Copyright © 1980 American Telephone and Telegraph Company
THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 59, No. 1, January 1980
Printed in U.S.A.

A Discrete-Event Simulation Analysis of Loop Network Assignment Operations

By H. T. FREEDMAN and T. R. HARMS

(Manuscript received August 2, 1979)

A computer simulation model is presented for customer movement in the loop plant. This model is used to analyze the effects of various customer movement parameters on loop operating costs and to find a strategy for the assignment of loop network facilities to customers which will minimize these costs.

I. INTRODUCTION

The local telephone network (known as the loop network) has been the object of many interesting studies encompassing problems such as the optimal design[1] and capacity expansion[2] of the network. In this study, our interest is not in modifying the network but in administering it. In particular, we study the impact of different strategies for assigning facilities to customers in differing environments. Since the movement of customers in the system and complexity of the system did not lend itself simply to analytical models, a discrete event simulation model was built to evaluate various assignment strategies.

The simulation model consists of three major parts: the randomly generated inward and outward movement of customers, the assignment (according to the strategy being studied) of a cable pair to a customer on an inward order, and the record keeping (which pairs are assigned to which customers and what costs have been incurred in the process of providing service to customers).

The effects of three environmental factors were studied with this simulation: growth (number of new customers per year), penetration (fraction of premises with telephone service at any given time), and abandonment (fraction of vacated premises which will never again be occupied). The assignment strategies studied include connect-through strategies with various holding times (see Section 3.2) and a reassignable policy. The simulation was also used to evaluate the effect of

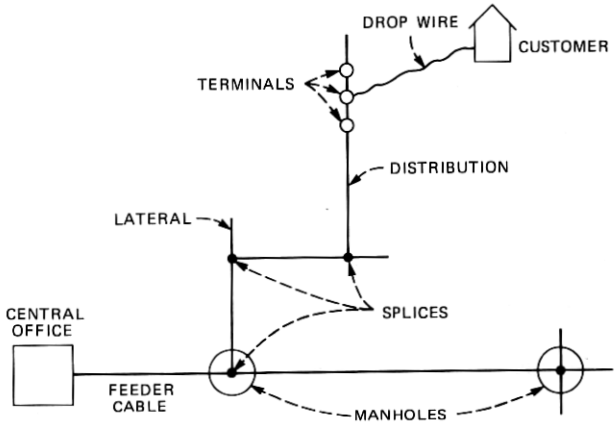

Fig. 1—Loop plant network.

some customers having two lines rather than one (second-line penetration).

In the next sections we present background information on the real system—the telephone loop plant network—and on the process we are interested in optimizing—the assignment of cable pairs to customers. Then the simulation model, which was written in SIMSCRIPT II.5,*[3] is presented along with the results of this simulation analysis. Other applications and extensions of this model are indicated in the final section.

II. THE LOOP NETWORK

2.1 Physical design

The physical system being studied is known as the loop network.[4] It is that portion of the telephone network which connects the customer to the central office, where the switching occurs. The connection is usually accomplished by a pair of copper wires (known as a pair, for short). Figure 1 shows the path of a typical pair from the central office to the customer. The pair leaves the central office in a large cable called a feeder cable. This feeder cable is located under the street in concrete or plastic conduits interrupted approximately every 500 feet by manholes to provide access to the pairs. The pair later leaves this cable through a splice to a lateral feeder cable and eventually branches into a distribution cable in the neighborhood of the customer. A distribution cable may be found either on poles or buried underground. Somewhere near the customer, the pair appears in one or more distribution terminals. The final link is provided by a drop wire from the customer's residence or business to the nearest serving terminal

---

* Registered trademark of CACI.

where it is connected to the pair. The interesting aspects of the assignment question arise when none of the pairs appearing in the nearest serving terminal are free to be used for a new customer.

2.2 Topological design

In general, it is difficult to forecast accurately where new customers will need lines. While one may feel reasonably confident about the number of new lines needed for a large area, it is much more difficult to pinpoint exactly where these new customers will appear. One approach to this problem is to design the network with flexibility built into it. One way to provide this flexibility is to have some of the pairs available in more than one terminal. A pair might appear in terminals on two different streets and thus be available for connection to potential customers on either street. Of course, at any given time it can only be used by one customer. This type of plant is called multiple plant (see Fig. 2). Other types of plant configurations also exist, but our study applies primarily to multiple plant, and we will restrict our presentation to this case. It should be noted that roughly two-thirds of the existing telephone plant can be simulated as multiple plant, and new plant is not designed in this manner anymore.

III. THE ASSIGNMENT PROCESS

Whenever a request for service (also known as an inward order) is received, a pair must be provided to connect the customer's telephone to the local central office. The provision of this pair (known as an assignment) may require only simple record changes (in the case where a pair is reserved for the customer) or may require some complex network activities to get a free pair where it is needed. Each activity has a cost based on the craft time and material required; these costs are known as the loop network operating costs. The question of interest here is the economical administration of the loop plant in the multiple configuration; that is, which pairs should be provided to which customers to minimize the operating costs over time? In this section, two assignment policies are described along with the set of activities available to be used in providing a pair under each policy.

