \geq l_i + d_{ij}$$

The length of the optical fibre required to connect each exchange site i to the metro node is less than or equal to a given constant λ : $l_j \leq \lambda$. The objective is to minimize the total length of cables:

$$\min \sum_{i \in \{0, \dots, N\} \wedge j \in \{1, \dots, N\}} d_{ij} \cdot x_{ij}$$

Table IV
COST RESULTS IN KMS.

s^e	s^m	No sharing	Sharing
155	16	17067	5776
145	16	16030	5630
135	16	14945	5435
155	14	17822	5749
145	14	16867	5680
135	14	15819	5528
155	12	18734	5881
145	12	17821	5640
135	12	16813	5546

We used this MIP approach for computing the last column of Table IV (Sharing), which presents the amount of cable spent if cable is shared for connecting metro nodes to local exchanges. In these experiments, we used 120Km as the value for λ . These values are compared with the ones obtained when cable is not shared (No Sharing). Consistent with the situation that we observed in Figure 5, we also see here a significant difference with respect to the case where the cable is not shared. Solving each branching instance took less than one second.

B. Constraint Programming Models

We model this routing problem in two different ways. In the models, N denotes the set of nodes, which is equal to $\{m\} \cup E$, where m is the metro node whose local exchanges we are linking and E is the set of local exchanges associated with m . We assume full connectivity of the network, so A (the set of arcs) is equal to N^2 . We use F to associate arcs with costs. There is also a global upper bound on the length of the paths associated with the nodes. We use λ to refer to that upper bound.

1) *Decomposing the Tree constraint using the Element constraint*: This model is based on the decomposition of the Tree constraint presented in Page 20 of [1]. The model relies on the following concepts:

- $\text{ELEMENT}(i, x, y)$ states that the i -th element in x is y (i.e., $x_i = y$).
- $\text{ELEMENT2}(i, j, x, y)$ states that the $((i, j)$ -th element in x is y (i.e., $x_{ij} = y$). In the model, we will avoid the usage of auxiliary variables by using this constraint as a function:

$$\begin{aligned} & \text{ELEMENT2}(i, j, x, y) \wedge R(\dots, y, \dots) \\ \equiv & R(\dots, \text{ELEMENT2}(i, j, x), \dots) \end{aligned}$$

The model is presented in Table III.

2) *Using the Tree Constraint*: In this model, in addition to the ELEMENT constraint, we use the Tree constraint. $\text{TREE}(s, l, \eta, \alpha)$ holds if the integer variables in s , which represent the successors of the nodes (i.e., s_i is the successor of node i), have been assigned values in such a way that they correspond to a forest with η trees, of which α are proper trees (i.e., trees involving more than one node). The root of

Table V
MODEL USING TREE AND ELEMENT2

Variables	
$s_i \in N$	an integer variable denoting the successor of the i -th node.
$l_i \in \mathbb{N}$	an integer variable denoting the length of the path associated with the i -th node.
Constraints	
$\text{TREE}(s, l, \eta, \alpha)$	The connection of customers to the exchange site is a tree
$s_e = e$	e is the only root
$\forall i \in C : i \notin \text{dom}(s_i)$	All customers need to be covered
$\forall i \in N : \text{dom}(l_i) = \{v \in \mathbb{N} v \leq \lambda\}$	The length of every path should be less than the given upper bound
$\eta = 1$	There is only one tree in the forest
$\alpha = 1$	The only tree in the forest is proper since it contains all the nodes
Objective	
$\min \sum_{i \in N} \text{ELEMENT2}(i, s_i, F)$	To minimize the sum of the cost of the arcs of the tree

each tree is a node whose successor is itself. The successor of a node corresponds to the parent of the node in the forest. l is an array of integer variables representing the length of the paths from the nodes to their roots (i.e., l_i is the length of the path from node i to the root of its tree)⁴. The model is presented in Table V.

3) *Evaluation*: We implemented the first CP model in Choco 2.1.4⁵. The second model was not implemented because the Tree constraint was not operational at the time of writing this paper⁶. While the evaluated CP approach is still quadratic with respect to the number of constraints, it let us anticipate the challenges to be faced in the implementation of the second CP model, which is indeed much more compact than the evaluated approaches.

In Figures 6 and 7, the metro node is Node 0 and the other nodes are local exchanges. In Figure 6 local exchanges are covered using an unbounded minimum spanning tree, which leads to very long paths (e.g., the path to reach local exchanges 7). In Figure 7 we limit the maximum length of the path to the maximum direct distance from a local exchange to the metro node. This restriction forces local exchanges 2,6,7 to be directed connected to the metro node.

When it comes to performance, the evaluated CP approach behaves very poorly with respect to the MIP approach (see Figure 8). Even though the optimal solution is computed relatively quickly (for the small cases that we tried), the time spent in proving optimality was quite high. One the reason for the model to behave so poorly is that no pruning is taking place at the level of the objective function. Another reason is that the determination of the successor variables is not taking into account the structure of the solution.

V. CONCLUSIONS

We have presented a sophisticated combinatorial optimisation approach to a large and challenging network design

⁴We present here a simplified version of the constraint. The full version of the constraint can be found in [5].

⁵<http://www.emn.fr/z-info/choco-solver/>

⁶http://sourceforge.net/tracker/index.php?func=detail&aid=3534214&group_id=96738&atid=615736

problem from a real-world setting. Our approach comprised two phases. The first phase seems suitable for mixed integer programming since the number of constraints is linear with respect to the number of nodes. The second phase seems less suitable since the number of constraints is quadratic. The extensive evaluation of the first stage indeed shows that MIP is appropriate for this. For the second phase we saw that our MIP approach comfortably handled the instances we are currently considering. However, as the number of constraints is quadratic with respect to the number of nodes in the network, we felt motivated to explore a more compact representation with CP. Our current CP results suggest that more work needs to be done both at the level of propagation and the level of labeling strategy.

REFERENCES

- [1] Nicolas Beldiceanu, Mats Carlsson, Pierre Flener, and Justin Pearson. On the reification of global constraints. Technical Report T2012:02, SICS, Kista, Sweden.
- [2] Nicolas Beldiceanu, Pierre Flener, and Xavier Lorca. Combining tree partitioning, precedence, and incomparability constraints. *Constraints*, 13(4):459–489, 2008.
- [3] J. M. Ho and D. T. Lee. Bounded diameter minimum spanning trees and related problems. In *Proceedings of the fifth annual symposium on Computational geometry, SCG '89*, pages 276–282, New York, NY, USA, 1989. ACM.
- [4] Barry O'Sullivan Luis Quesada Linda Doyle Marco Ruffini, Deepak Mehta and David B. Payne. Deployment strategies for protected long-reach pon. *Journal of Optical Communications and Networking*, 4:118–129, 2012.
- [5] Thiago F. Noronha, Andréa C. Santos, and Celso C. Ribeiro. Constraint programming for the diameter constrained minimum spanning tree problem. *Electronic Notes in Discrete Mathematics*, 30:93–98, 2008.