

Optimization of Relay Location and Connectivity in Wireless Broadband Distribution Networks

Vic Grout

Faculty of Technology and Computer Science
University of Wales, NEWI
Plas Coch Campus, Wrexham, UK
E-mail: v.grout@newi.ac.uk

Abstract: A basic technique is presented for determining the initial seeding points for a wireless broadband distribution network. Variants of the method are then applied to different practical scenarios.

1. INTRODUCTION

In the race to provide broadband access to the home or small business user, the wireless option has emerged as one of the frontrunners [1]. In a wireless system, a number of distribution points act as relays between the fibre backbone and the terminal subscriber nodes. These relays may be omni-directional, serving a number of nodes in an area or steered, linking relay to node or node to node in a point-to-point or point-to-multipoint fashion. The minimal configuration in such a system is a tree with the primary distributor (PD) at its root. Figure 1 gives an example configuration.

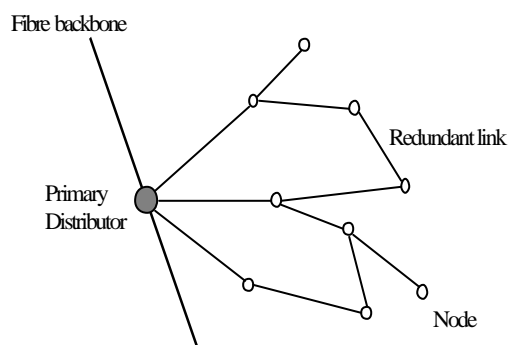


Figure 1. A Broadband Distribution Network

Redundant circuits may be added for fault tolerance and reliability. At the same time, some links may not be viable if line-of-sight (LOS) connection is absent between the two end points. The PD and other nodes chosen to act as these relay points in the initial design of the network are its *seeds*. Their position is vital as it is from them that the network will evolve. This paper discusses how best to locate these seeds.

2. FUNDAMENTAL ALGORITHM

This is an optimization problem. The input to the problem, to the algorithm that solves it, will be the subscriber node distribution requiring connection to the backbone. This may take a variety of forms and its solution subject to a number of constraints as outlined in the following sections. For any more than a dozen or so nodes, an exhaustive search approach will be prohibitively inefficient and approximation methods become necessary. If, for n nodes, the complexity of the problem is $C(n)$ and $C(n)$ is too large then we have two options: simplifying the problem using heuristics - replacing C by $C\epsilon$ so that $C\epsilon(n)$ is manageable, or reducing n to $n\epsilon$ so that $C(n\epsilon)$ is manageable - that is solving the original problem, without assumptions, but for a reduced number of nodes. The algorithm described here uses an iterative form of *representative reduction* to bring the problem down to an acceptable scale. It works as follows:

Let each node i ($1 \leq i \leq n$) be described by its co-ordinates (x_i, y_i) and a *weight*, w_i , the purpose of which will emerge as the algorithm is developed. A *distance matrix*, $D = (d_{ij})$ is also defined where d_{ij} represents the distance between nodes i and j , to be defined appropriately. An *active* number of nodes, $n\epsilon$ is initially defined: $n\epsilon = n$. We seek, at each stage, to reduce the active number of nodes by one by replacing the two 'closest' neighbours among all the nodes by a single representative having position, distance and weight characteristics reflecting its components. Thus if nodes p and q are found such that $d_{pq} \leq d_{ij}$ ($1 \leq i, j \leq n\epsilon$) and are to be replaced by a new node r , then r will have values of

$$x_r = \frac{w_p x_p + w_q x_q}{w_p + w_q}, \quad (1)$$

$$y_r = \frac{w_p y_p + w_q y_q}{w_p + w_q} \quad (2)$$

$$\text{and } w_r = w_p + w_q. \quad (3)$$

We denote this single reductive step by **[R]**. The affected elements of the distance matrix, D , are also recalculated according to specification (see Section 3.2). However p and q may not be representatively replaced in this way if certain constraints are violated as outlined below. Once a valid replacement is made, $n\mathfrak{C}$ is set to $n\mathfrak{C} - 1$ and the reduction continues, step-by-step until a stopping condition is matched, also specified in the following sections.

3. PROGRESSIVE SCENARIOS

The actual implementation of the algorithm is developed in stages, each offering additional levels of refinement to the basic method. The various constraints are independent, however, and may be applied separately or in any appropriate combination.