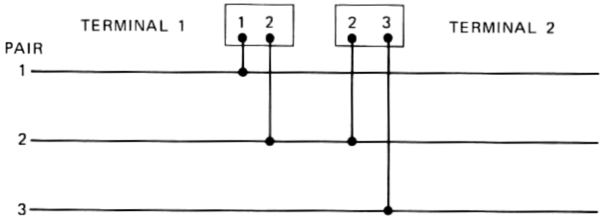

Fig. 2—Multiple plant.

3.1 *Reassignable policy*

Under a reassignable policy, when a customer discontinues service, the drop wire connecting his phone to the network is disconnected; thus, any pair not actually serving a customer (working) is available for assignment (spare).

Consider an inward order for residential service at a given address. If one or more pairs in the terminal associated with this address is spare, one will be chosen for assignment to the new customer. The connection is completed by having a drop wire connected from the customer's premises to the spare pair at the serving terminal. These two operations, assign a spare and connect a drop wire, are the minimum effort required to provide service in this case.

If no spare is available in the designated serving terminal, additional operations are required to provide service. Figures 3a and 3b illustrate one possibility known as a line and station transfer (LST). In Fig. 3a we have the following situation: In the customer's terminal a working pair exists not connected to a customer served out of this terminal, but connected through multipling to a customer in a different terminal, and this second terminal has a spare pair. The trick is to transfer the

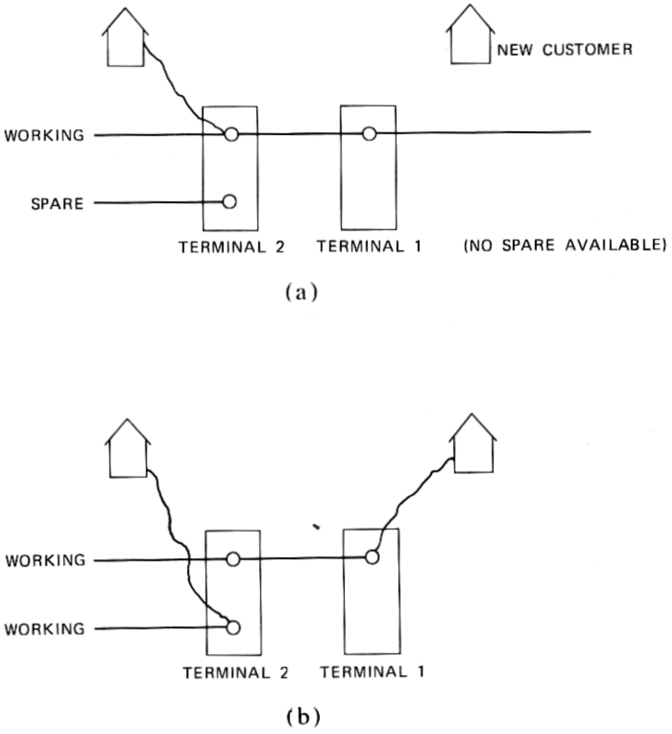

Fig. 3—(a) LST before. (b) LST after.

second customer to this spare, thereby making a spare available for the first customer (see Fig. 3b).

An LST thus involves moving a drop wire from one pair to another (a move which must be carefully coordinated with changes in the central office) on top of the usual work required to connect a spare. This operation can be repeated over several terminals to provide a pair; it is then known as a multiple-stage LST.

Another possibility is an operation simpler than an LST. It consists of stringing a drop wire from the customer's premises to a spare at a terminal different from the one it is associated with (see Figs. 4a and 4b). This is known as a wire out of limits (WOL). A WOL involves extra effort to place a longer and thus costlier drop wire. WOLs are also trouble-prone and unsightly.

Finally if none of the above operations is possible, the order is referred to the engineering department (RE), which will take more costly measures to provide service.

In summary, we have four basic activities: spare assignment, LST single or multiple stage, WOL, and RE. Each of these possible ways of providing service involves not only a different cost, but also changes

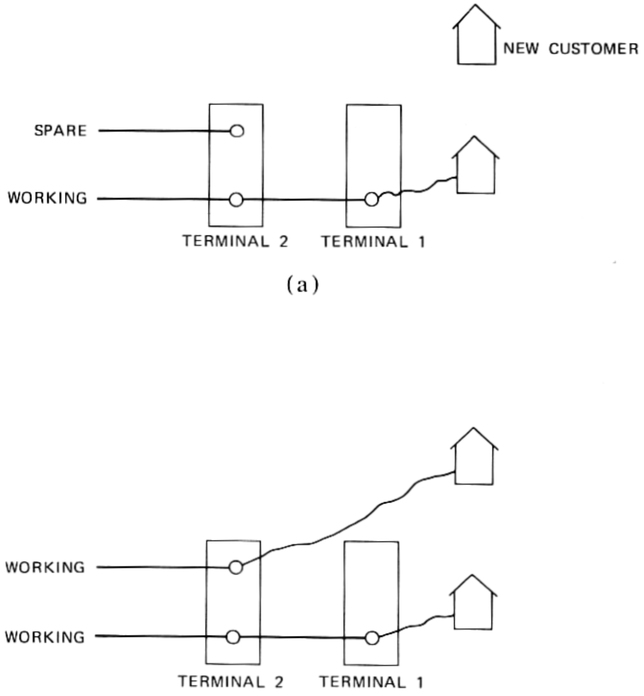

(a)

(b)

Fig. 4—(a) WOL before. (b) WOL after.

the state of the network in different ways, thus affecting the activities that will be necessary to provide service to subsequent customers.

