

Designing Resilient Long-Reach Passive Optical Networks

Deepak Mehta and Barry O’Sullivan and Luis Quesada

CTVR, Cork Constraint Computation Centre
University College Cork, Ireland

Marco Ruffini and David Payne and Linda Doyle

CTVR, University of Dublin
Trinity College, Ireland

Abstract

We report on an emerging application focused on the design of resilient long reach passive optical networks using combinatorial optimisation techniques. The objective of the application is to determine the optimal position and capacity of a set of metro nodes. We specifically consider dual parented networks whereby each customer must be associated with two metro nodes. An important property of such a placement is resilience to single node failure. Therefore excess capacity should be provided at each metro node in order to ensure that customers can be redistributed amongst the metro sites. Our application, as well as finding optimal node placements, can compute the minimum level of excess capacity on all metro nodes. In this paper we present three alternative approaches to optimal metro node placement. We present a detailed analysis of the impact of different placement approaches on the distribution of excess capacity throughout the network. We show that preferential distributions occur in practice, based on a case-study in Ireland. Finally we show that load and excess capacity provision are independent of each other.

Introduction

Long Reach Passive Optical Networks (LR-PONs) are gaining increasing interest as they provide a low cost and economically viable solution for fibre-to-the-home network architectures (Payne 2009). The LR-PON by-passes local exchanges and replaces the metro/back haul transmission systems that would previously exist between the old local exchange sites and the metro-node or outer core node. A major fault occurrence would be a complete failure of the metro-node that terminates the LR-PON, which could affect tens of thousands of customers. The *dual homing* protection mechanism for LR-PON enables customers to be directly connected to two metro nodes, so that whenever a single node fails all customers are still connected to a back-up or protection node (Hunter, Lu, and Gilfedder 2007). This is similar to a simple protection solution for IP routers known as double or redundant protection (Maeschalck et al. 2002). For such a protection mechanism, additional capacity needs to be provided at each metro node so that the capacity of each metro node is not less than that required by the sum of its primary and secondary customers that are switched to it

following the failure of their primary nodes. Providing such additional capacity at each metro node incurs additional cost and increased energy consumption. Therefore, we are concerned with minimizing the additional capacity required at each metro site while ensuring that the overall network is capable of withstanding any node failure.

Our solutions rely on the application of standard combinatorial optimisation techniques to the design of resilient LR-PONs. The advance achieved in the work presented in this paper is that we have developed a constraint-based application for designing long reach passive optical networks. The objective of the system is to determine the optimal position and capacity of a set of metro nodes. This paper focuses specifically on the impact of (customer) load distribution over metro nodes on the optimal excess capacity (over-provisioning) required in a dual-homing LR-PON architecture. In particular, we study two distribution models for initial customer load: one distributing load uniformly over the metro nodes, while the other iteratively places customers at nodes in proportion to their existing load, inspired by the Barabasi-Albert model for preferential attachment in networks (Barabási and Albert 1999). We show that the distribution of customers in Ireland follows that produced by our preferential model, thereby providing us with a tool for studying the scalability of our approach. We also show that a consequence of this model is that loads follow a very skewed distribution, with knock-on consequences for how capacity over-provision should be distributed in order to ensure that our PONs are resilient to single-node failures.

Related Work. There has been a significant amount of work on the use of combinatorial optimisation techniques in the design and management of optical networks. Most papers in this area have focused on one of three problem classes: placement and location problems (Skorin-Kapov, Skorin-Kapov, and Boljuncic 2006); wavelength and routing assignment (Simonis 2009); and resilience and protection (Palmieri and Fiore 2006). Due to the computational complexity of optimally managing traffic on optical networks approximation algorithms for traffic grooming have been proposed (Antonakopoulos and Zhang 2009).

The most relevant related work to our contribution in this paper is the work on dual-homing protection using integer linear programming (Wang et al. 2004) and local search (Lee and Koh 1997). While these papers, and others, formulate

