4. Provides a second level of redundancy by using two sets of lines for offices served by a folded ring;

5. Includes a third level of redundancy by providing one extra DS1 for every seven working DS1s on the port side in a central office;

6. Determines the number of rings to be built and the sequences of nodes on the ring;

7. Allows the user to run the model for a single ring, thereby enabling the user to trace the cost calculations through the logic of the model;

8. Maps the nodes subtending a particular host or tandem; and

9. Provides the following reports for each ring: (a) transport cost results for all of the rings; (b) transport configuration of all of the rings; and (c) universal service transport cost on a per line basis.

Some fundamental assumptions underlie the transport calculations. First, the model connects all remote offices to their respective host offices via SONET rings (if there is only one remote, a folded ring is assumed). Likewise, the model connects all host offices to their respective tandems through the use of SONET rings. Second, the model assumes unidirectional SONET deployment.

5. Signaling

Signaling costs for use in developing per line investments for BCPM 3.1 are provided through a user input table that reflects the cost of building a modern SS7

network. The input table provides investments for residence and business lines for small, medium and large companies. The signaling cost for a wire center is based on a weighted average of residence and business lines associated with that wire center. Values in the input table are developed by running the BCPM Signaling Module for portions of the U S WEST territory.

6. Support Plant

Once the model calculates the loop, switching and interoffice plant (excluding land and building) needed for each grid, it uses user-adjustable investment ratios to load in the support investments. Support investment represents those plant items not directly used in provisioning basic service.

BCPM 3.1 produces estimates of total investment less support investment in the loop module. It derives estimates for support investment through the application of support factors, the values of which are specified by the user. These factors represent the ratio of support investment in various accounts to total investment, less support, land and building investment.

7. Capital Costs

The BCPM 3.1 Capital Cost Module develops a series of annual charge factors for depreciation, rate of return and tax rates that, when applied to individual investment categories developed in other modules, produce capital costs for use in developing universal service fund costs. The module incorporates all of the methodologies that are currently in practice today, including: deferred taxes; mid-year, beginning-year, and end-year placing conventions; Gompertz-Makeham

Survival curves; Future Net Salvage Values; Equal Life Group methods; and many others. The module also incorporate separate cost of debt and equity rates, along with the debt to equity ratio.

8. Operating Expenses

BCPM 3.1 estimates operating expenses through application of user-adjustable expense factors. The model allows the user to specify maintenance expenses as either a per access line amount or as a percent of investment.

9. The Report Module

The Report Module provides the final step in the process of developing universal service support levels. In this module, the model combines costs factors -- including depreciation, return and taxes -- with operating expenses to generate monthly costs. The model then uses monthly costs to calculate universal service support for a given benchmark. These support levels are available at the grid, wire center, company, Census Block Group ("CBG") or state level.

The user has the option of either running BCPM 3.1 with user-adjustable inputs, the default inputs, or state-specific specific inputs.

B. BCPM Satisfies The Commission's Criteria In The Universal Service Order While The HAI Model Does Not.

BCPM 3.1 satisfies each of the ten criteria in the Commission's Universal Service Order. By contrast, the HAI model fails to comply with several of the criteria.

Criterion 1: *BCPM 3.1 Uses Least-Cost, Reasonable Technology*

BCPM 3.1 "incorporates least-cost, most efficient and reasonable technology." In doing so, the BCPM 3.1 establishes an optimal grid size that is determined by adhering to sound engineering practices that reflect forward looking, least cost technology for providing basic service. Thus, the use of reasonable technology is assured by observing sound engineering practices, and the technology also meets the requirements of efficiency and being forward looking.

BCPM 3.1 builds a forward looking, least cost network that is superior to that built by the HAI model. By using standard engineering practices based on CSAs and DAs, BCPM assures that the network will be able to meet service requirements for all customers. By relying on non-standard engineering practices, the HAI model builds a network that delivers inferior service to customers served by long loops.

The technology that is used in BCPM 3.1 is reasonable and least-cost:

Switching

For large wire centers, BCPM 3.1 uses a switch curve based on Lucent 5ESS and Nortel DMS-100 digital switches. The model has separate switch curves for host, remote, and stand-alone switches for both vendors to support current and forward-looking deployment practices. For small wire centers, BCPM uses a default switch curve that includes Nortel, Siemens Stromberg-Carlson, Lucent, and Mitel switches.