3.2 Connect-through policy

Obviously, the costs of disconnecting a pair when a customer discontinues service and of reconnecting a pair for a new customer can be avoided if the line is left connected upon discontinuation of service, in anticipation of a new customer moving in at that same location and requiring telephone service. This is called the connect-through policy, and the pair thus reserved is called a CT. However, sometimes it may be quite a while before the new customer moves in and requests telephone service. The drawback is that reserving this idle pair for a future customer may force more complex activities (e.g., LST or WOL) to provide service for a present customer who could have used this pair. Furthermore, due to clerical errors and other environmental factors, the CT may in some cases not be usable for the new customer, thus negating any potential savings. The net savings from the CT policy must be traded off against the cost of the LSTs and WOLs necessary to avoid using the idle CTs. Other tradeoffs must be considered, including the effect of a CT policy on capacity expansion costs. In this paper, we concentrate only on the operating cost tradeoffs. Historically, the decision rule used has been to reserve a CT for service at the same location for a fixed amount of time called the holding time (60 days was the most commonly used holding time). If it is reused during that holding time, the savings due to the CT are obtained; if it is not reused during the holding time, it then becomes available for assignment, if needed, for other locations. Note that when a CT pair is used at another location, the cost of breaking the connection at the original location and reconnecting it at the new location is incurred; this is less than the cost of an LST or WOL. This activity is known as breaking a CT (or BCT).

Under a CT policy, we now have six activities: reuse CT, spare assignment, LST, WOL, BCT, and RE. What should the preference ordering among the activities be? Under which conditions is the CT policy better than the reassignable policy? What is the optimal value for the CT holding time? It is with these questions in mind that our simulation was built.

3.3 Assignment strategies studied

More formally, the following assignment strategies were evaluated:

(*i*) Break as Needed: A connect-through policy where CT pairs are broken as needed. This is equivalent to having a holding time equal to zero. The preference ordering among activities is from cheapest to most expensive (reuse CT, spare assignment, BCT, LST, WOL, and RE).

(*ii*) Nonzero Holding Time: A connect-through policy with a nonzero holding time. The preference ordering is as above, except if the CT is of an age under the holding time, then BCT is just ahead of RE. The effect of the holding time on the operating costs was evaluated.

(*iii*) Break as a Last Resort: A connect-through policy where CT pairs are broken only as a last resort. This is equivalent to having an infinite holding time. The preference ordering is reuse CT, spare assignment, LST, WOL, BCT, and RE.

(*iv*) Reassignable: A policy without connect-throughs. The preference is then simply spare assignment, LST, WOL, and RE.

The environment in which the network operates is also an important factor affecting operating costs. Therefore, the simulation was built to allow the following three neighborhood (customer movement) characteristics as input variables for study: the average length of time a premises is left without telephone service (vacancy time), the rate of growth in the number of new customers in the neighborhood, and the fraction of CTs which become unusable (abandonment rate). These characteristics are described more fully in the next section.

IV. THE SIMULATION MODEL

In this simulation, n pairs are available for assignment in the neighborhood. These pairs are terminated at a fixed number, m, of equally sized terminals in a manner specified as input (i.e., the multipling scheme or lack of it is reflected in the way the pairs are terminated). At each terminal, a number D of potential customers equal to the number of terminations in the terminal is assumed to exist. The terminals are sized for the ultimate number of customers. Three major sections of the simulation are customer movement, pair assignment, and record keeping. Figure 5 illustrates the processing flow in the simulation.

4.1 Customer movement

The simulation models customer demand in two ways. One is to model growth in the number of customers, the other one is for modelling the inward and outward movement of customers (churn).