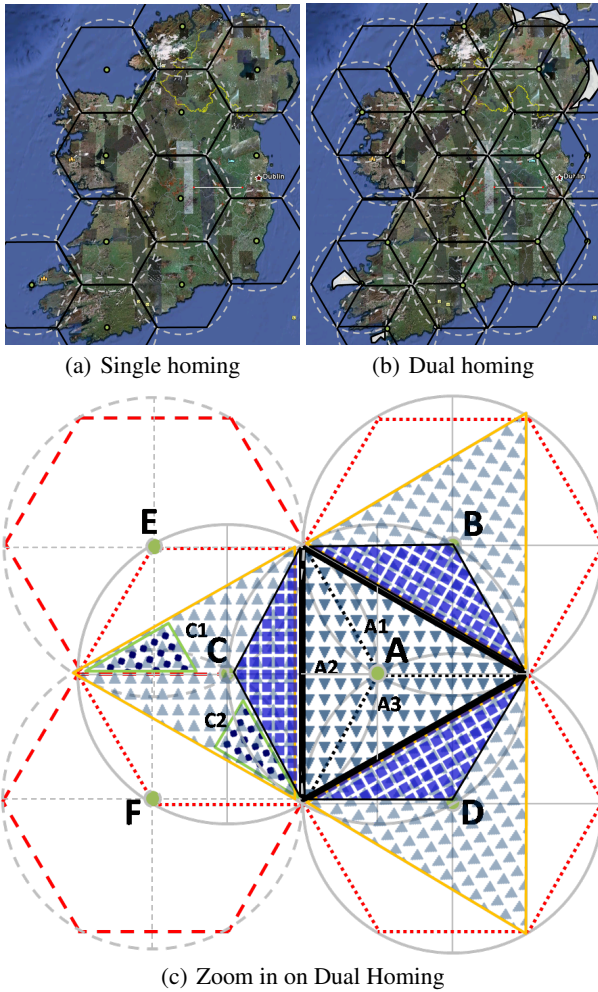


Figure 1: A geometric (honeycomb) model for designing resilient LR-PONs.

the dual-homing problem as an optimisation problem, none focus on the optimisation of the amount of over-capacity needed throughout the network, which we focus on here, amongst other things.

A Geometric Resilience Model

We present a coverage plan that determines over-provisioning levels for protection equipment (Ruffini, Payne, and Doyle 2010). The part of the LR-PON that we want to protect is between the metro node and the exchange site. The load of an exchange site is the number of customers to which it is connected. The load of a metro node is the sum of the loads of all the exchange sites to which it is connected. Each metro node can offer fiber access to the area inside a circle centered at the node, with radius equal to the optical reach, typically 60km, divided by the routing factor.

In order to provide protection coverage to a circular area we need to provide additional metro nodes whose areas overlap those of the node we want to protect. We can approximate each circle with its inscribed hexagon. Therefore, an

arbitrary area can be covered with a honeycomb structure, similar to that of a GSM network. Figure 1(a) show an example for Ireland, where a basic unprotected coverage can be achieved with 12 nodes.

Dual Homing. Dual homing can be achieved by overlapping a minimum number of non-concentric circles, as shown in Figure 1(b). If we zoom in on a part of the dual homing system as shown in Figure 1(c), we can see that each metro node (e.g. *A*) offers primary coverage to the exchange sites that are closer to it than to any other adjacent metro nodes (e.g. *B*, *C* and *D*). This area is the equilateral triangle covered with triangle shading. Since the area of the triangle is half the area of the hexagon, each node requires half the IP equipment for primary coverage compared to the case where there is no protection. In addition, since each node can physically cover an area as large as the hexagonal cell, it can provide protection for those areas within the hexagon but external to the triangle. Clearly, the advantage of this dual homing scheme based on a honeycomb structure is that if any node, say *A*, fails, its load can be spread amongst the adjacent nodes, *B*, *C* and *D*, which protect, respectively, sectors *A1*, *A2* and *A3*. There are issues that can arise from dual-homing based on honeycomb structure, e.g. border effects whereby the most external areas of the structure only provide primary coverage. This can be overcome by tuning the distance between metro sites. In this paper we are not focused on this problem.

Distributing Load Across the Network. In order to handle the additional load placed on the metro nodes in the case of a neighboring node failure, metro nodes can also pass on to their adjacent nodes *part of their primary traffic*, thus reducing the level of over-provisioning required. This is a form of pre-emptive redistribution of traffic in response to a major network failure. This method significantly reduces the IP routing resources required for protection purposes at each node, leading to significant cost savings.

Optimisation Model for Over-provisioning

We have modelled the problem of reducing the over-provisioning capacity over the metro nodes subject to the constraint that it can accommodate the load generated by the failure of any single node as an optimization problem.

