6

The Design of Forward Looking Cost Models for Local Exchange Telecommunications Networks

William W. Sharkey[1]
Federal Communications Commission

Abstract - In the last few years great strides have been made in the development of algorithms that design telephone networks. As computational capabilities improve it is possible to produce better results, both from an engineering and an economic standpoint. This paper considers the design issues that cost model developers have addressed successfully. Many of these issues are illustrated by a detailed description of a model developed by FCC staff, known as the Hybrid Cost Proxy Model (HCPM). HCPM is capable of utilizing very precise customer location data. From these data, the model uses clustering algorithms to identify serving areas that satisfy appropriate engineering constraints. Within each serving area, the model uses a modified minimum-cost spanning tree algorithm to connect actual customer locations to a serving area interface. The same tree algorithm connects each interface point to a switch. Within each path, the model performs intensive integer searches to find the cost minimizing, yet engineering-feasible, choice of technology and electronics for that path. The result is a-low cost, feasible network plan that gives an appropriate estimate of the forward looking cost of providing wireline telephone service to a particular area. This estimate should prove particularly useful in the ongoing debate about the size and make-up of the Universal Service Fund, and for other regulatory purposes such as the pricing of interconnection and unbundled network elements.

Engineering process (cost proxy) models have been developed in recent years as an alternative to more traditional econometric and accounting approaches to cost measurement. Because econometric models rely on assumptions of (smooth) functional forms, engineering process models offer a more detailed view of cost structures than is possible using econometric data. In addition, engineering models are better suited for modeling forward looking (long-run) costs because they rely much less on historical data than econometric models.

In the telecommunications industry, cost proxy modeling can play a particularly significant role for three reasons. First, the very rapid technological change in the industry compounds the standard econometric difficulties in using historical data to estimate a forward looking cost function. Second, many applications of cost studies in telecommunications require highly disaggregated levels of detail, which can capture regional variations in cost and the costs of specific network components. Finally, a growing awareness of the benefits of deregulation in the public utility sector industries generally, has given rise to a demand for a new set of tools that can be used to promote competition and advance the pace of deregulation in the industry.

The deregulatory goals for telecommunications are explicitly set out in the Telecommunications Act of 1996,[2] in which Congress sought to establish "a pro-competitive, de-regulatory national policy framework" for the U.S. telecommunications industry. In the two full years following this act, the Federal Communications Commission has undertaken proceedings on universal service,[3] interstate access charge reform,[4] and local exchange competition[5] to overhaul its current regulations in light of the 1996 Act. In these proceedings, the Commission has examined, in varying degrees, the use of forward looking economic cost methodologies as a basis for determining universal service support levels, cost-based access charges, and pricing for interconnection and unbundled network elements. The 1996 Act has fundamentally changed telecommunications regulation by replacing the framework of government-recognized monopolies with one in which federal and state governments work in tandem to promote efficient competition and to remove entry barriers and regulations that protect monopolies. The 1996 Act, when fully implemented, should greatly reduce the legal, regulatory, economic, and operational barriers to entry in the local exchange and exchange access market, while preserving and advancing enhanced universal service goals.

The local competition provisions of the Act confer three fundamental rights on potential competitors to incumbent local exchange carriers (LECs): the right to interconnect with other carriers' networks at rates based on cost, the right to obtain unbundled network elements at cost-based rates, and the right to obtain an incumbent LEC's retail services at wholesale discounts in order to resell those services.[6] The Act also requires a fundamental restructuring of current regulatory mechanisms for funding universal service goals. For this purpose, the FCC must first define the services to be supported by federal universal service mechanisms, and then determine a mechanism to estimate the cost of such services in a manner that is "explicit and sufficient to preserve and advance universal service."[7] In its recently initiated access reform proceeding, the Commission also seeks to reform its system of interstate access charges to make it compatible with the competitive paradigm established by the 1996 Act and with state actions to open local networks to competition.[8]

Forward looking economic computer-based cost models could enable regulatory authorities to estimate the forward looking cost of network facilities and services without having to rely on detailed cost studies, prepared by incumbent LRCs, that would otherwise be necessary. In addition, a publicly available cost proxy model could be useful to regulators by providing an independent check on the accuracy of incumbent LEC cost studies. During the course of the model development process, several industry-sponsored models were submitted to the FCC for evaluation. These include the Benchmark Cost Proxy Model (BCPM) sponsored by US West, Sprint and Bell South and the HAI model sponsored by AT&T and MCI. Simultaneously, staff members of the FCC worked on an internal model known as the Hybrid Cost Proxy Model (HCPM), which incorporated elements of both of the industry models in addition to a set of new loop design and clustering algorithms developed internally. In October 1998 the Commission adopted a synthesis model consisting of the HCPM clustering and loop design modules in combination with HAI switching, transport and expense modules as a platform for the forward looking mechanism for determining high-cost support for non-rural LECs.[9] Subsequently, the FCC initiated an investigation of appropriate input values for use in the model. In May 1999, tentative input values were released for public comment. A final order recommending the use of the model and a set of input values for the purposes of determining high-cost support is expected in September 1999.

1. CRITERIA FOR EVALUATING THE UTILITY OF ECONOMIC COST MODELS

This section briefly discusses the criteria the FCC uses for evaluating forward looking economic cost models.