Feeder Equipment

BCPM 3.1 has two DLC categories, each of which provides multiple size options for remote and central office terminals. These options allow placing small DLCs in CSAs that serve a relatively small number of customers. Large DLCs are assumed to be integrated DLC systems, while the small DLCs utilize universal systems. In addition, where appropriate, the model captures the efficiencies that large DLCs provide. The model determines whether to use a small or a large DLC based on the number of lines the DLC can serve. Given an engineering fill factor of 90%, the model places a small DLC if the CSA serves less than 216 lines, i.e. 240 times 90%. This engineering fill factor is a user-adjustable input.

A typical DLC remote cabinet size for a large DLC, such as the "Litespan-2000," can serve only up to 1,344 lines. Whether more DLCs are placed in that CSA depends on whether sound engineering practices call for another DLC or whether it is optimal to divide a grid further, into smaller ultimate grids, each representing a CSA. For example, it is possible for a single CSA to serve 5,000 customers if a large number of customers are located in a single office complex. In this case, multiple DLC systems would be installed to provision the 5,000 lines.

The large DLC Remote Terminal ("RT") used in BCPM is the DSC Litespan LSC-2030 Remote Terminal Outdoor Cabinet which supports up to 1344 lines. BCPM assumes that the Litespan RPOTS channel unit is used in the RT except in cases where distribution cable lengths exceed CSA standards. In these cases, a RUVG2 or REUVG channel unit is recommended per DSC Litespan Practice OSP 363-20-010 Issue 6, July 1997 at 5.3.2. The BCPM sponsor's transmission

engineers use the REUVG card in actual networks. The REUVG is used on extended range loops in BCPM 3.1 because for the modest increase in cost, it provides superior performance and significantly greater flexibility in application.

Feeder Cable

The type of cable the model uses in the feeder system is determined based on the specified copper/fiber breakpoint. The copper/fiber breakpoint is a user adjustable input. The default input for the copper/fiber breakpoint is 12,000 feet. A copper/fiber breakpoint of 12,000 feet requires placing copper in the feeder if the maximum loop length from the wire center to all customers within an ultimate grid is less than 12,000 feet. If the loop length for any customer in the ultimate grid exceeds 12,000 feet, fiber is placed in the feeder to serve all customers in the ultimate grid. For all loops, cable beyond the DLC site is copper.

Feeder cables are sized to accommodate the number of working lines based on total residential, business, and special access lines. The size of feeder cables is based on the number of actual working lines adjusted by a variable engineering fill factor. For example, at an 85% engineering fill factor, a 400 pair cable can accommodate 340 working pairs before increasing the cable size. The default assumes a 65% engineering fill factor for the lowest density zone, a 70% engineering fill factor for the next two lowest density zones, and a 75% engineering fill factor for the remaining six density zones. These engineering fill factors for feeder cable are user adjustable inputs.

The total capacity for a fiber feeder segment is the sum of the required large DLC fiber strands and required small DLC fiber strands. BCPM 3.1 determines the number of maximum size fiber cables and the size of the additional fiber cable to meet the capacity needs of the segment. The fiber feeder cable sizes available in the model are 12, 18, 24, 36, 48, 60, 72, 96, 144, and 288 strands.

The feeder cable is connected to distribution cable at a Feeder Distribution Interface ("FDI"). The FDI connects many distribution cables to a feeder cable.

Distribution Equipment

The BCPM distribution technology is designed not to impede the deployment of advanced services. All customers within a copper/fiber breakpoint distance of 12,000 feet of the central office are served with copper feeder facilities. Customers beyond this distance are served from a DLC system connected to the central office by fiber facilities.

In determining the number of FDIs to install in an ultimate grid, the model reviews the cable sizing used in the grid. When the distribution cable sizing exceeds 1,200 pairs, the model places an FDI at the road centroid within each populated distribution quadrant. Thus, the FDI is placed at the center of the DA. If there are no roads, and therefore, no population located within a particular distribution quadrant, no distribution plant is placed in that distribution quadrant. Feeder cable, consisting of horizontal and vertical connecting cable, links the DLC to the FDI within non-empty quadrants.

When the distribution cable sizing does not exceed 1,200 pairs, the model allows for cost savings from placing fewer FDIs. More precisely, for CSA/ultimate grids that are served by distribution cables totaling less than 600 pairs, the algorithm essentially computes the cost of placing a single FDI within those ultimate grids. This is tantamount to co-locating the FDI with the DLC. In such cases, horizontal and vertical connecting cable[1] is placed from the ultimate grid road centroid to the road centroid of a non-empty quadrant's road reduced cluster. Within the model there are a number of rules that are used to select specific pieces of equipment to be used in the distribution plant. Among those rules with the most impact are:

1. Within a grid, if the length of copper from the DLC to the last lot in a quadrant is less than 11,100 feet, 26 gauge cable is used to serve all customers. In those circumstances where the distance from the DLC to the last lot is greater than 11,100 feet, 24 gauge wire is used in all cables to and within the distribution quadrant. Where distances exceed 13,600 feet, extended range plug-ins are installed on lines that exceed 13,600 feet.