3.1 Small Scale - Unrestricted Distribution

We begin by assuming that we can identify target subscribers individually so that each node has equal weight in the initial configuration. $w_i = 1$ ($1 \leq i \leq n$). As clusters develop, their superior weights will pull satellite nodes towards them rather than vice-versa. If there are no connection restrictions or other geographical features to be taken into account then the matrix, D , simply takes the form of the Euclidean distance between each node pair. $d_{ij} = E(i,j) = [(x_i - x_j)^2 + (y_i - y_j)^2]^{1/2}$ ($1 \leq i, j \leq n$). After each reductive step, with nodes p and q replaced by r according to **[R]**, the p^{th} and q^{th} rows and columns are removed from D and a new row and column, r , added, with $d_{ir} = E(i,r)$ and $d_{ri} = E(r,i)$ ($1 \leq i \leq n$) calculated as before.

In this freest form of optimization, there are no constraints preventing the closest node pair being replaced at each stage. This being the case, the reduction process if left unchecked would continue until $n\mathfrak{C} = 1$. In fact this would be a remarkably inefficient method for determining the location of this single node r^* , the calculations

$$x_{r^*} = \frac{1}{n} \sum_{i=1}^n x_i \quad (4)$$

$$\text{and } y_{r^*} = \frac{1}{n} \sum_{i=1}^n y_i \quad (5)$$

giving an immediate and superior solution. However the value of this approach is that it provides a bottom-up, distribution-dependent method of *partitioning* the nodes and locating the seeds accordingly. In practice the algorithm would terminate at a predetermined value of $n\mathfrak{C} > 1$ with the $n\mathfrak{C}$ remaining nodes giving the positions of the $n\mathfrak{C}$ seeds required. $n\mathfrak{C}$ may be originally calculated based on the ratio of the number of subscribers to the area covered, the *subscriber density*. If necessary to conclude, the $n\mathfrak{C}$ seeds, now at calculated, *greenfield* sites, may be relocated to their closest *real* node from the original n .

3.2 Small Scale - Restricted Distribution

Even allowing for this cut-off in the reduction process, the situation is idealised possibly to the point of being unrealistic. In this section three practical variations are introduced, each or any of which could be applied in isolation or combination. The first acknowledges the very real likelihood that many node pairs will not have line-of-sight contact. This can be reflected in the distance matrix, D . If LOS is not present or achievable between nodes s and t then we can set $d_{st} = \infty$. As a variation, if connection can be established through an intermediary, at extra cost, then this can be factored in via a modification to the Euclidean calculation: $d_{st} = kE(s,t)$ or $d_{st} = E(s,t) + K$, for example with k, K constants. Either way, properly connected node pairs will be clustered in preference. Forced connections may be chosen as a last resort. Invalid links will never be selected.

How D is recalculated after each reduction step is flexible. If it is possible to establish whether LOS is present between the new node r and other nodes then this will be the most accurate way to proceed. Alternatively, for each node i ($1 \leq i \leq n$), d_{ri} may be approximated, pessimistically, as $\max(E(r,i), d_{pi}, d_{qi})$ or, optimistically, as $\min(E(r,i), d_{pi}, d_{qi})$ and symmetrically for d_{ir} . In other words, r 's LOS connection is 'guessed' from its components, p and q . There may be subtler geometric approximations or calculations to be applied but this goes beyond the scope of this paper!

The second constraint places an upper bound on the cluster size. It will be undesirable for a single node to effectively act as a concentration point for too many subscribers. There may be a technical limit (number of directional transmitters/receivers, etc.) and almost certainly reliability issues to be considered. On the basis of the initial weights and their reductive combination, it will be a simple exercise to prohibit the clustering of two nodes with a combined weight greater than some determined maximum, \underline{w} . The step **[R]** will only be permitted if $w_p + w_q \leq \underline{w}$. This in turn introduces a further terminating condition. The reduction process continues until either $n\mathfrak{C}$ reaches its lower bound or no further valid pairings are possible. Thirdly, and for similar technical and/or operational reasons, it may be necessary to restrict the area covered by a cluster. At each stage, we can

insist that $d_{pq} \leq d$, a determined maximum, for [R] to proceed. This provides a final stopping condition, the reduction being halted if all remaining node pairs are too far apart.

3.3 Large Scale Distribution

Of course we may not be able to identify individual subscribers or there may be too many to allocate initial nodes with unit weights. An alternative description of the subscriber density is given by a subscriber distribution. The situation is illustrated in Figure 2.