1.1 Use of Forward Looking Economic Cost as a Basis for Pricing

In dynamic, competitive markets, firms base their actions on the relationship between market-determined prices and forward looking economic costs. Forward looking economic costs are the costs that would be incurred if a new element or service were provided, or that could be avoided if an existing element or service were not provided, assuming that all input choices of the firm can be freely varied. This is often referred to as long-run economic cost. This "long-run" approach ensures that rates recover not only those operating costs that vary in the short run, but also fixed investment costs that, while not variable in the short term, are necessary inputs directly attributable to providing the element or service. If market prices exceed forward looking costs, new competitors will efficiently enter the market and bring pressure to bear on prices. If forward looking economic costs exceed market prices, new competitors will not enter and incumbent firms may decide to exit. These voluntary actions by firms produce efficient resource allocation by adjusting price and output until the value to consumers of additional output is just equal to the cost of the resources required to produce it. In contrast, basing prices on embedded costs would fail to establish the critical link between economic production costs and market prices, and would be inconsistent with the goal of efficient competition.[10] Pricing based on forward looking costs enables efficient providers to cover their costs and make new investments, while facilitating efficient market entry by potential competitors.

1.2 Use of Proxy Models for Multiple Objectives

For the purposes of determining universal service support levels, the FCC determined that a cost proxy model should, at a minimum, be able to estimate the full stand-alone cost of the minimum set of network elements capable of delivering traditional voice telecommunications service and narrowband data services, at currently acceptable quality levels, to customers of the public switched network and to private line users. Because incumbent local exchange carriers may choose to construct network facilities capable of providing services that require higher transmission speeds ("broadband" services), it is also necessary that a cost proxy model be able to model a network capable of providing these services if the model is to be used for the purpose of setting prices for unbundled network elements.

1.3 Consistency with Independent Evidence

It may be possible to obtain independent estimates of the costs of some network elements as a check on the validity of model estimates. For example, it may be feasible to compare estimates of loop costs with competitive bids for installing loops or the costs that cable networks incur in installing similar networks. Econometric studies might also provide a check on model results.

1.4 Potential for Independent Evaluation

The algorithms in a proxy model should be clearly identified and explained so they can be independently evaluated by state or federal regulators. It must be recognized, however, that this criterion may be in partial conflict with the overriding goal of obtaining accurate cost estimates. For example, a model that utilized only publicly available information would allow full independent evaluation, but might be less accurate than a model that used proprietary information (such as vendor pricing data).

1.5 Flexibility

Some jurisdictions may possess detailed information about important model inputs (such as discount prices offered by switch vendors) that model designers could only estimate. In addition, these jurisdictions may possess detailed information on local conditions, such as zoning restrictions and labor rates, that they may wish to add as inputs to a model. As a general rule, cost proxy models should permit parties to utilize such information where available.

2. UNDERLYING STRUCTURE OF COMPUTER-BASED COST PROXY MODELS

An economic cost proxy model for estimating the cost of network elements starts with an engineering model of the physical local exchange network, and then makes a detailed set of assumptions about input prices and other factors. Such models estimate the total monthly cost for each network element. This section examines both model design and the use of variable input factors for network investment, capital expenses, operating expenses, and common costs.

2.1 Preliminary Modeling Issues

2.1.1 Existing Wire Center Approach

Each of the models submitted to the FCC for evaluation was based on an assumption that wire centers will be placed at the incumbent LEC's current wire center locations. Subject to this constraint, all remaining network facilities are assumed to be provided using the most efficient technology currently in use. The constraint to use existing wire center locations is not fully consistent with a forward looking cost methodology, and it should be recognized that over time, an existing wire center model may become less representative of actual conditions faced by new entrants and incumbents. For example, after existing wire center locations were chosen by incumbent LECs, larger capacity switches and fiber/digital loop carrier technologies became available. Both of these factors have significantly altered the fundamental trade-off between switching and transmission in the design of an optimal communications network. Because of ongoing advances in technology, facilities-based new entrants and incumbent LECs may in the future choose a much different network topology that will result in different forward looking costs than today's network. A closely related issue is whether the placement of remote switching units should be restricted to existing wire center locations, and if the models should assume that every wire center includes a (host or remote) switch. Similarly, in the future, wireless technologies may be capable of providing narrowband telecommunication services at a lower cost than wireline technologies.

2.1.2 Specification of Demand

An accurate estimate of the cost of serving a wire center or a serving area within a wire center depends on a reliable forecast of customer demand patterns within the area, and the number of residential and business lines. Proxy models rely on census data to determine residential demand. However, because census data do not report the number of business lines,[11] model designers must use indirect methods to estimate business demand. The potential for error in estimating business and residential demand creates certain difficulties. First, as noted below, fill factors or utilization rates may be expected to vary between business and residential lines.[12] Second, loop lengths are typically shorter for business lines than for residential lines. Thus, unless the differences in costs associated with different fill factors for business and residential areas happen to offset exactly the differences in costs associated with differences in loop lengths, the cost of serving an area will depend on the ratio of business to residential lines. An understanding of the magnitude of these competing effects, however, requires an accurate estimate of the number of business and residential lines in a particular area.[13]

The HAI model and the FCC synthesis model incorporate access line demand data from the Operating Data Reports, ARMIS 43-08, submitted to the FCC annually by all Tier 1 LECs.[14] These models incorporate data on the number of: 1) residential access lines, both analog and digital; 2) business access lines, which include analog single lines and multi-line analog and digital lines, PBX trunks, Centrex trunks, hotel and motel long-distance trunks, and multi-line semi-public lines; and 3) special access lines. The number of residential lines in each Census block is computed by multiplying the number of households in a block by the ratio of total residential lines, as reported by ARMIS, to the total number of households in a study area. The number of business lines in a wire center is determined by multiplying the number of employees, as reported by Dun and Bradstreet, by the ratio of business lines to employees, as determined from ARMIS data. Some refinements to this process have been made that take account of the different demands for telephone use per employee. For example, service industry demand for telephone service is most likely greater than demand in the manufacturing sector.