2. The mix of aerial, buried and underground facilities is determined by terrain and density specific to that grid.

---

[1] While this is typically considered distribution cable, the Model has fixed the classification of this cable as feeder. In a future release of BCPM, this cable will be classified differently.

3. Drop terminals are provided at each point where drops connect branch cables and are sized for the number of connecting drops.

4. Indoor building terminals are placed on each multi-tenant building and are sized for the number of lines terminated at that location.

5. Different NIDs are used for business and residence locations.

6. Branch cables are sized to the number of pairs for housing units and business locations.

Transport

The BCPM transport module is based upon SONET, which is a set of standards for optical (fiber optic) transmission. SONET was developed to meet the need for transmission speeds above the T3 level (45 Mbps) and is generally considered the standard choice for transmission devices used with broadband networks. Technologies like T3 are likely to be replaced by new services offered through a SONET platform. By way of comparison, OC-1 can carry over 30 times more data than DS1.

SONET enables more efficient use of installed fiber; it taps the latent capacity already in the network. SONET allows new network configurations, including ring networks, which have a greater degree of survivability than traditional mesh networks. The transport module has three different size/bandwidth SONET terminals (OC3, OC12, OC48). The model's algorithms

select the appropriate terminal size/bandwidth based on traffic demands, making it an efficient model while building in redundancy to the network.

The BCPM 3.1 transport module uses manual digital cross connect systems as opposed to automated cross connect systems. Automated digital cross connects are typically associated with the provisioning of dedicated special services. In modeling basic service, BCPM 3.1 provides the cost of interoffice transport connections of umbilical switching trunks to a remote. For universal service purposes, the BCPM 3.1 sponsors advocate the use of manual cross connect technology as a more cost effective solution since these switched umbilical and interoffice trunks are not rearranged frequently. The use of automated digital cross connect technology at every node location would cause a cost increase in the interoffice transport element for universal service that is not warranted.

Cable Type

The type of cable used in the feeder system is determined based on the specified copper/fiber breakpoint. The copper/fiber breakpoint is a user adjustable input; the default input for the copper/fiber breakpoint is 12,000 feet. A copper/fiber breakpoint of 12,000 feet requires placing copper in the feeder if the maximum loop length from the wire center to all customers within an ultimate grid is less than 12,000 feet. If the loop length for any customer in the ultimate grid exceeds 12,000 feet, fiber is placed in the feeder to serve all customers in the ultimate grid. For all loops, cable beyond the DLC site is copper.

Unlike the HAI model, BCPM 3.1 does not use the obsolete T-1 carrier system in its feeder and sub-feeder. T-1 carrier on copper cable is unlikely to be an economical choice if all relevant costs are considered, because it requires specialized design and cable conditioning for each loop, an extremely expensive proposition. The HAI model apparently does not include these extra design and conditioning costs for the T-1 carrier. T-1 carrier is also likely to require coarser gauge cable than the HAI model assumes.

Within a grid, if the length of copper from the DLC to the last lot in a quadrant is less than 11,100 feet, 26 gauge cable is used to serve all customers. In those circumstances where the distance from the DLC to the last lot is greater than 11,100 feet, 24 gauge wire is used in all cables to and within the distribution quadrant. Where distances exceed 13,600 feet, extended range plug-ins are installed on lines that exceed 13,600 feet.

Host-Remote Switching Configurations

The Commission tentatively concluded that the model should enable the placement of host switches in certain wire centers and remote switches in certain wire centers. BCPM 3.1 meets these requirements. It has separate switch models for host, remote and stand-alone switches. BCPM 3.1 places hosts and remotes based on the nature of the switch that is currently in that switching node, according to the Bellcore's Local Exchange Routing Guide ("LERG").