Model Description. The topology of a LR-PON is modelled as a directed graph, (V, E) , where V is a set of metro nodes and E is a set of directed edges. An edge from metro node i to j , $\langle i, j \rangle$, represents that there is one or more customers that are covered by both i and j such that i is their primary node and j is their secondary node and i can pass some or all of those customers to j , if required. Each node $i \in V$ is associated with an initial load (or the initial number of customers) denoted by Q_i . Each edge $\langle i, j \rangle \in E$ is associated with a constant U_{ij} , which is an upper bound on the load that can be transferred from i to j . The shortest path distance between each pair of nodes i and j is denoted by SP_{ij} , and h denotes the maximum number of hops (or distance) from the failed node that are allowed for the spreading of the load.

Table 1: Model for Overprovisioning Capacity

Objective :	$\text{Min } \sum_{i \in V} \alpha \times (M_i - Q_i) + \sum_{\forall i,j,k \in V} T_{ijk}$
such that:	
$C_1:$	$I_{ik} = \sum_{\langle j,i \rangle \in E} T_{jik}$
$C_2:$	$O_{ik} = \sum_{\langle i,j \rangle \in E} T_{ijk}$
$C_3:$	$I_{kk} = 0$
$C_4:$	$O_{kk} = Q_k$
$C_5:$	$SP_{ki} > h \Rightarrow I_{ik} = 0$
$C_6:$	$F_{ik} = Q_i + I_{ik} - O_{ik}$
$C_7:$	$M_i \geq F_{ik}$

An integer linear programming formulation of this problem as presented in Table 1 is described below.

Variables. For each triplet of metro nodes i , j and k , an integer variable T_{ijk} is used to represent the number of customers that are off-loaded from i to j when k fails. The domain of T_{ijk} is $\{0, \dots, U_{ij}\}$. For each pair of metro nodes i and k three integer variables are used: I_{ik} , O_{ik} , and F_{ik} . Here, I_{ik} denotes the sum of incoming loads that i receives from its neighbours when k fails, O_{ik} denotes the sum of outgoing loads that i passes to its neighbours when node k fails, and F_{ik} denotes the final load of i that includes the over-provision capacity that is required when k fails. For each metro node i an integer variable M_i is used to represent its maximum of the final loads over all possible failures.

Constraints. When a metro node k fails, the incoming and outgoing loads of a metro node i different from k are I_{ik} (C_1) and O_{ik} (C_2) respectively. When a metro node k fails, its incoming load is zero (C_3) and its outgoing load is equal to its initial load Q_k (C_4). Additionally, for each metro node i if $SP_{ki} > h$ then the incoming load of i is also 0, i.e., $I_{ik} = 0$ (C_5). The over-provision capacity of a metro node i when a metro node k fails is the difference between its incoming load and outgoing load. Therefore, when a node k fails, the final load of node i is the sum of its initial load and the over-provision capacity (C_6). The load capacity of a node i has to be greater than or equal to the maximum of the final loads for all possible node failures k (C_7).

Objective. The objective is to minimize the total amount of IP over-provision capacity required over all metro nodes. It is also desirable to minimize the number of customers that are affected by the redistribution of load. Therefore, objective is as follows: $\sum_{i \in V} \alpha \times (M_i - Q_i) + \sum_{\forall i,j,k \in V} T_{ijk}$, where α is any constant that is greater than $\sum_{\forall i,j,k \in V} T_{ijk}$.

Model Properties. Each failure of a node k can be associated with a digraph where the set of the directed edges is the set of $\langle i, j \rangle$ such that $T_{ijk} > 0$. We call this the *load transfer graph* of metro node k . This digraph is acyclic since we are also minimizing the load transfer. It is important to emphasize that the shortest path constraint (C_5) is used to restrict the inclusion of the metro nodes in the load transfer graph that are not reachable within a given distance h from a failed node k . Nevertheless, it does not avoid the existence of a path in the graph such that the distance (or number of hops) between two metro nodes is greater than the given h . This is




Figure 2: A load transfer graph

illustrated in the example shown in Figure 2.

Figure 2 shows an example of a load transfer graph associated with the failure of node a when the number of hops is restricted to 1. Each node in the graph is labeled with an interval representing the initial load of the node and the maximum load that it can handle. Each edge $\langle i, j \rangle$ is labeled with an interval whose upper bound is the number of exchange sites that have i as primary node and j as secondary node, and an integer representing the number of exchange sites that are transferred given the failure. In case of a failure of i , all those exchange sites having i as primary node are transferred to their secondary nodes.