4.1.1 The growth model

In the growth model, the system is empty at the beginning of the simulation, and growth is modeled by having each potential customer enter the system (make his first request for service) at some random time (uniformly distributed) between 0 and T, where T is chosen so that the expected number of new customers entering the system per year is G, an input variable representing the growth rate. Since there are mD eventual customers, $T = mD/G$. Once a customer requests service, he uses that service for a random time (exponentially distributed) whose mean value is an input parameter known as the occupancy time, T_o. When a customer moves out, another customer moves into

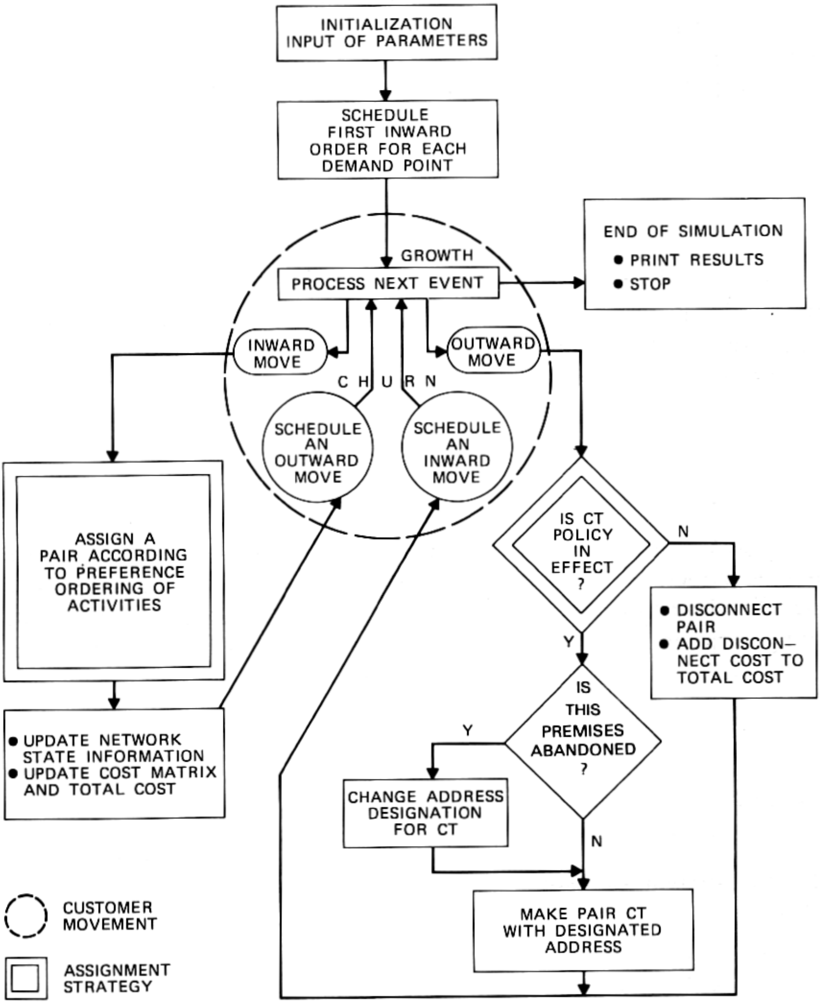

Fig. 5—Simulation flow chart.

the same location after a random interval (exponentially distributed), whose mean value is an input parameter, known as the vacancy time T_v. Other distributions (e.g., beta type II) have also been used, but the same results were obtained. The occupancy and vacancy times are characteristics of the neighborhood reflecting the amount of churn. These are often combined into the penetration parameter:

$$p = T_o/(T_o + T_v).$$

Penetration reflects the fraction of premises in the system expected to be working (in-service) at any given time.

Each simulation starts at time $t = 0$ with no working or assigned pairs, and thus the expected number of working pairs at time t is:

$$E[W(t)] = pGt.$$

The simulation continues until the time of working exhaust (when all available pairs are connected to customers) is reached.

4.1.2 The no-growth model

In the no-growth model, the simulation starts at time 0 with the steady-state expected number of customers present in the system:

$$w(0) = pmD, \quad \text{where} \quad w(0) < n.$$

The other $(1 - p)mD$ remaining potential customers enter the system in an exponentially distributed time with mean T_V. Once a customer is in the system, his movement is governed by the same model as for the growth case. Since all mD customers remain in service for an average time T_o and out of service for an average time T_V, the expected number of customers at any time is:

$$E[w(t)] = [T_o/(T_o + T_V)]mD = pmD = w(0).$$

The simulation, in this case, is terminated at a set time or if working exhaust is reached.

4.1.3 Abandonment

Abandonment is another demand phenomenon which the simulation can model. Abandonment is the unreusability of some CT pairs due to either physical abandonment of a premises or to changes or errors in address designation which cause plant assignment procedures to ignore reuse possibilities. This is modeled in the simulation by taking each outward order and making it an abandonment with probability A, where A is an input parameter. If abandoned, the pair is left CT and a new inward order is generated as before, but the customer's address is changed so that the CT can never be reused. The record-keeping system (see Section 4.3), however, does not maintain information on which CTs are abandoned since this information is missing in the real world.

4.1.4 Customer movement effects studied

In the simulation analysis, several different scenarios were used for the customer movement characteristics. These included three values for the growth rate, two for penetration and two for abandonment. These were:

High Growth = An average of seven new customers per year per 100 available pairs.