2.2 Loop Plant

The largest portion of a network's investment consists of its investment in loop plant. It is therefore vitally important that models estimate accurately the cost of loop plant sufficient to satisfy demand. "Loop plant" consists of all network facilities, including wires, telephone poles or conduits, drops, etc., connecting the end office switch and customers' premises.

All cost proxy models include assumptions regarding feeder and distribution utilization rates (also called "fill factors"). In each model, lower utilization rates increase total loop investment because the increase in capacity associated with lower fill factors increases the amount of loop plant used to deliver telecommunication services. Thus, the choice of fill factor can have a significant effect on total cost. While all models allow user inputs for these quantities, it is not obvious what levels should be used as inputs. In a well-engineered network, it is necessary to include unused capacity when constructing loop plant to reduce the likelihood of outages in the case of breakages and to account for growth in demand. Furthermore, optimal fill factors should vary over the service life of the plant, increasing as demand grows until more plant is put into service.

Fill factors may also differ between business and residential markets. In residential markets, LECs traditionally place multiple wire pairs per home in order to be able to provide a second or third line to premises without incurring construction costs. Thus, fill factors that are less than 50 percent may be reasonable for residential markets. In business-dominated wire centers, the rate of utilization depends on the proportion of businesses using Centrex service rather than PBX terminal equipment, because PBXs serve to concentrate traffic between the customer and the central office. Customers using PBX equipment therefore require fewer lines than customers using Centrex service. Depending on the relative use of Centrex to PBX equipment, and LECs' plans for marketing Centrex services, business fill rates could be either lower or higher than residential fill rates.

The forward looking cost of installing loop plant includes both the cost of cable and the cost of building or obtaining access to structures that support the loop plant. With respect to cable investments, all three models use default input prices to estimate the cost of loop plant, but allow users to specify different input prices. Structures for cable plant consist of aerial, buried, and underground (i.e., cable in conduit) facilities, and the plant mix assumptions used by a proxy model can have a significant effect on estimated model costs. A crucial variable is the proportion of plant that is installed in new developments (where installation costs are relatively low) to plant installed for existing business and residential users. Different assumptions about the sharing of structure costs can also have a significant effect on estimated model costs.

One issue that has been raised with regard to forward looking cost studies generally is the treatment of existing sunk investments in structure facilities. While the use of existing structure investments in a forward looking cost model would superficially appear to be related to the fixed wire center assumption that most models have adopted, there are important reasons for rejecting the use of embedded costs for structure investments. For example, the deployment of feeder plant depends critically on the trade-offs in the costs of cable and electronics, which can be accurately described only in a fully forward looking context.

2.3 Switching Investment

After determining the number of lines assigned to a wire center, a proxy model must determine the number and size of the switches to be placed in these wire centers. The HAI model and the FCC synthesis model determine the investment in switches and interoffice transport based on the number of lines and DEMs, along with Bellcore assumptions on busy hour call attempts. These models use data from a McGraw-Hill study of the central office equipment market to derive average per-line prices for switching investment, including separate costs for the buildings, land, and other inputs to determine investment in switching. Since switch vendors typically grant carriers substantial discounts when selling switches, and require carriers to sign nondisclosure covenants that they will keep confidential the actual prices for which switches are sold, proxy models must rely on indirect evidence about the magnitude of such discounts, which can be expected to vary with the size of the purchasing carrier.

Another important issue in the modeling of switching costs is the ratio of traffic-sensitive to non-traffic sensitive cost in a switch. This ratio may be specific to the particular switches designed by different vendors. If the non-traffic sensitive costs are not constant across all switches, one would expect, since switches depreciate relatively quickly, that cost-minimizing carriers would install switches whose costs are largely traffic or non-traffic sensitive depending on the type of traffic that will be switched in an area. For example, in an area that switches a large amount of traffic with long holding times, it may be cost minimizing to install a switch whose costs are largely non-traffic sensitive.

2.4 Expenses

Cost proxy models are deisgned to produce an estimate of the annual or monthly cost of producing a set of services, including local exchange and access to intra- or inter-state toll services. Annualized cost consists of the sum of the return on equity, taxes, interest, depreciation, network operations and support expense, customer operations, and corporate overhead. The expense side of a model can have a significant effect on the final cost estimate.

2.4.1 Capital Expenses

Capital expenses are computed as the sum of a return on investment, taxes, and depreciation. In the HAI model, the return on investment is equal to the net investment base (gross investment minus accumulated depreciation) multiplied by a rate of return equal to a weighted average of the cost of equity and the cost of debt, with weights equal to the corresponding percentages of equity and debt in total investment. Taxes in the model are equal to the product of the net investment base, the percentage return on equity, the percentage share of equity, and a "tax gross up" factor determined by the following equation:

$$Taxes = \%Equity \times \%Return\ on\ Equity \times Investment\ Base \times \frac{Composite\ Tax\ Rate}{(1 - Composite\ Tax\ Rate)}.$$

For each category of plant, the capital cost is computed for each year of the economic life of the plant and the resulting stream of returns is "levelized" through a net present value calculation to give a constant annual cost of capital for that category of investment.[15] Aggregate capital costs are then computed as the sum of the capital costs for each category of plant.

The second component of a capital expense computation is a model's choice of depreciation rates. Since higher levels of depreciation lead to lower levels of investment base, and consequently lower annual expenses associated with return on investment and income taxes, changes in annual capital costs caused by changes in depreciation rates will automatically be mitigated to some extent by offsetting changes in return and taxes.