Figure 2. Subscriber Density

Here, the geographic area in question of dimensions $A \times B$ is divided into an $a \times b$ grid with a value z_{ab} giving the projected number of subscribers for the cell (a, b) ($1 \leq a \leq a$, $1 \leq b \leq b$). Placing a single node at the centre of each cell gives $n = ab$ initial nodes with

$$x_i = \frac{A}{a} \left(a - \frac{1}{2} \right), \quad (6)$$

$$y_i = \frac{B}{b} \left(b - \frac{1}{2} \right) \quad (7)$$

and $w_i = z_{ab} \quad (8)$

for $1 \leq i \leq n$, counting (arbitrarily) west-to-east and south-to-north, where $\mathbf{a} = ((i-1) \text{ modulo } a) + 1$ and $\mathbf{b} = \lfloor (i-1)/a \rfloor + 1$ ($\lfloor \cdot \rfloor$ being the integer part). Subsequently the reduction proceeds as described previously, but with a minor refinement. As in this entirely regular initial configuration, many node pairs will be equidistant, weights may be used to break the tie. Amongst all nodes separated by equal distance, p and q are chosen so that $w_p + w_q \leq w_{ic} + w_{jc} (d_{iq} \leq d_{ij}; 1 \leq i, j \leq n)$. Low density nodes are combined early in the process to avoid being left isolated when the reduction terminates and nodes with zero weights are relocated entirely to non-zero

neighbours by [R]. A special, simplified calculation is necessary when combining two nodes of zero weight.

$$x_r = \frac{x_p + x_q}{2}, \quad (9)$$

$$y_r = \frac{y_p + y_q}{2} \quad (10)$$

and $w_r = 0. \quad (11)$

Including cells of zero weight in the reduction process gives a solution that allows for growth in these new areas. An obvious variant would be, for further simplicity, not to allocate nodes to such cells; this removes the special case above.

3.4 Forced Seeding Points

Finally in attempting to deal with real-world problems, it may be necessary to fix a number (m) of seeding points at known nodes. This could be appropriate on the basis of their proximity to the backbone carrier or having key equipment already in place. An obvious requirement would be LOS between the each fixed seed and the PD. The optimization should proceed as normal on the remaining nodes, accepting the set seeds as unalterable.

In fact this variation is achieved quite easily by initially setting $w_{m+1} = \infty$ for all known seeds $m+1 \leq i \leq m$. The reduction process will draw all paired nodes towards these seeds while allowing independent pairing to proceed normally. A modification to the maximum weight condition is necessary of course since all seeded combinations will exceed the set limit. It is simply necessary to define a second variable, representing the combined cluster weight *excluding* the seed node. Two seed nodes may *never* be paired.

4. RESULTS

Test samples have been drawn from two sources. Firstly, randomly generated subscriber distributions of various sizes and, secondly, known network distributions, both past and current. One aspect of testing is difficult in both cases. Whereas, for small networks, it is possible to generate the 'true' optimum by exhaustive search, and thus make an accurate comparison, this is not a practical approach for larger cases and subtler approaches become necessary.

For small values of n (up to about 20), where exhaustive search comparison is possible, results are good with the results from the reduction algorithm largely matching those from exact calculation. Our heuristic produces the optimum solution in around 94% of cases with the mean excess less than 1%. A few larger (20 $\leq n \leq 40$) examples with particular characteristics (symmetry, redundancy, etc.) have had these features exploited but otherwise tested exhaustively. Again,

results are good – similar to the smaller cases - but more testing is necessary. For larger values of n , we must resort to theoretical analysis of the method itself. The work in [2] offers a foundation and further investigation is continuing.

Cost calculations may be based upon total link length connecting relocated nodes or number of relays for a variety of constraint combinations. The number of relays reflects the true cost of the network; however, link length offers a more accurate measure of deviation from the optimum. Finally, the complexity of the process, being a linear sequence of matrix searches, is bounded by $O(n^3)$.

5. EXTENSIONS

There is a natural extension to this algorithm that cannot be taken further in a paper of this length. If we regard the configuration produced by the reduction/optimization process as an *initial* solution then it may be possible to improve upon it by a process of local search. Such techniques are well-established [3] and may be adapted easily to this problem. There are three distinct places in which perturbations to the current solution could be applied: the choice of relay nodes, the inter-relay links and the node-relay links. Options are suggested in Figure 3.




Figure 3. Relay/connection perturbation

6. CONCLUSION

Providing broadband access to subscribers through wireless links is a prominent technology. Furthermore, the relay approach to the problem of connectivity is attractive one. The problem of determining the optimal such connection strategy is thus very real. This paper has outlined a method for determining the seeds for this network and the interconnection among seeds and subscribers. It is flexible: permitting a range of real-world constraints to be applied. This is work in progress but the results are encouraging. The reduction/optimization approach combined with the local search extension may reasonably be expected to provide even better results than those achieved to date.

REFERENCES

- [1] Tim Fowler, "Mesh Networks for Broadband Access". IEE Review, January 2001, pp17-22, 2001.
- [2] V.M. Grout, "Optimisation Techniques for Telecommunication Networks". PhD Thesis, Plymouth Polytechnic, UK, 1988.
- [3] D. S. Johnson & L. A. McGeoch, "The Traveling Salesman Problem: A Case Study in Local Optimization", in 'Local Search in Combinatorial Optimization' E. H. L. Aarts and J. K. Lenstra (editors), John Wiley and Sons, Ltd., pp. 215-310, 1997.