Evaluation over Artificial Networks

Our focus on over-provision capacity is unique in this area, so we present an evaluation that shows the factors that effect the level of such capacity required in various settings. We first focused on randomly generated networks, but inspired by the Irish data, for our empirical study. As an empirical parameter in our experiments with respect to the routing factor of 1.6, we varied the coverage radius of metro nodes from 50 km to 90 km in steps of 10, giving rise to different packing densities of metro nodes. We only plot data for radii 50 and 90 since these are the extreme values between which all other data are contained. We also assigned 5000 units *uniformly* at random, or following a preferential distribution (Barabási and Albert 1999). According to the latter, units of load were iteratively added to the metro nodes such that at each iteration a unit of load was assigned to a node with a probability proportional to the existing load on that node. At the limit such an approach generates load distributions that follow a power-law similar to those found in the degree distributions of scale-free networks. Therefore, we present many of our results using *survival functions* in which a point shows the probability of a load (over-provision) exceeding the corresponding value on the x -axis. Uniform distributions can be observed as plots with tails that fall almost vertically. Distributions that tend to lift upwards and rightwards, and which are beginning to straighten, correspond to heavy-tailed distributions; a pure power-law presents as a straight-line when the axes are both logarithmic. Figure 3 presents the load and over-provision distributions for different values of the coverage radius of nodes in the Irish context, which involves 1100 exchange sites each serving several hundred customers. In Figures 4(a) and 4(b) we show the same distributions for our uniform, denoted with a u and a radius, and preferentially loaded random networks, de-

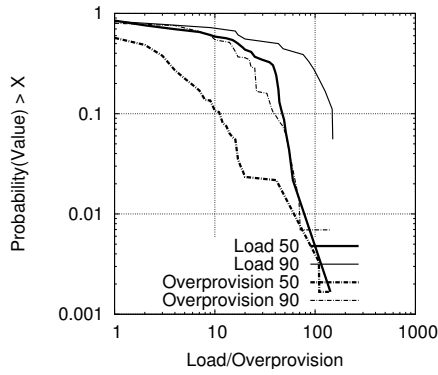



Figure 3: Ireland: Load/over-provision.

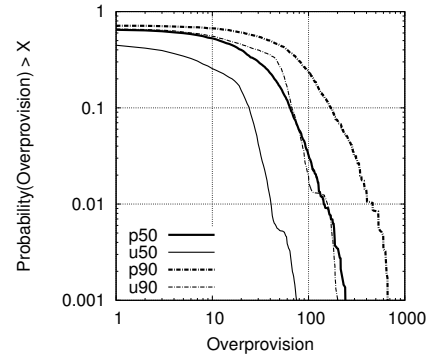
noted with a p and a radius. The distribution in Figure 3 is very similar to the distribution one obtains from the preferential model we have studied (Figures 4(a) and 4(b)), and quite different to that from the uniform model (same figures), since the tail of the survival function in the load plot in the latter case falls much more quickly. This supports the claim that our preferential model is representative of real-world networks.

Figure 4(d) presents the effect of the coverage radius on the average number of metro nodes needed in the random networks we considered, as well as the amount of over-provision capacity needed in both the uniform and preferential load models. As expected, as coverage radius increases, the number of metro nodes decreases. However, there is little effect on the level of over-provisioning needed in either the uniform or preferential settings, apart from much more being necessary in the latter case due to the demands placed by the very few but heavily loaded metro nodes on their neighbours from a robustness perspective.

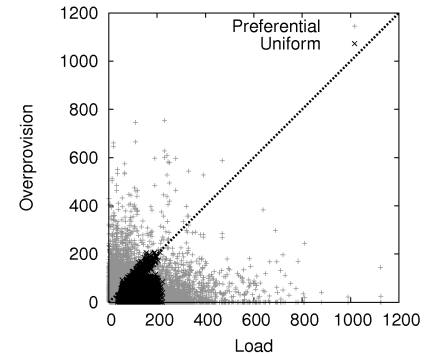
In Figures 4(a) to 4(c) we present a detailed analysis of how loads and over-provision capacities are distributed across nodes under the two load distribution models described above as metro node covering radius is also varied. As expected, in the uniform case the load (Figure 4(a)) is evenly distributed across the metro sites; this is illustrated by the fast fall in the tail of the uniform plots (denoted u). We tend to need to add a similar level of over-provision capacity across the nodes in this setting (Figure 4(b)). On a per node level, there is little variation in terms of load/over-provision (Figure 4(c)). However, in the preferential case (denoted p) the distribution of load across the metro sites is heavily skewed, with very few highly loaded nodes (Figure 4(a)). Similarly, the over-provision is also distributed in a very skewed manner - note that the tail on the right of the over-provision plot extends over a very large range (Figure 4(b)). At a node level there is evidence of a negative correlation between the level of load and the amount of over-provision (Figure 4(c)). If a node is lightly loaded, it will often have to be heavily over-provisioned to handle the few, but heavily loaded, nodes. Clearly load distribution has a significant effect on over-provision capacity characteristics.