Depreciation schedules specified in a forward looking proxy model should be based on forward looking costing principles and should reflect projected economic lives of investments rather than physical plant lives. For the reasons described above in Section 2.2, the reported plant lives for loop-plant structures, such as conduit, manholes, and poles, are particularly important. Because of the relatively large investment needed to construct such facilities, inaccurate estimation of the expected economic lives of such facilities may result in a significant under- or overestimation of the forward looking costs of these facilities.

2.4.2 Operating Expenses

All proxy models use annual cost factors to calculate non-capital-related expenses. An annual cost factor is the ratio of expense booked to a specific account and the gross investment booked to the same account. Typically, the expense associated with investment is the product of the model-generated investment and the associated annual cost factor. Annual cost factors are used by models, as well as by companies in individual cost studies, because methods for developing forward looking expenses are complex and contentious.

2.4.3 Joint and Common Costs

If proxy models are used to estimate forward looking economic costs, the question of joint and common costs must be addressed. In the case of unbundled network element pricing, costs that are common to a set of network elements can be allocated among the individual elements in that set. For example, shared maintenance facilities could be allocated to the elements that benefit from those facilities. Common costs also include costs incurred by the firm's operations as a whole. Given these joint and common costs, setting prices for individual network elements based on forward looking incremental costs alone would not recover the full forward looking cost of the network. Consequently, in order to recover the full forward looking cost, a reasonable measure of joint and common costs should be included in the prices for interconnection and unbundled network elements.

3. OVERVIEW OF THE HYBRID COST PROXY MODEL[16]

The HCPM consists of two independent modules: a customer location module and a loop design module. The customer location module first groups individual geographic locations of telephone customers into clusters, based on engineering considerations. Next, the customer location module determines a grid and microgrid overlay for each cluster, and places each customer location into the correct microgrid cell. The loop design module determines the total investment required for an optimal distribution and feeder network by building loop plant to the designated customer locations represented by populated microgrid cells. The number of microgrids in a grid can vary from 4 to 2500. When used with a source of geocoded customer locations and a maximum copper reach of 18,000 feet, a uniform microgrid size of 360 feet can be maintained. All customer locations can therefore be determined with an error of not more than several hundred feet. All modules are written in high-level programming languages, and compiled versions can be supported on a number of computing environments. In the following sections, a more detailed description of each of the modules is presented.

3.1 The Customer Location Module

This section describes a method of modeling customer location based on cluster analysis. This approach is designed to accept as input a set of geocoded locations for every customer. It then performs a "rasterization" procedure that assigns customers to small microgrid cells.[17] A cluster algorithm then groups the set of raster cells into natural clusters. Finally, a square grid is constructed on top of every cluster, and all customer locations in the cluster are assigned to a microgrid cell for further processing by the loop design module. The cluster module can also accept Census block-level data as an input. In this case, every block that is larger than a raster cell is broken up into smaller blocks, each no larger than a raster cell, and the population of the original block is distributed uniformly among the new blocks.

One of the primary tasks faced by the HCPM is to identify clusters of customer locations. Each customer in a particular cluster, or serving area, will then be connected to the feeder system through a single interface, the serving area interface or SAI.

The clustering task is difficult because both engineering constraints and the general pattern of customer locations must be considered. There are two main engineering constraints. First, a serving area is limited to a certain number of lines by the capacity of a remote terminal. Second, a serving area is limited to certain geographic dimensions by current technology, because as distance increases beyond a critical value, service quality is degraded.

Given the engineering constraints, one could create feasible serving areas by simply placing a grid containing cells of an appropriate dimension over the entire wire center. For this to be a cost-effective approach, however, customers would have to be located in a relatively uniform pattern across the entire wire center. But, people do not tend to live that way. They tend to live clumped together in towns and communities. This tendency creates areas of varying population density throughout the wire center. Under these conditions, a gridding approach may divide a natural grouping of customers into different serving areas when a single serving area would be more cost-effective.

The objective of a clustering algorithm is to create the proper number of feasible serving areas. Unfortunately, this is not a well-defined objective, because of the existence of both fixed and variable costs associated with each additional serving area. A fixed cost gives a clear incentive to create a small number of large clusters, rather than a larger number of smaller clusters. On the other hand, with fewer clusters, the average distance of a customer from a central point of a cluster, and consequently the variable costs associated with cable and structures, will be larger. In moderate- to high-density areas, it is not clear, *a priori*, what number of clusters will embody an optimal trade-off between these fixed and variable costs. However, in low-density rural areas, it is likely that fixed costs will be the most significant cost driver. Consequently, a clustering algorithm that generates the smallest number of clusters should perform well in rural areas.

While statisticians have studied a wide variety of clustering algorithms,[18] there are two basic approaches to clustering: the agglomerative or bottom-up approach, and the divisive or top-down approach. Each approach starts with an initial state where each customer location belongs to a particular cluster. The initial state is then improved upon according to some rule until no more improvements can be made. The clustering module for the HCPM contains three alternative algorithms that represent implementations of both of the above approaches.

In the initial state for the default divisive approach, each location belongs to a single parent cluster. This initial state is improved upon by dividing the parent cluster into a new parent cluster and a child cluster. This step increases the total number of clusters by one. The improvement step is repeated until every cluster is feasible from an engineering standpoint.[19] A child cluster is created from the parent cluster by choosing the customer location furthest from the parent's line-weighted center as an initial child cluster member. Then, customer locations that are closer to the center of the child cluster than they are to the center of the parent cluster are reassigned in an iterative manner, recalculating the cluster centers at each step. Customer locations are added to the child cluster until it is full, i.e., until no more locations can be added without violating engineering constraints.