Low Growth = An average of two new customers per year per 100 available pairs.
Zero Growth = No new customers in the system at any time.
Low Penetration = An average of 80 percent of premises which had service at any time currently have service (this corresponds to a vacancy time of one year with an occupancy time of four years).
High Penetration = 97.5 percent (this corresponds to a vacancy time of 0.1 year with an occupancy time of four years).
Zero Abandonment = 100 percent of all outward orders may have customers eventually returning to the same address.
High Abandonment = 90 percent of all outward orders may have customers returning to the same address. The other 10 percent will not.

4.2 Pair assignment

When a customer requests service at a terminal, an assignment algorithm is used to find an available pair. The algorithm searches sequentially through 13 possible operations for assigning a pair (the six previously described plus various combinations of them—see Table I) until it finds one that can be used. The order in which these operations are searched is specified as an input, so that different orderings can be compared, as was described in Section III. Table I lists these operations in order of increasing complexity and cost. Note that when it is specified in the input that a CT plan is not used, the operations involving CT pairs are ignored and a cost of disconnection is assessed every time a customer leaves the system.

Table I

1. Reuse a CT pair at the same address.
2. Assign a spare.
3. Break an overage* CT at the same terminal.
4. Break any CT at the same terminal.
5. Break an overage CT at a different terminal.
6. Break any CT at a different terminal.
7. Perform a 1-stage LST.
8. Perform a 1-stage LST by breaking an overage CT.
9. Perform a 1-stage LST by breaking any CT.
10. Perform a 2-stage LST (may include breaking CT).
11. Wire out of limits (assign to any spare pair).
12. Wire out of limits by breaking any CT.
13. Referral to Engineering. This is modeled as having a special operation performed to provide service, such as putting in an electronic "pair gain" system. However, it is removed as the customer leaves and the customer is not counted in the working fill.

---

* An overage CT is a CT that has been idle longer than the holding time specified as input.

4.3 Record keeping

Two types of records are maintained: the network state (which pairs are assigned to which customers in which terminal) and cost tracking (which activities have been used when).

4.3.1 Network state

At all times, the simulation maintains a complete current record of the network. Each pair in the network has a status—working, idle CT, or spare. If it is spare, it is available at all terminals in which it appears in multiple. If it is working, the address of the customer is recorded. For an idle CT pair, the address of the last customer to which it was connected is maintained; in addition, the date of discontinuation is kept for comparison with the holding time as needed. The configuration of the network is also updated as LSTs and WOLs occur. At all times, the number of working pairs is monitored to indicate the network fill. The fill of the network is the ratio of the number of working pairs to the total number of available pairs and is used as a measure of network congestion. In general, more complicated operations will be required more often as fill increases. The number of CT pairs is also monitored to provide a measure of CT penetration. The operating cost at any time will generally be a function of the fill and CT penetration.

4.3.2 Cost tracking

When an operation is selected in the assignment phase of the simulation, it is tallied in a matrix, N, whose rows represent network fill and whose columns represent the operations. Thus, $N(f, i)$ represents the number of times operation i was used when the fill was f. For each fill level, the probability of requiring operation i can be estimated by

$$\pi_i(f) = N(f, i)/\sum_i N(f, i).$$

The variance of the estimate can be estimated[5] by:

$$\mathrm{Var}[\pi_i(f)] = \pi_i(f)(1 - \pi_i(f))/\sum_i N(f, i).$$

The average cost per inward order of customers into the system at a given fill is found from

$$\mathrm{Cost}\ (f) = \sum_i C_i\pi_i(f),$$

where C_i is the cost of operation i.

Both the total cost function and its components given by the probability functions for each operation are of interest in the simulation analysis. The next section presents the results of our analyses.

V. SIMULATION RESULTS

The first part of this section considers the effects of customer movement for one assignment strategy—a break-as-needed policy. The second part analyzes different assignment strategies under various customer movement scenarios.

The effects of network configuration are not studied here. A 90-pair group was used for these results with 12 terminals of size 15 where each pair appears in exactly two terminals according to a standard multipling scheme. A small size pair group was chosen to keep the computation time low, since analytic models[6] suggested that the pair