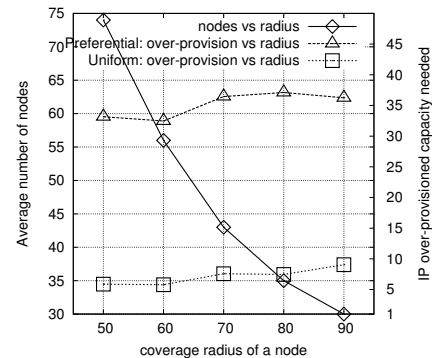
(a) Load survival function.



(b) Over-provision survival function.



(c) Over-provision versus load.



(d) Coverage and loads.

Figure 4: Empirical results on random networks.

Table 2: Comparison of distance and over-provision for different approaches to placing metro nodes

#mn	road-distance			over-provision			average degree			diameter			tree proportion		
	HC	DIST	TRANS	HC	DIST	TRANS	HC	DIST	TRANS	HC	DIST	TRANS	HC	DIST	TRANS
8	297,987	115,723	172,898	1,079,146	1,423,688	338,075	1.87	2.00	3.00	0.75	0.75	1.50	100.00%	75.00%	0.00%
12	182,727	92,701	129,487	366,085	1,361,836	289,168	2.08	2.50	3.58	1.83	0.75	1.83	75.00%	83.33%	8.33%
16	143,581	77,348	106,382	533,854	1,351,484	275,474	2.00	2.12	3.56	1.44	0.75	1.94	100.00%	87.50%	12.50%
20	148,711	68,185	97,483	560,429	1,303,144	320,245	2.39	2.45	3.50	1.75	0.75	2.00	70.00%	80.00%	15.00%
24	107,892	61,683	77,443	565,414	1,144,976	325,798	2.12	2.66	3.54	1.75	0.79	2.00	87.50%	83.33%	4.17%

Evaluation over the Ireland Network

In this section we study three approaches for placing the metro nodes in Ireland and in the subsequent section we shall study their impact on the over-provision capacity.

The first approach that we have already studied is based on the honey-comb structure where intersection points amongst the hexagons define the locations of metro sites. This way of placing metro nodes will be denoted as HC. For the remaining two approaches the positions of k metro nodes are selected from the set of the positions of the given exchange sites. The second approach selects the positions of the metro nodes such that the sum of the distances between the exchange sites and their corresponding primary and secondary metro nodes is minimized. Minimizing the road-distance can minimize the length of the cable required for connections. This method is denoted as DIST. The third approach selects the positions of the metro nodes such that the maximum of the loads that can be transferred between pairs of metro nodes is minimized. Here the hypothesis is that when less load is transferred from one metro node to another then less over-provision capacity will be required on the metro node. This approach will be denoted as TRANS.

DISTANCE Approach. We describe the Constraint Optimization Problem (COP) formulation for choosing positions for k metro nodes for the DIST approach.

Constants and Variables. Let E be a set of exchange sites whose locations are fixed. Let $d(i, j)$ denote Euclidian distance between exchange sites i and j . Let k be the number of metro nodes whose positions are to be determined. Let M be a vector of integer variables. $M(j)$ refers to the position of a metro node j . The domain of each $M(j)$ is $\{1 \dots |E|\}$. If $M(j)$ is assigned r then it means that the position of metro node j is the position of exchange site r . Each $i \in E$ is associated with two integer variables: P_i and S_i . The domain of P_i and S_i is $\{1 \dots k\}$. If j is the primary (resp. secondary) metro node of exchange site i then $P_i = j$ (resp. $S_i = j$).

Constraints. The primary and secondary nodes of a site i are different: $P_i \neq S_i$. For each pair of nodes i and j , if $i < j$, $M(i) < M(j)$ is enforced to avoid symmetries.

Objective. The objective is to minimize the sum of the distances between exchange sites and their metro nodes, i.e. $\min \sum_{i \in E} d(i, M(P_i)) + d(i, M(S_i))$.

Solution Technique. We modeled this problem using mixed integer linear programming (MILP) and used CPLEX to solve it. MILP problems are solved by a branch and bound search mechanism. An LP relaxation of the problem is solved at each node of the search tree. If the optimal value of the relaxation is greater than or equal to the value of the best can-

didate solution found so far, the search backtracks. Otherwise, if all variables in the LP solution are integral, then it becomes a candidate solution. If one or more variables are non integral, the search branches on one of the non integral variables by splitting its domain. Cutting planes are commonly added at the root node and possibly at other nodes, resulting in a branch and cut method (Rossi, van Beek, and Walsh 2006). The results that we obtained using this technique are presented later in the paper.