Alternatively, in the initial state for all agglomerative approaches, each location belongs to its own unique cluster. This initial state is improved upon by merging the two closest clusters together, reducing the total number of clusters by one. The improvement step is repeated until merging is no longer feasible from an engineering standpoint. The clustering module contains two agglomerative algorithms that differ only in the way in which they measure the distance between clusters. In the standard agglomerative algorithm, distance is measured from the line-weighted center of one cluster to the line-weighted center of another. In the nearest-neighbor algorithm, distance is measured from the two customer locations, one in each cluster, that are closest together.

Once one of the clustering algorithms has been run, it has been found that the initial result can generally be improved by reassigning certain customer locations to different clusters. The clustering module contains two optimization routines that perform these reassignments. As a final step, the cluster module computes potential locations for either one SAI or pair of SAIs. The location for a single SAI is simply the line-weighted center of the cluster. The locations for a pair of SAIs are determined by dividing each cluster into a parent and child. The module then reports the line-weighted centers of the parent and child as potential locations of a pair of SAIs. The actual number of SAIs used is determined within the loop design module.

Output from the cluster module is illustrated in Figure 1. Each cluster, which represents a single feasible serving area, is represented by a set of customer locations connected to the cluster center by a straight line.

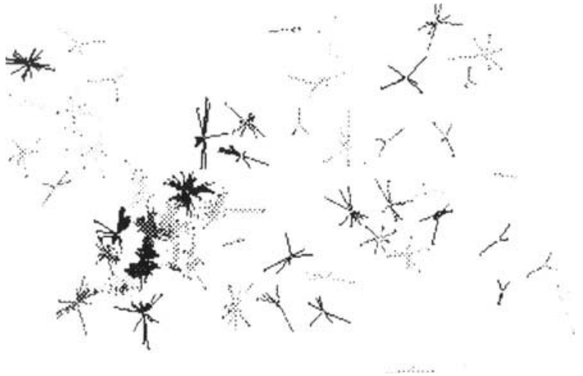

Figure 1: Clusters for the Wire Center GNSNCOMA

The final step in the customer location module is to convert the data from the cluster algorithm into a form that can be utilized by the loop design module. For wire centers that are sufficiently small, it would be possible to build plant to the exact customer locations determined by the geocoded data as processed by the cluster module. For larger wire centers, which may have 20,000 or more individual customer locations, it would be extremely time consuming to build distribution plant directly to each individual location. As in the clustering module, an acceptable compromise between absolute accuracy and reasonable computing time can be achieved by defining a grid on top of every cluster and assigning individual customer locations to microgrid cells within the grid. Customers within each microgrid cell are assumed to be uniformly distributed within the cell. If multiple customer locations fall in a single microgrid cell, the cell is divided into lots, as explained in the next section. Loop plant can therefore be designed specifically to reach only populated microgrid cells and the individual customer lots within each microgrid. For large wire centers, the number of populated microgrid cells will be much less than the number of customer locations, even when a relatively small microgrid size is specified. Therefore, with this approach it is possible to place an upper bound on computing time, while simultaneously placing a bound on the maximum possible error in locating any individual customer.

3.2 Loop Design Algorithms

The customer location module will report, for each cluster, the bottom-left and top-right coordinates of the overlay grid, the number of microgrid cells, the number of lines associated with each cell, and coordinates for the possible locations for the SAIs that could serve the cluster. Given the inputs from the customer location module, the logic of the loop design module is straightforward. Within every microgrid with non-zero population, customers are assumed to be distributed uniformly. Each populated microgrid is divided into a sufficient number of equal-sized lots and distribution cable is placed to connect every lot. These populated microgrids are then connected to the nearest concentration point (SAI) by further distribution plant. During this phase of the loop design algorithms, the heterogeneity of microgrid populations, and the locations of populated microgrids are explicitly accounted for. Finally, the SAIs are connected to the central office by feeder cable. On every link of the feeder and distribution network, the number of copper or fiber lines and the corresponding number of cables are explicitly computed. The total cost of the loop plant is the sum of the costs incurred on every link.

Distribution consists of all outside plant between a customer location and the nearest SAI. Distribution plant consists of backbone and branching cable, where branching cable is closer to the customer location. Feeder consists of all outside plant connecting the central office main distribution frame, or fiber distribution frame, to each of the SAIs. Feeder cable consists of main feeder, subfeeder and sub-subfeeder routes.

3.2.1 Distribution Plant Design

The distribution portion of the loop design module determines the cost of distribution plant for each cluster in isolation (ignoring information from all neighboring clusters). The algorithms described in the following sections compute the cost of all plant that is required to connect each customer within the cluster to the nearest SAI.

Each microgrid is divided into lots based on microgrid population. Distribution cable is built to touch every lot in the cell. Backbone cables, which connect cells to the SAI, are assumed to run horizontally and branching cables within a cell are assumed to run vertically along the grid lines of the cell. Branching cable is assumed to follow every other vertical lot boundary. Drop cable is designed to serve groups of four properties whenever possible.