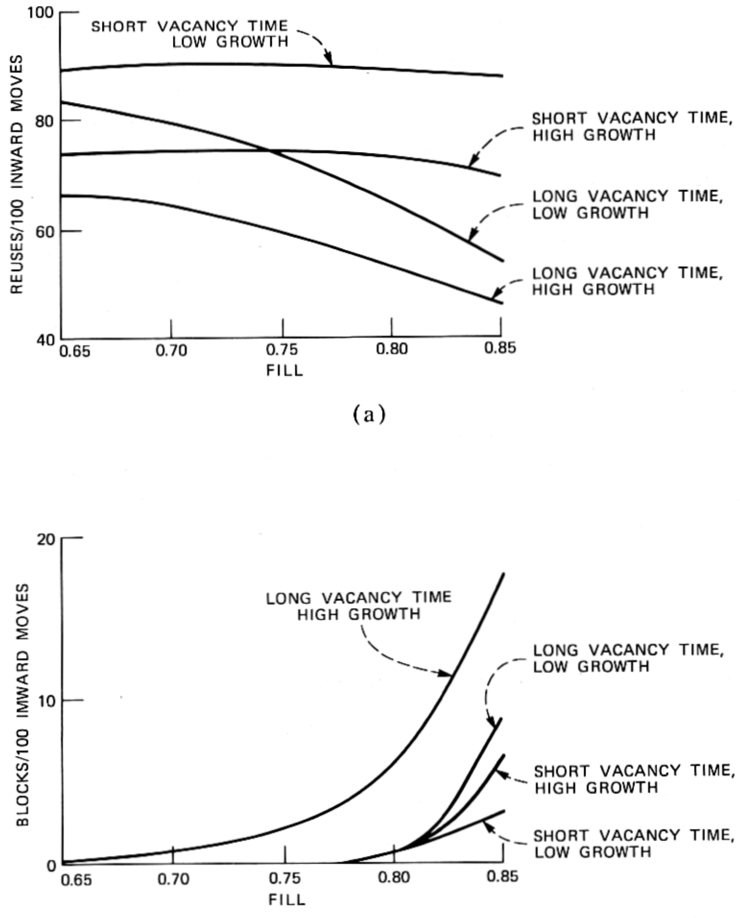

(a)

(b)

Fig. 6—(a) Effect of vacancy time and growth on reuse rates. (b) Effect of vacancy time and growth on blockage rates (*continued*).

group size is not an important factor for operating costs. The number of potential customers is 15 per terminal, for a total of 12 × 15 = 180 potential customers.

In this section, the reuse, BCT, and blockage rates are expressed as occurrences per 100 inward orders. The costs are expressed as average dollars per inward order. A blockage is defined as any inward order not served by a reuse, spare assignment, or BCT. The BCT rate includes CTs broken in conjunction with an LST or WOL.

All these costs and rates are highly dependent on the congestion of the network. At higher fills, there is less probability of having a reuse and consequently higher probability of having LSTs, WOLs, and REs,

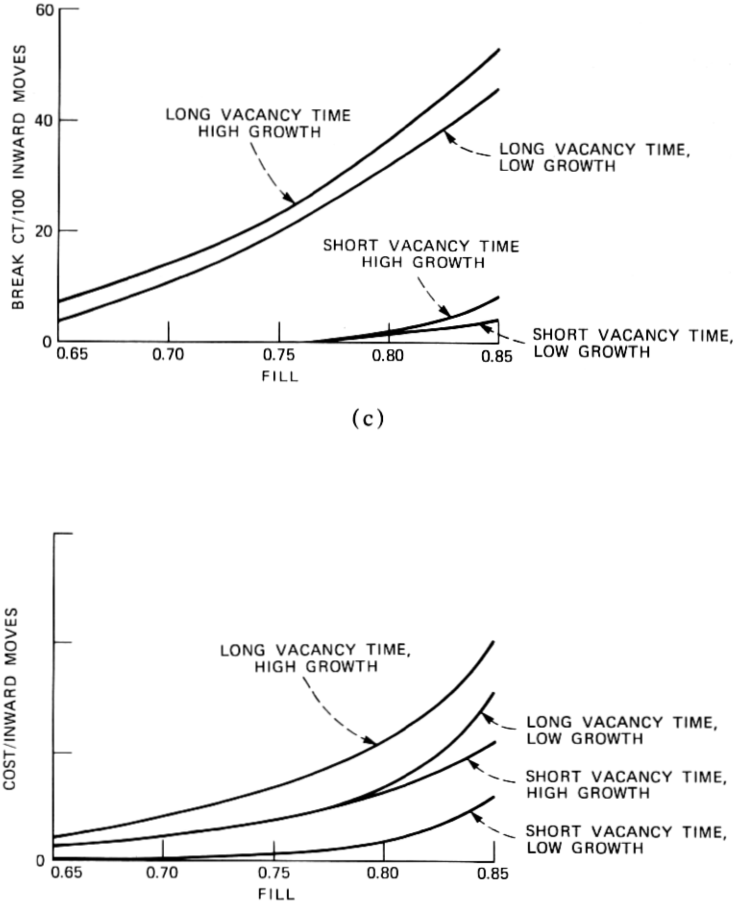

(c)

(d)

Fig. 6 (*continued*)—(c) Effect of vacancy time and growth on break CT rates. (d) Effect of vacancy time and growth on cost.

and the cost per inward order is higher. As a consequence, the cost and rate criteria are plotted as a function of fill and the curves are used for comparison. These curves are only shown up to fills of 0.85, since relief (capacity expansion) is generally provided before the network reaches higher fills.