TRANSFER Approach. The COP formulation for choosing positions for k metro nodes for the TRANS approach is described below. The variables and constraints are strictly a superset of the COP formulation of the DIST approach. Therefore in the following we only mention those constants, variables and constraints that are not already mentioned.

Constants and Variables. Let l_i be the load of $i \in E$ which is equivalent to the number of customers associated with $i \in E$. Let T_{ij} be a variable that is used to denote the load that can be transferred from a metro node i to another metro node j . Let T_{max} be the maximum load that can be transferred between any pair of metro nodes i and j .

Constraints. If j is the primary metro node of $i \in E$ then there does not exist any other metro node r such that the distance between the positions of i and r is less than the distance between the positions of i and j : $(P_i = j) \rightarrow \forall r((r \neq j) \Rightarrow d(i, M(j)) \leq d(i, M(r)))$. If j is the secondary metro node of an exchange site i then there does not exist any other metro node r such that r is not the primary node of i and the distance between the positions of i and r is less than the distance between the positions of i and j : $(S_i = j) \rightarrow \forall r((r \neq i \wedge r \neq P_i) \Rightarrow d(i, M(j)) \leq d(i, M(r)))$. The load transferred from a metro node i to another metro node j is equal to the sum of the loads of all exchange sites whose primary metro node is i and secondary metro node is j : $T_{ij} = \sum_{P_k=i \wedge S_k=j} l_k$. The following constraint is enforced for each pair of metro nodes i and j : $T_{max} \geq T_{ij}$.

Objective. The objective is to minimize the value of T_{max} , that is to minimize the maximum of the loads transferred between any pair of metro nodes T_{ij} .

Solution Technique. A local search algorithm is designed to solve this problem. Initially k positions are selected randomly for metro nodes out of set E . A move is defined by changing the position of one of the metro nodes. The size of the neighbourhood is bounded by $2(n - k)$, where n is the number of exchange sites. The first improving move is always selected. When all moves are tried and if none of them improves the objective function then a random solution is selected again. The search stops when either a given time