Two algorithms are used to determine the correct amount of cable and structures that are necessary to connect each microgrid to the nearest SAI. In a fully optimizing mode, the model computes the cost of distribution plant for all clusters using both approaches, and selects the approach giving the lower cost.[20]

The first algorithm is most appropriate in densely populated clusters, in which the proportion of populated micrograms to total microgrids is relatively large. Backbone cables run along every other cell boundary and connect with the distribution plant within a cell at various points, as illustrated in Figure 2. The second algorithm generally gives a more efficient distribution network for clusters with a lower population density, where the number of populated microgrids is smaller. In this case, the construction of an optimal distribution network within a cluster is closely related to the problem of constructing an optimal feeder network for the entire wire center, and the same algorithm is used to provide a solution.

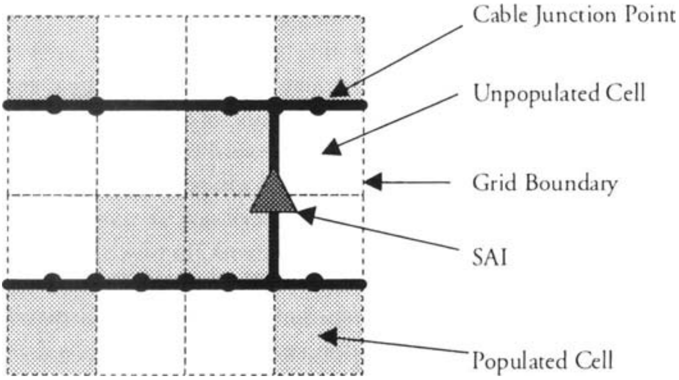

Figure 2: Connection of Cells to the Closest SAI

3.2.2 Feeder Plant Design

In previous versions of the HCPM and in other models of the local exchange, feeder plant was deployed in a "pine tree" network in which four main feeder routes emanate from the central office along east-west and north-south routes. Subfeeder routes perpendicular to the main feeder routes were then used to bring the feeder system closer to individual SAIs. This design proved to be highly efficient in terms of creating opportunities for the sharing of structure costs among feeder cables serving different SAIs. Lower structure costs made possible by increased sharing, however, came at the expense of longer feeder routes and correspondingly higher cable costs. In order to balance these two opposing tendencies, the HCPM examined a large number of possible feeder systems having different number of subfeeder routes, and chose the configuration giving the lowest cost.

The current version of HCPM uses a variant of an explicit optimization algorithm, discovered by Prim in 1957, to determine the trade-off between structures and cable costs. [21] This algorithm is based on some well known mathematical principles of network design based on techniques of discrete mathematics and graph theory. An abstract network consists of a single "supplier" node, a set of customer nodes, a cost function specifying the cost of connecting any two nodes, and a set of pair-wise traffic demands between any two nodes. In the application of the Prim algorithm to the feeder network, the supplier node is the central office for a given wire center, and the customer nodes are the remote terminals, or SAIs, that define the interface points between the feeder and distribution portions of the network. The algorithm can also be applied to determine cable routes for distribution networks within a cluster, as noted previously. In this case, the supplier node is an SAI within a cluster, and the customer nodes represent individual subscriber locations that are to be connected to that SAI.

In both the feeder and distribution portions of the network, the objective of the telecommunications engineer is to minimize the cost of connecting each customer node to the supplier node. While in general this is an extremely difficult problem to solve, there are several special cases in which efficient algorithms exist that define a fully optimal network solution. One special case of interest is the construction of a "minimum-distance spanning tree network" in which the sole objective is to minimize the aggregate length of communications links within the network. Such a network would be approximately optimal when traffic demands are sufficiently low that the actual cost of each link in the network is largely determined by the cost of structures (which depend only on distance).

A minimum-distance network can be constructed using the Prim algorithm in the following way. Beginning with a network consisting only of the supplier, find the nearest customer node that is not yet attached to the network and attach it. The network then consists of the supplier and one customer. The algorithm proceeds step-by-step in attaching customer nodes to the network on the basis of minimum distance. At any point in the algorithm, it chooses from the set of all nodes that have not yet been attached, a node that is closest to some node in the existing network. Prim demonstrated that this simple algorithm necessarily leads to a minimum-distance network.[22] In other words, when the algorithm is completed, there is no possible way to reconfigure the network so as to lower the aggregate distance of all links in the network.

As long as structure costs are significantly larger than cable costs, the original Prim algorithm provides a satisfactory solution. In the design of both feeder and distribution networks, however, a minimum-distance spanning tree network is not generally optimal.[23] While it minimizes the total distance of all links in the network, it does not minimize the distance between any particular node and the supplier. For example, if there is significant demand at a particular remote terminal for access lines to the central office, then the actual cost on the network between this node and the central office may need to be accounted for in the minimum-distance optimization problem. There are no simple algorithms that can be applied to give a completely optimal solution to the general problem. (A "star" network, which minimizes the cost of connecting each node to the central office, would not be optimal because it does not take advantage of potential sharing of structure costs in the network.) However, it is possible to modify the Prim algorithm to take account of the effects of traffic in the network on total cost, and generally to improve the performance of the algorithm computationally.

The HCPM makes two fundamental modifications to the Prim. First, the model creates a set of potential junction points that follow the location of the main east-west and north-south feeder routes in previous versions of the model. The algorithm permits, but does not require, these potential junction points to be used in creating the feeder network. Junction points create additional opportunities for the sharing of structure costs, and in some circumstances they can also reduce the distance between a terminal node and the central office.

The second modification of the Prim algorithm is in the rule used to attach new nodes to the network. Rather than minimizing the distance from an unattached node to the existing network, the algorithm minimizes the total cost of attaching an unattached node, and of constructing all of the lines required to carry traffic from that node back to the central office. A heuristic description of the algorithm is given below.[24]

Step 1: Begin with a network consisting of the central office alone.