5.1 Customer movement effects

Figures 6a to 6d show the significant effect that the vacancy time and growth parameters have on the graphs of costs and activity versus fill. For reuses (Fig. 6a), growth rate is the more important factor at low fills (since no CTs will need to be broken at low fills, all nongrowth inward orders will be reuses), while the penetration determines how quickly the reuse rate falls as fill increases. Blockages (Fig. 6b) increase with increasing growth or decreasing penetration (increasing vacancy time), with penetration having the larger effect. The BCT rate (Fig. 6c) is determined almost entirely by penetration. This is due to the fact that high penetration means low vacancy time, and CTs which are reused quickly cannot be broken. The cost per inward order curve (Fig. 6d) combines the above results, using the appropriate cost factors. At low fills, there are no blockages or BCTs, so the effects are the same as for reuses, with growth the key variable. At high fills (above 0.75), blockage effects predominate, and penetration is most important.

Another customer movement characteristic of interest is abandonment. The effects on cost of a 0.10 abandonment probability is illustrated for the high penetration/high growth case in Fig. 7a. Although abandonment causes fewer reuses and many more BCTs, the blockage rate is unchanged. When customers return to an abandoned premise, the unreusable CT can always be broken to provide service, so blockage rates are unaffected. Since BCTs are low-cost operations, the difference in cost due to abandonment is small. For other assignment strategies where CTs may not be broken, abandonment can be expected to have a larger impact.

The remaining characteristic studied here is second-line penetration. The total growth rate was the same as the high growth used before, with 95 percent of this growth consisting of first lines. This will be considered a second-line penetration of 5 percent. The first lines had the same short vacancy times and occupancy times used before. Second lines, however, had an occupancy time equal to half the first line occupancy time and a vacancy time 40 times larger. The resulting costs are compared with the one-category high growth/high penetration case in Fig. 7b. Although the reuse rate is slightly lower and the BCT rate higher when second lines are present, the effects on total cost are small.

In summary, the penetration and growth characteristics have a large effect on the cost per inward order (and the activity rates that comprise

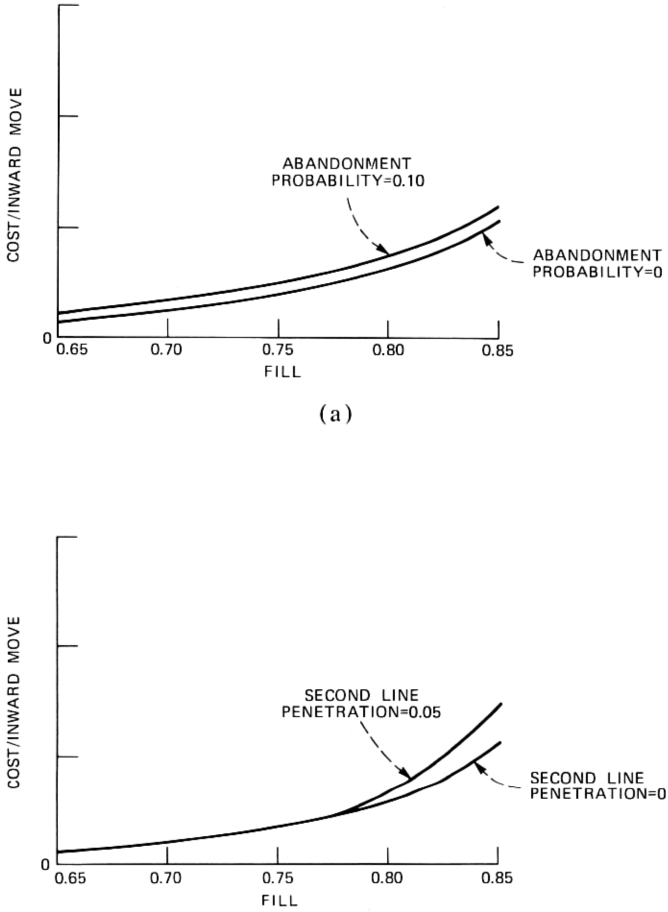

(a)

(b)

Fig. 7—(a) Effect of abandonment on cost (short vacancy time, high growth). (b) Effect of second-line penetration on cost (short vacancy time, high growth).

cost); the abandonment rate and second-line penetration have a small effect. These results are most important for use in the development of analytic models of operating cost. They also provide some guidance in assessing what scenarios need to be developed for the assignment strategy comparison described in the next section.

5.2 Assignment strategies effects

The four different strategies described in Section 3.3 were investigated: Break as needed, nonzero holding time, break as a last resort, and a reassignable policy. The results for the CT strategies under two different environments (long vacancy time/low growth, and short vacancy time/high growth) are illustrated in Figs. 8 and 9. A longer

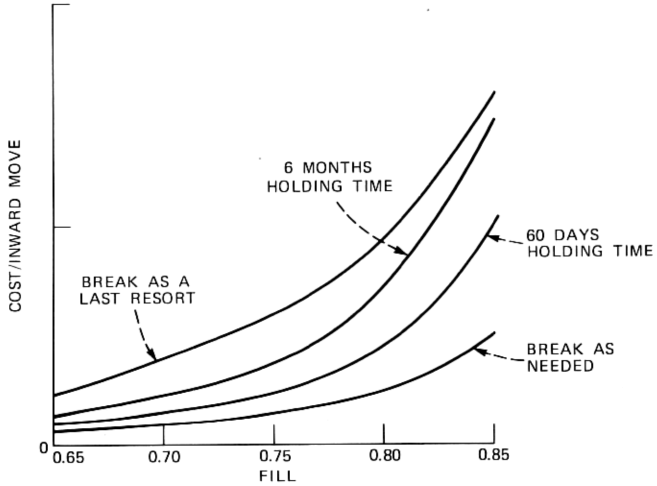

Fig. 8—Effect of CT holding time on cost (long vacancy time, low growth).