The BCPM 3.1 Switching Module has a detailed method for allocating costs of the switches on the basis of functional categories of investment, so that customers

in the host-remote relationship pay for the cost of the functions they use. Switching investments are allocated among customers as follows: The processor investment per line is determined by a three-step process that allocates the host processor investment across all switches on the host/remote complex. The first step is to divide the total USF processor investment for all switches on the complex by the total number of lines on the complex. This produces a host processor investment per line. The second step is to divide the processor investment for each remote switch by its associated number of lines. This produces a remote processor investment for each remote. The final step is to compute the total processor investment per line for each switch. For standalone switches, this is simply the processor investment from step one. For hosts and remotes in the same rate center, the per line investment is the weighted average of the host investment for the host and the host plus remote investments for each remote. This produces a single processor investment per line for all switches in the rate center. For remotes located outside the host rate center, the processor investment is the sum of the host processor investment per line and the remote processor investment per line.

The trunking and SS7 host office investments must be allocated by complex, since remotes are assumed not to have these facilities and use the trunking and signaling resources of the host. For each complex, BCPM divides the host USF trunking investment by the local trunk usage for all switches on the complex. SS7 investments are handled similarly.

2. BCPM 3.1 Uses Existing Wire Centers As The Center Of The Loop Network.

The starting point of the BCPM 3.1 design is the existing central office locations. The model uses the wire center vertical and horizontal coordinate location information from LERG to locate the central office within the wire center. Feeder routes are designed to begin at this point, and move out to cover the area within a wire center in the manner described above.

3. The Loop Design In BCPM 3.1 Supports The Provision Of Advanced Services.

The loop design of the BCPM 3.1 supports the provision of advanced services. The voice grade service that the design would provide includes the capability to support currently available modems for dial up access. The model does not use loaded loop plant.

The BCPM 3.1 is designed carefully to observe required limits for loading and resistance by limiting copper loop lengths to twelve kilofeet. The BCPM 3.1 design is based on 26 gauge cable in the feeder and 26 and 24 gauge cable in the distribution. This allows the design to meet both the 1500 ohm supervisory limit of today's digital switches and the 900 ohm powering limit of digital loop carrier line cards, without requiring the use of much more expensive extended range cards. By avoiding the use of bridge taps, the BCPM 3.1 design also removes capacity concerns.

The BCPM sponsors have provided empirical evidence that the 12 Kft maximum copper loop length is more cost-effective in almost every case than an

18Kft loop, while preserving the service quality needed for advanced services.[2] The 12Kft standard proves to be more cost-effective because at 18Kft, a heavier 24-gauge cable and more expensive extended range DLC line cards must be used to support transmission standards for advanced services.

4. Wire Center Line Counts In BCPM 3.1 Match U S WEST's Actual Wire Center Line Counts.

The wire center line counts in BCPM 3.1 match U S WEST's actual wire center line counts.

Criterion 2: *Each network function in BCPM 3.1 has an associated cost.*

The model develops costs for the local loop, including costs for the drop, distribution, feeder, and the switch, along with costs for transport signaling, support plant, and associated capital costs and operating expenses. The algorithms the model uses ensure that the model provides sufficient plant and equipment. These algorithms are clearly documented and verifiable within the model software and methodology documentation.

Criterion 3: *BCPM 3.1 incorporates the forward-looking cost of purchasing and operating known and proven facilities, equipment and technologies.*

While switch (i.e., wire center) locations are assumed to be fixed, no equipment or technology is assumed to be embedded or fixed; all equipment is assumed to be variable and avoidable. Forward-looking costs are based on material

---

2 "Analysis of 18Kft and 12Kft Runs", Submission of the BCPM3 Model by BellSouth Corporation, BellSouth Telecommunications, Inc., US WEST, Inc., and Sprint Local Telephone

prices net of discounts rather than list prices for equipment and material. The model does not rely upon embedded costs for facilities, functions or elements.

Criterion 4: *The rate of return should be either the authorized federal rate of interstate services, currently 11.25 percent, or the state's prescribed rate of return for intrastate services.*

BCPM 3.1 allows the user to select a rate of return or to utilize the FCC's recommended rate of return of 11.25%.

Criterion 5: *Economic lives and future net salvage percentages used in calculating depreciation expense should be within the FCC-authorized range and use currently authorized depreciation lives.*

BCPM 3.1 includes two different sets of inputs for depreciation expense. The first set of inputs consists of default values that use economic lives and future net salvage percentages that are within the FCC's authorized range. These values comply with this criterion of the FCC's checklist.

The second set of inputs uses economic lives that the BCPM sponsors deem appropriate. U S WEST believes these lives more accurately reflect forward-looking, economic lives than do the lives used in the FCC's range. The economic lives U S WEST advocates are:

a. Aerial and Underground Cable Accounts: 15 year life;

b. Buried Cable Account: 20 year life;

c. Digital Switching Account: 10 year life;

d. Digital Circuit Account: 10 year economic life; and

Companies, CC Dockets Nos. 96-45 and 97-160, December 11, 1997.

e. Non-Metallic Cable Account: 20 year economic life.