Step 2: From the set of unattached nodes, find the node for which the average cost per line, including the cost of structures, cable and terminal electronics, is lowest for connecting that node to the existing network.

.

.

.

Step k: At any step in the algorithm, choose from the set of unattached nodes the node for which the average cost is lowest for connecting that node to the existing network. This cost will depend on the particular node in the existing network that is selected for connection, and will include the structure cost of connecting to that node as well as the incremental cable cost of carrying traffic from the new node to the central office along the currently existing network.

The algorithm terminates when all nodes have been attached. Unlike the Prim algorithm, it may be possible to lower total network cost by rearranging some of the links in the network after the algorithm terminates. However, the general optimization problem is computationally intractable, while the above algorithm is highly efficient. FCC staff have found that this modified Prim algorithm leads to lower feeder cost estimates than the unmodified Prim algorithm and the more traditional pine tree feeder route designs. Furthermore, the modified Prim algorithm provides a good approximation of the way in which real world engineers are likely to design the feeder network because the network grows naturally from the central office, by adding new nodes on the basis of minimum attachment cost as new communities are established.

In the construction of the feeder network, the HCPM allows the user to determine whether to use airline distances between nodes or rectilinear distances. The model also applies a road factor to all distance computations in the feeder network. This road factor is intended to convert distances determined by the distance function into actual route distances, which must account for the existing road network and other terrain factors. In principle, the road factor should be determined empirically for each region of the country by comparing actual feeder route distance to model distance computation. Clearly, a different factor should be applied to airline distance than to rectilinear distance computations. Some empirical evidence on the appropriate value for the road factor is given in Love et al. (1988).

Figure 3 illustrates the operation of the Prim algorithm and the HCPM modifications to it. A set of 15 locations representing one wire center (source) and 14 customer nodes were randomly generated. Each customer node was assigned a demand of one unit. A highly simplified cost function of the form *Cost per link* = $(P_S + T * P_C) * distance$, where P_S represents the price per foot of structures, P_C represents the price per foot of cable, and T represents the traffic carried on the link, was examined. In Figure 3a, we set $P_S = 1$ and $P_C = 0$. Since the only cost is the distance-related cost of structures, this example illustrates the outcome of the unmodified Prim algorithm. In Figure 3b, we set $P_S = 0$ and $P_C = 1$. In this case, the algorithm seeks to minimize the distance of every node from the central office, resulting in a "star" network. In Figure 3c, we set $P_S = 1$ and $P_C = 1$. This example illustrates a balanced network that would be constructed if cable costs and structure costs are both significant cost drivers. Figure 3d illustrates a balanced network assuming rectilinear distances rather than airline distance. Figures 3e and 3f represent the effect of creating possible junction points along the north-south and east-west axes emanating from the central office.

Figure 4 illustrates the feeder network constructed by FEEDDIST for a wire center in Montana. Solid circles represent SAIs, open circles represent junction nodes, and diamonds represent the center points of all populated microgrid cells.

Based on the feeder and distribution algorithms, the HCPM uses input data for the cost of structure, cable and electronics, as well as other engineering inputs, to determine a level of forward looking total investment in loop plant. Other model components of the synthesis model compute comparable investments in switching and transport plant. An expense module then converts these investment costs into annual and monthly cost estimates following the general procedures outlined in the previous section.

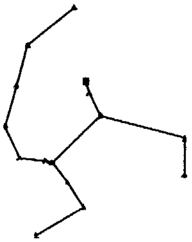

Figure 3a: Minimum Structure Distance Network

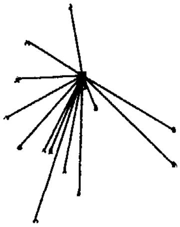

Figure 3b: Star Network

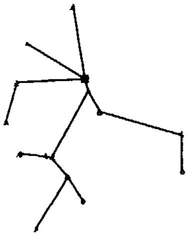

Figure 3c: Balanced Network

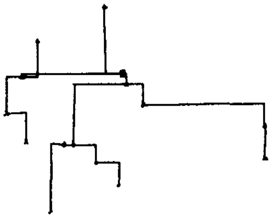

Figure 3d: Balanced Network with Rectilinear Distance

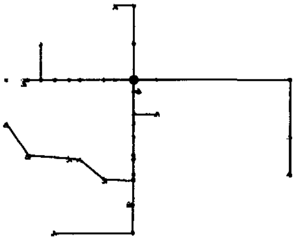

Figure 3e: Balanced Network with Junction Nodes and Airline Distance

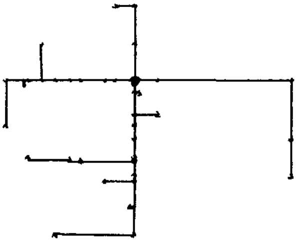

Figure 3f: Balanced Network with Junction Nodes and Rectilinear Distance

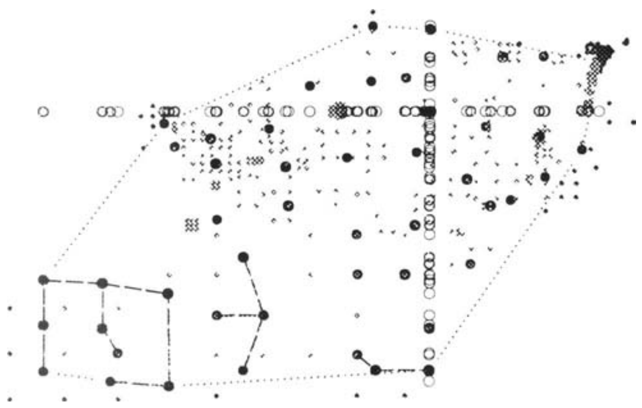

Figure 4: Feeder Network for the Wire Center ABSRMTXC

NOTES

[1] The opinions expressed in this paper are those of the author and do not necessarily reflect the views of the Federal Communications Commission or any of its Commissioners or other staff.