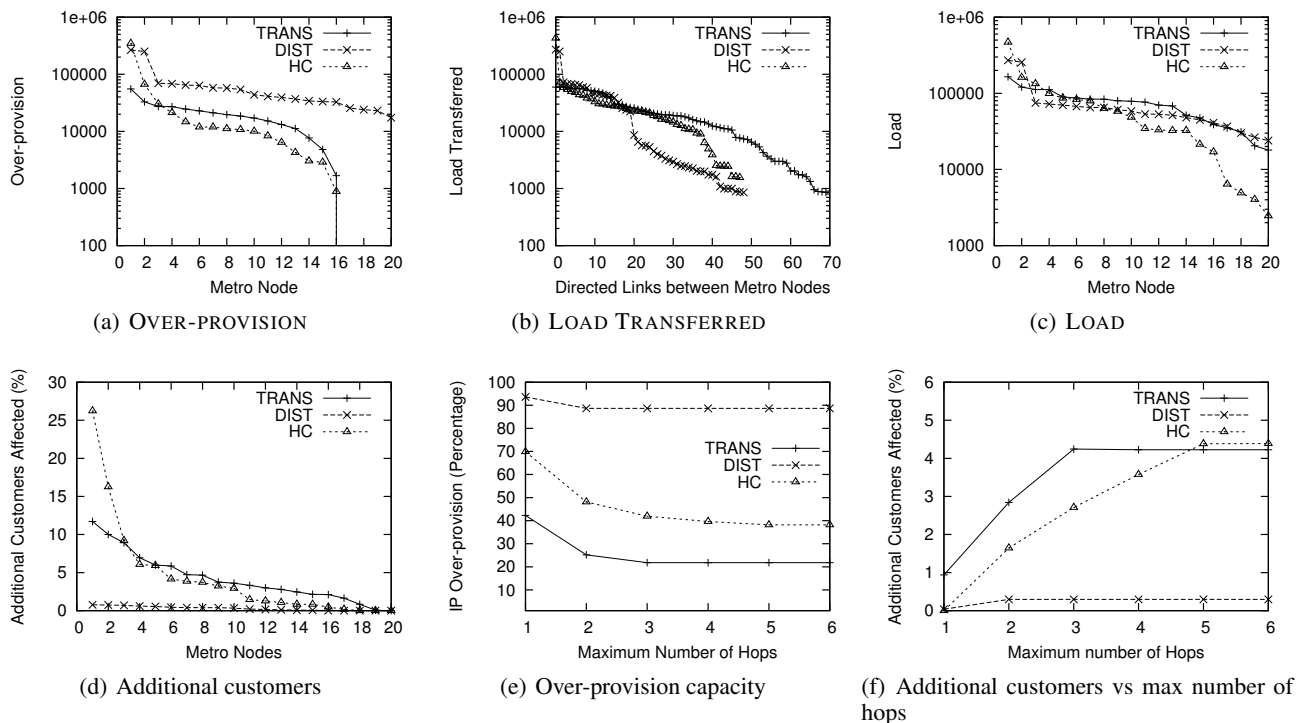


Figure 5: Placement of 20 metro nodes in Ireland.

runs out or the average gradient of improvement of the objective value per iteration is greater than a given threshold. The worst-case complexity of a single iteration is $\mathcal{O}(nk^2)$.

Impact of Placement Approaches

We investigate the impact of placing metro nodes in Ireland using different approaches to determining the over-provision capacities of the metro nodes. In particular, we focus on HC, DIST and TRANS approaches. The amount of optical fiber to connect two network points is usually larger than their Euclidean distance because fiber paths usually follow the road-lay-out. Therefore, we have applied a routing factor of 1.6. The cost parameters of a dual-parented LR-PONs network solution depends on *the number of metro nodes*, the *total road-distance* between the metro nodes and the exchange sites, and the *over-provision capacity* required on the metro nodes for the resiliency purpose.

In our experiments for Ireland, we varied the number of metro nodes from 8 to 24 in steps of 4. All the placement approaches were evaluated based on the total road-distance and the total over-provision capacity. The results in terms of time are not presented since these are off-line approaches. In any case, none of the approaches took more than 2 hours computation time for solving any instance. We remark that the two techniques for solving the different optimization problems considered in this paper, namely mixed integer linear programming and local search, are well known approaches in the AI community (Russell and Norvig 2003).

One can observe in Table 2 that DIST is the best in terms of road-distance. This is obvious since the nodes are placed

in such a way that the road-distance is minimized. It is interesting to see that TRANS is always better than HC in terms of road-distance. In HC, fixing the position of a single metro node determines the positions of the other nodes. Thus, it restricts the possibilities of decreasing the road distance by placing the metro nodes in more convenient locations, which makes HC perform worst in terms of road-distance.

The over-provision distribution for 20 metro nodes is shown in Figure 5(a). As shown in the graph, the over-provision capacities of almost all the metro nodes for DIST are more than that required by the other approaches. The reason is that DIST is the worst approach in terms of load capacity. The over-provision capacities of most of the metro nodes for HC is less than that required for TRANS. However, in a very small number of cases, it is higher by an order-of-magnitude. Therefore, in terms of the total over-provision capacity TRANS is better than HC.

Each placement approach can be associated with a digraph where a vertex represents a metro node and an edge $\langle i, j \rangle$ represents that i can transfer non-zero load to j in case of a failure. As TRANS tries to minimize the maximum of the loads that can be transferred from one node to another, it ends up creating more edges. As shown in Figure 5(b) there are 70 directed edges for TRANS and less than 50 for the other approaches. Thus, the average degree of the digraph associated with TRANS is higher than that of the other approaches as shown in Table 2. Notice that the directed graphs of HC have the smallest average degree since each node in this type of graph is limited to have at most 3 neighbors. When the average degree is higher it means that load will

be shared by more nodes, which may result in requiring less over-provisioning capacity. The variance of the load distribution for HC is higher than the other approaches as shown in Figure 5(c), and therefore the variance of the over-provision distribution is also higher as shown in Figure 5(a). Interestingly, the load distribution for DIST is very similar to that of TRANS. However, the over-provision capacities of all the metro nodes for DIST is significantly more than that required for TRANS. This is because the average degree for TRANS is higher than that of with DIST.

When nodes are well over-provisioned they are unlikely to pass load to their neighbors when receiving load from a failed node. This in turn will affect fewer additional customers. Here by an additional customer we mean a customer whose primary node has not failed but it is passed to its secondary metro node. This is consistent with the results shown in Figure 5(d) for DIST. Figure 5(e) reports on the percentage of over-provisioning required for different approaches when the number of hops are restricted to different values. Figure 5(f) shows the percentage of additional customers that are affected for different number of hops.