U S WEST believes that the FCC's range does not reflect true economic lives for several categories, and, therefore, it supports Ameritech Michigan's request for a waiver of compliance with the FCC's range.

Criterion 6: *BCPM 3.1 complies with this criterion by including multi-line business services, special access lines, and multiple residential lines.*

The inclusion of these services and lines causes the model to reflect the economies of scale referred to in this criterion.

Two factors ensure that BCPM 3.1 accounts for the cost of providing service for all business and households within a geographic region. First, the model uses the most current information and best possible techniques to identify and locate housing units and businesses in the wire center area. The methodology used to accomplish this step is described above.

Second, the model provides two methods to develop the service needs of the households and businesses in the wire center. With the first method, the user can directly input wire center line count information. As an alternative, the model uses a residence line multiplier, a single business line multiplier, and a special access line multiplier to reflect the line needs in the wire center. The residence line multiplier is a factor developed at a state level from ARMIS and NECA data and is applied to the number of housing units to produce the number of residence lines served in each grid. The single line business multiplier is also a state level factor developed from ARMIS and NECA data and, when applied to the number of total business lines, produces the number of single business lines in the wire center. The

special access line multiplier is a factor developed from BCPM sponsor studies and, when applied to the number of total business lines, produces the number of special access and private lines in each grid.

Criterion 7: *BCPM 3.1 allows the user to input either a common cost factor or expenses on a per line basis.*

This is a user-adjustable input.

Criterion 8: *All underlying data, formulae, computations and software used in the study are available for review and comment.*

Over the last 18 months, the BCPM model has evolved through a series of enhanced versions, with each version improving on the reasonableness of inputs and the plausibility of outputs. This progressive refinement of the model has been driven by a variety of factors, including a series of field tests comparing results with actual data; challenges to the model in workshops with state regulators, challenges to the model by FCC staff members, critiques of the model by advocates for other proxy models; and numerous analytical studies. This intensive process has led to the use of reasonable inputs in BCPM and outputs that clearly are plausible.

On March 2 and 3, 1998, BCPM sponsors filed ex parte documentation with the FCC discussing the results of tests relating to the most recent output runs of both the BCPM and HAI models. The test demonstrated that:

- As would be expected, the BCPM model produces higher funding in less dense western states and lower funding in dense eastern states. The HAI model generated the non-intuitive results that eastern states

require greater universal service funding than the more rural western states.

- Using road length in the model (as BCPM does) avoids overbuilding in dense areas and underbuilding in sparse areas.

- BCPM wire center cable route mileage does not exceed actual road mileage.

- BCPM grid areas do a good job of modeling the actual wire center area. HAI exceeds wire center area in dense areas and falls short of actual wire center area in less dense areas.

- BCPM is very accurate in its line counts. The HAI special access line count is questionable.

- BCPM has a high correlation between predicted customer locations and actual locations in tested rural wire centers.

- The BCPM model is sensitive to changes in key variable values where HAI's sensitivity is very questionable.

- The BCPM uses industry standard CSA engineering design.

Criterion 9: *All of the underlying data, formulae, computations and software used in BCPM 3.1 are available for review and comment.*

The availability of this information allows users of the model to examine the critical assumptions and engineering principles. All critical assumptions and engineering principles are outlined in the BCPM methodology and are available for critical analysis.

Criterion 10: *BCPM 3.1 provides for efficient targeting of universal service support by disaggregating investment calculations down to the individual grid.*

This disaggregation provides for a level of deaveraging that is substantially more granular than the level of deaveraging available with CBGs. BCPM 3.1 also can aggregate to the CBG or wire center level.

ATTACHMENT B

HAI Distribution Cable Requirements

Issue: Whether the distribution plant modeled by HAI 5.0a is adequate to serve customers in their "actual" locations as identified by PNR and Associates (PNR).

Finding: The distribution route miles modeled by HAI 5.0a are too few to serve the customers in the *convex hull* clusters of geocoded and surrogate locations that underlay the rectangular clusters. The rectangular clusters are used in HAI 5.0a in the design of the network.

Hence, HAI 5.0a's estimate of the required investment in rural, low-density areas is too low.

Discussion: The customer locations assumed by HAI 5.0a for the purpose of "building" plant are inconsistent with the "actual" locations in the underlying polygon (convex hull) clusters.