[2] 47 U.S.C. §§ 151 *et seq.*

[3] *In the Matter of Federal-State Joint Board on Universal Service*, CC Docket No. 96-45.

[4] *In the Matter of Access Charge Reform*, Notice of Proposed Rulemaking, FCC No. 96-488, CC Docket No. 96-262 (rel. Dec. 24, 1996).

[5] *Implementation of the Local Competition Provisions of the Telecommunications Act of 1996*, CC Docket No. 96-98, FCC 96-325 (released August 8, 1996), Order on Reconsideration *Implementation of the Local Competition Provisions of the Telecommunications Act of 1996*, CC Docket No. 96-98, 11 FCC Rcd 13042 (1996), *petition for review pending sub nom. and partial stay granted, Iowa Utilities Board v. FCC*, No. 96-3321 and consolidated cases (8th Cir., Oct. 15, 1996), *partial stay lifted in part, Iowa Utilities Board et al. v. FCC*, No. 96-3321 and consolidated cases (8th Cir., Nov.1, 1996).

[6] 47 U.S.C. §§ 251(c)(2)-(4), 252(d)(1).

[7] 47 U.S.C. § 254.

[8] In providing interstate long-distance service, interexchange carriers use local telephone companies' facilities to originate and terminate calls. The use of local telephone company facilities to originate and terminate long-distance calls is referred to as access service. LECs receive access charges for providing interexchange carriers with access to the local exchange carrier's customers.

[9] CC Docket 96-45 and CC Docket 97-160.

[10] *Local Competition Order*, para. 706. The Commission adopted a forward looking incremental cost methodology known as total element long-run incremental cost (TELRIC) for use in setting interconnection and unbundled network element prices. *Id.* at para. 672. This provision was stayed by the Eighth Circuit Court of Appeals and later reversed by the United States Supreme Court.

[11] Dun and Bradstreet report data on the number of daytime employees by CBG.

[12] Fill factors or utilization rates of loop plant are the percentage of a loop plant's capacity that is used in the network. Utilization rates are necessarily less than 100 percent so that capacity is available for growth or, in the event of breakage, to avoid outages. Lower utilization rates mean that carriers deploy more unused capacity, which increases the cost of loop plant.

[13] Alternatively, the differences in fill factors and loop lengths, and thus the cost of providing service to a particular area, may depend upon the density of customers, not the type of customers, in a particular area.

[14] Tier 1 local exchange carriers are companies having annual revenues from regulated telecommunications operations of $100 million or more. *Commission Requirements for Cost Support Material To Be Filed with 1990 Annual Access Tariffs, Order,* 5 FCC Rcd 1364 (Com. Car. Bur. 1990).

[15] Economic lives are specified for each of thirteen categories of plant.

[16] This section contains a condensed version of the HCPM model documentation, which is available on the internet at http://www.fcc.gov/ccb/apd/hcpm

[17] The size of a microgrid cell is specified by the user. The recommended size is a square with sides equal to 500 feet, but smaller cells can be chosen at the expense of increased computing time.

[18] For information about different clustering methods see: Everitt, Brian S. (1993).

[19] This and other stopping rules in the divisive algorithm can be adjusted to increase performance.

[20] In order to generate approximately optimal results using less computing time, the user has the option of computing distribution costs using both approaches only for the lowest-density grids.

[21] See Prim, R.C. (1957) for a description of an efficient algorithm for computing minimum distance networks. A computed coded version of the Prim algorithm, and some extensions, is contained in Gower, J.C. and G.J.S. Ross (1969).

[22] In fact, one can start with any initial node and be assured of reaching a minimum-distance network using the algorithm.

[23] The unmodified Prim algorithm is, however, used to connect multiple SAIs within a grid and for connecting drop terminal nodes to SAIs. In the mathematical literature on network design, networks that allow for the creation of junction points are known as "Steiner networks." [See Sharkey (1995)]. A Steiner network must always have a cost at least as low as a minimum-distance spanning tree network.

[24] This modification is due to Jeff Prisbrey.

REFERENCES

Atkinson, J. et al. 1997. *The Use of Computer Models for Estimating Forward-Looking Economic Costs: A Staff Analysis.* Federal Communications Commission, DA 97-56, released January 9, 1997.

Bush, C.A. et al. 1998. "The Hybrid Cost Proxy Model: Customer Location and Loop Design Modules," available from FCC internet site at http://www.fcc.gov/ccb/apd/hcpm/.

Everitt, Brian S. 1993. *Cluster Analysis,* Third Edition, London: Arnold.

Gower, J.C. and G.J.S. Ross. 1969. "Minimum Spanning Trees and Single Linkage Cluster Analysis." *Applied Statistics,* 18, 54-64.

Love, R.F., J.G. Morris and G.O. Wesolowsky. 1988. "Mathematical Models of Travel Distances." *Facilities Location: Models and Methods.* Amsterdam: North Holland.

Prim, R.C. 1957. "Shortest Connection Networks and Some Generalizations," *Bell System Technical Journal,* 36, 1389-1401.

Sharkey, W.W. 1995. "Network Models in Economics." Chapter 9 in M.O. Ball et al., eds., *Network Routing: Handbooks in OR and MS.* Vol. 8. Amsterdam: North Holland.