In Table 2, we also present the results of studying the structures of the transfer load graphs associated with each approach for placing metro nodes. For each category of instances we compute the average diameter (i.e., the length of the longest shortest path between any pair of nodes) of the corresponding load transfer graphs, and the proportion of the graphs that are trees. We observe that DIST load transfer graphs have very low diameter. We also observe that, for both HC and DIST, the proportion of trees is very high. If the load transfer graph is a tree, a diameter equal to one means that no additional clients are affected. This is consistent with the situation we are observing in Figure 5(f) where DIST is showing a very low client affectation.

Overall, results suggest that HC is inferior to TRANS in terms of road-distance and over-provision capacity. In terms of over-provision TRANS is the best but in terms of road-distance DIST is the best. In order to compare these two approaches we need cost models for the number of metro nodes, road-distance and over-provision capacity. Although we have a cost-model for the road-distance this model does not take into account the number of metro nodes and the over-provision capacity. Depending upon the cost-models it may be possible to develop hybrid approaches that minimize the total cost.

Conclusions

We have presented the first study of the effects of customer distribution and metro node coverage in dual parented LR-PONs that are robust to single node failures. In order to compute the minimum additional capacity required at each metro site, we modeled the problem using mixed integer linear programming.

We proposed and studied a preferential load distribution model that occurs in practice, as demonstrated by an analysis based on Irish data. We showed that node load and capacity over-provision are heavily determined by load distribution and not on each other alone in the HC approach. Our study

also shows that the IP over-provisioning capacity requirement for the chosen protection scheme is independent of the reach of the nodes of the PON.

We also studied three approaches to node placement: one minimized distance from metro-nodes to customers (DIST), the second placed a geometrical honeycomb structure with nodes at the vertices (HC), and a third where the offloaded traffic transferred after the event of a single node failure is minimized (TRANS). The study has shown that of these approaches the TRANS method of node assignment requires the minimum over-provision capacity.

Acknowledgement. Supported by Science Foundation Ireland Grant 08/CE/I1423.

References

- Antonakopoulos, S., and Zhang, L. 2009. Approximation algorithms for grooming in optical network design. In *INFOCOM*, 1548–1556. IEEE.
- Barabási, A.-L., and Albert, R. 1999. Emergence of scaling in random networks. *Science* 286:509–512.
- Hunter, D. K.; Lu, Z.; and Gilfedder, T. H. 2007. Protection of long-reach PON traffic through router database synchronization. *Journal of Optical Communications and Networking* 6(5):535–549.
- Lee, C. Y., and Koh, S. J. 1997. A design of the minimum cost ring-chain network with dual-homing survivability: a tabu search approach. *Comput. Oper. Res.* 24:883–897.
- Maesschalck, S. D.; Colle, D.; Groebbens, A.; Develder, C.; Lievens, A.; Lagasse, P.; Pickavet, M.; Demeester, P.; Saluta, F.; and Quagliatti, M. 2002. Intelligent optical networking for multilayer survivability. *IEEE Communications Magazine* 40(1):42–49.
- Palmieri, F., and Fiore, U. 2006. A low cost and effective link protection approach for enhanced survivability in optical transport networks. In *ISPA Workshops*, 368–376.
- Payne, D. B. 2009. FTTP deployment options and economic challenges. In *Proceedings of the 36th European Conference and Exhibition on Optical Communication (ECOC 2009)*.
- Rossi, F.; van Beek, P.; and Walsh, T., eds. 2006. *Handbook of Constraint Programming*. Elsevier. chapter 15.
- Ruffini, M.; Payne, D. B.; and Doyle, L. 2010. Protection strategies for long-reach PON. In *Proceedings of the 36th European Conference and Exhibition on Optical Communication (ECOC 2010)*.
- Russell, S., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach*. Prentice Hall, second edition.
- Simonis, H. 2009. A hybrid constraint model for the routing and wavelength assignment problem. In Gent, I. P., ed., *CP*, volume 5732 of *LNCS*, 104–118. Springer.
- Skorin-Kapov, D.; Skorin-Kapov, J.; and Boljuncic, V. 2006. Location problems in telecommunications. In Resende, M. G., and Pardalos, P. M., eds., *Handbook of Optimization in Telecommunications*. Springer. chapter 20, 517–544.
- Wang, J.; Vokkarane, V. M.; Qi, X.; and Jue, J. P. 2004. Dual-homing protection in wdm mesh networks. In *Optical Fiber Communication Conference*, TuP5.