The figure below shows a hypothetical convex hull cluster of geocoded and surrogate locations. The rectangle shown is derived from the North-South, East-West aspect ratio and area of the convex hull cluster. Specifically, the rectangle has the aspect ratio of the rectangle that just covers the convex hull cluster (a *minimum*

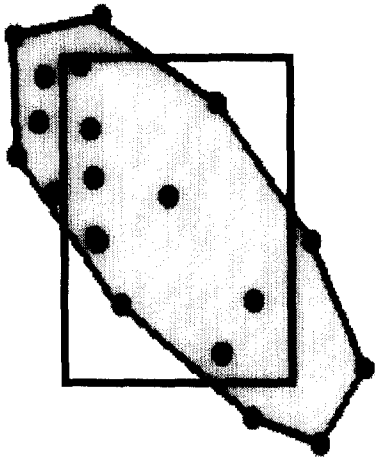

bounding rectangle) and the

area of the convex hull cluster itself. The rectangle cluster is what is directly used by HAI 5.0a in its design of the network.

HAI 5.0a assumes that customer locations (i.e., lots) are evenly distributed within the rectangular cluster. For simplicity, assume there are 9 locations. This yields the following figure.

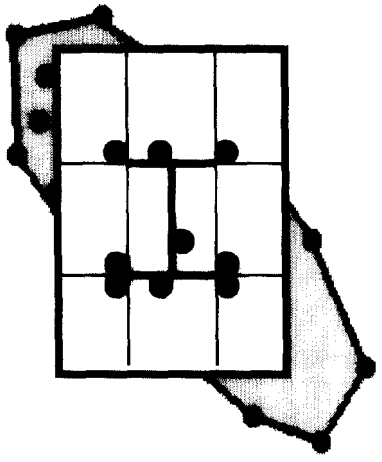

HAI 5.0a subtracts off two lot depths from the cluster North-South length to determine the length of the backbone cable. It also subtracts off two lot widths from the East-West cluster length to determine the length of the branch cable. In the figure shown above, there are two branch cables. Backbone and branch cable is laid in only the middle lot. A drop serves the house in each lot.

Since the default drop length in the lowest density area is 150 feet, the house in each lot must be 150 feet from a branch cable. That is, the houses are concentrated toward the center of the rectangular cluster as indicated in the figure.[1]

This has an important implication for whether the model is providing for a realistic amount of cable. Assume that the area of the convex hull is 15 square miles. Hence, the area of the rectangle is the

[1] As modeled by HAI 5.0a, it is only the distance from the cluster center to the edge of the middle lot (in this example) that matters for determining whether multiple DLCs are needed.

same and the area of each lot is roughly 1.67 square miles. Lots are assumed to be twice as deep as they are wide. Each lot is 1.83 miles deep (9,640') and 0.91 miles wide (4,820'). Thus, the total distance of cable, including the 150' drops, in this cluster = 9,640' + 2*4,820' + 9*150' = 20,630' or 3.91 miles.

Examining the underlying convex hull cluster of geocoded and surrogate locations strongly suggests that this amount of cable is much too little to serve customers in their "actual" locations. That is, the placement of customers for determining cable lengths within the rectangular clusters is inconsistent with where PNR locates customers in the underlying polygon clusters. In reality, customers are more widely dispersed. Not only will more cable be required but also the 18-kft copper criterion will likely be violated more often, thus requiring additional electronics.

Analysis: A determination of whether HAI 5.0a is not modeling enough distribution plant in its rectangular clusters can be made in the following manner. First, the distribution plant route miles modeled by HAI 5.0a for a specific rectangular cluster is found. Then, the "minimum spanning tree" distance in the underlying polygon cluster is calculated.[2] If the amount of distribution plant route miles modeled by HAI 5.0a is less than the minimum spanning tree amount, then we conclude that HAI 5.0a is not building enough plant to reach customers in the "actual" locations identified in the polygon clusters.[3]

Theoretical Examples: **Example #1**

HAI 5.0a groups a set of "actual" customer points into a *cluster*, according to a set of aggregation rules. The two key aggregation criteria are that no customer in the cluster be more than 2 miles from its nearest neighbor and that no customer is more than 18-kft from the centroid of the cluster, measured rectilinearly. Below is shown a hypothetical cluster that meets these criteria.

[2] A minimum spanning tree distance is the mathematically determined shortest distance that connects all of the customers within a given area.

[3] Actual is in quotes to indicate that this refers to PNR's location of customers using geocoding or its surrogate methodology. The surrogate locations likely are not customers' true spatial location.

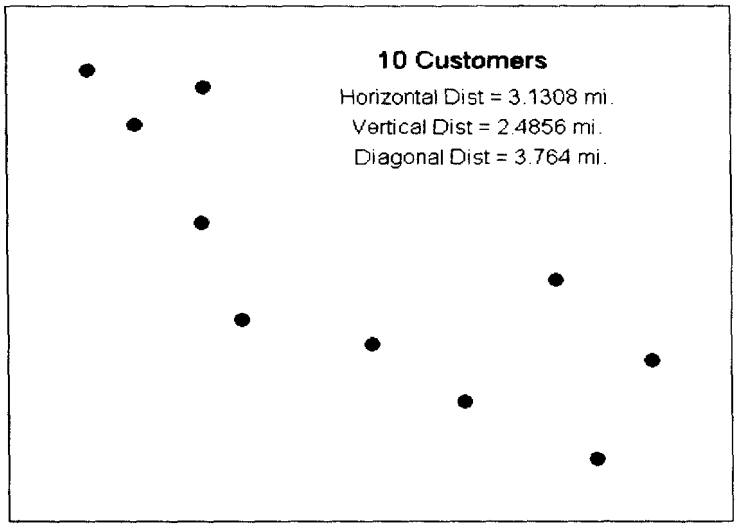

The *minimum spanning tree* for these points – the mathematically shortest connection possible for these points – is 5.88 miles.

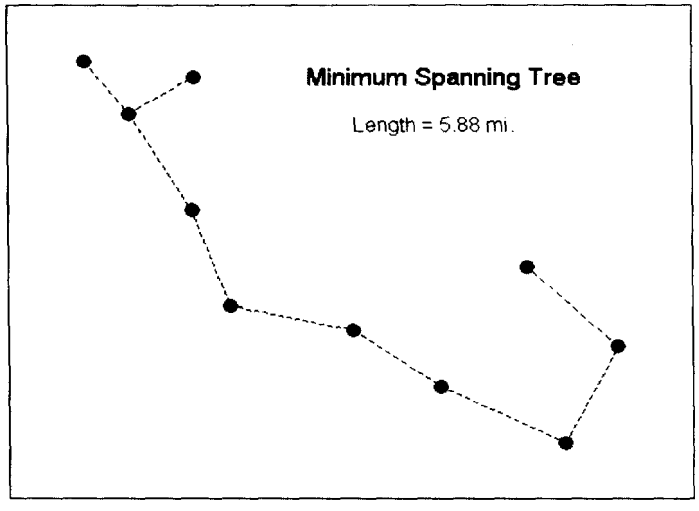

When HAI has determined the set of points that constitute a cluster, it logically draws a *convex hull* around those points, and determines its area.

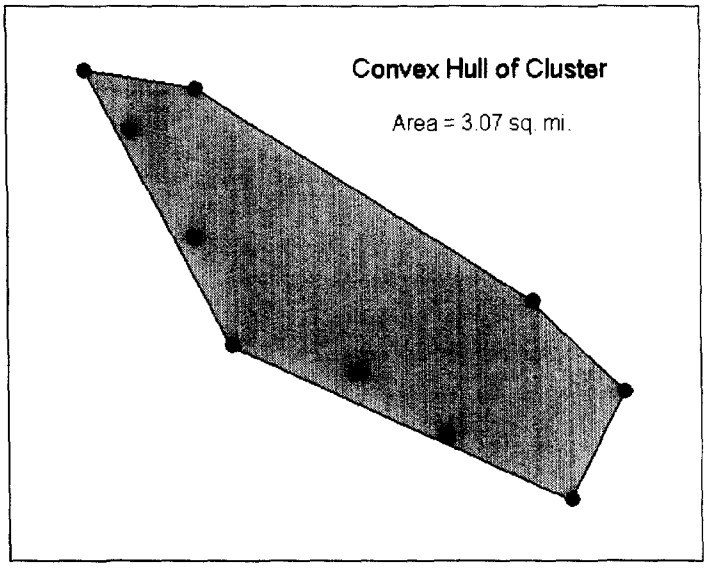

HAI then logically constructs a *minimum bounding rectangle* – oriented North-South-East-West – which exactly bounds the cluster's points. HAI then determines the *aspect ratio* of that rectangle (that is, the ratio of the rectangle's height to its width) ... in this case, 0.8.

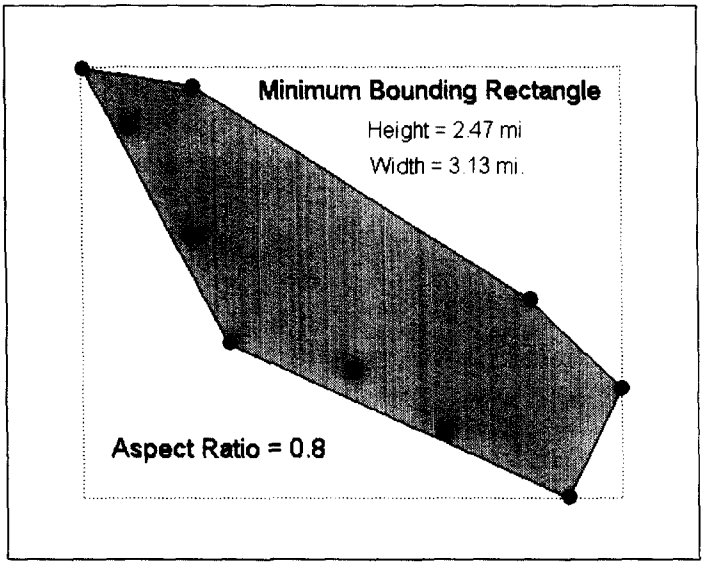

HAI then constructs a *rectangle* with the above aspect ratio; the *size* of that rectangle is determined by its *area* ... and that area is

set to be the *area of the convex hull* ... in this case, 3.07 square miles.

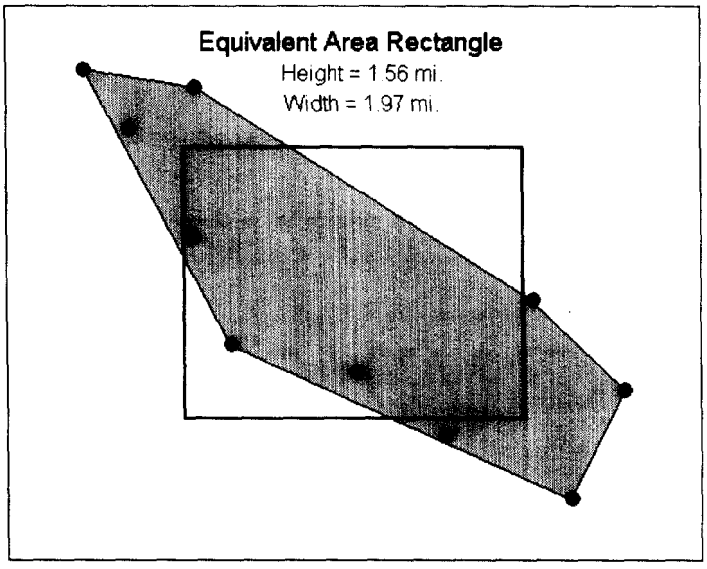
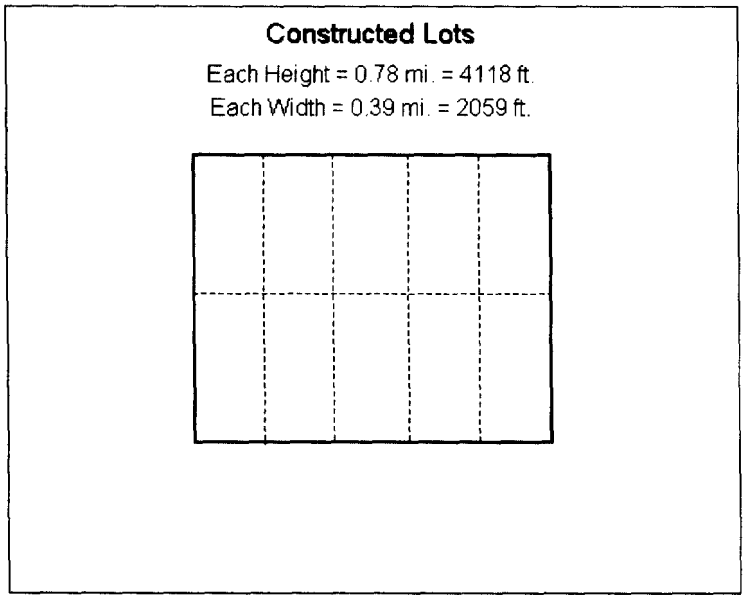
HAI then constructs *lots* within this constructed rectangle. Each lot is twice as high as it is wide.