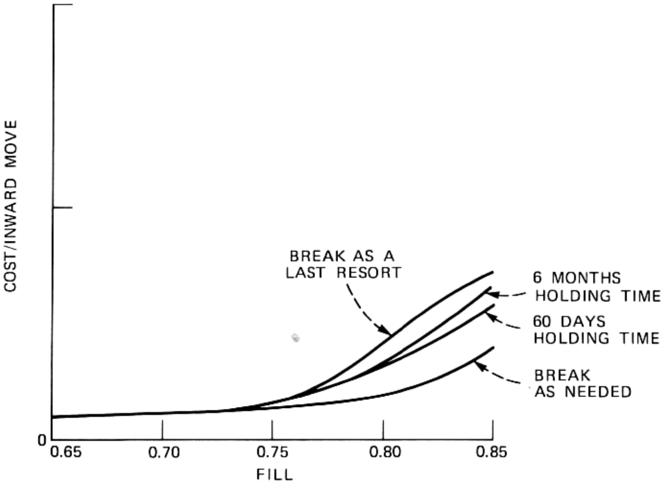

Fig. 9—Effect of CT holding time on cost (short vacancy time, high growth).

holding time tends to increase the CT reuses and decrease the break CTs, but also increases the number of blockages. In the long vacancy time case, the blockage effect dominates, so the cost is lowest with the break as needed policy. The cost difference between policies becomes more pronounced as the fill increases. For the short vacancy time case, the reuse savings with a holding time policy are more substantial, and at low fills, the cost curves are about the same. However, at higher fills, the blockage effect again dominates, and again the break-as-needed policy becomes the lowest cost policy.

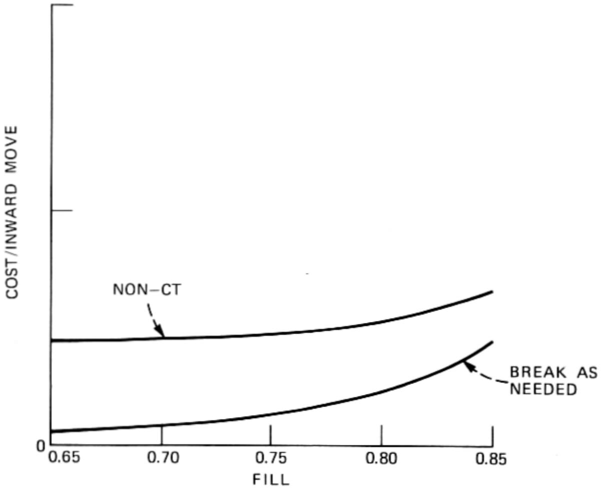

Fig. 10—Effect of non-CT policy on cost (long vacancy time, low growth).

It appears, therefore, that a CT policy of breaking as needed is the best one to use. But how does it compare to a non-CT policy where all lines are disconnected on an outward order? Figure 10 shows this comparison for the long vacancy time, low growth case. Even though reuses are low due to the long vacancy time, the CT policy is clearly better at all fills. The greatest savings from using a CT plan are at the low fills when the reuse rate is highest. The savings are due mostly to the avoidance of connect and disconnect costs, although the blockage rates are slightly lower under the CT plan.

VI. DISCUSSION

The simulation model described in this paper has been put to a number of uses in studies of the loop network. In addition to providing an analysis of the operating costs tradeoffs in the pair assignment problem, it was subsequently used for analyzing assignment policies in areas with seasonal demand and areas with mixed apartment and single-family housing. Analytic models of loop operating costs have also been developed[6] for use in an analysis of optimal timing for capacity expansion. These analytic models were based in part on the results of our simulation analyses, particularly the customer movement effects described in Section 5.1.

Another area in which this simulation (with appropriate modifications) might be used would be in a study of policies for controlling defective (unusable) pairs in the loop plant. Defective pairs are generated both from environmental causes and as an accidental byproduct of other network rearrangements. However, fixing a defective pair is one way of providing an available pair to serve a customer. Thus our simulation model could be used to determine where the repairing of a defective pair should fall in the assignment preference ordering.

The results of our simulation model were validated using actual field data and assignment personnel. Those interested in our validation procedures and results should consult Ref. 7.

REFERENCES

1. D. B. Luber, "Serving Area Sizing for Conversion Projects," IEEE Trans. Commun. *COM-27*, No. 5 (May 1979).
2. J. Freidenfelds, "A Simple Model for Studying Feeder Capacity Expansion," B.S.T.J., 57, No. 4 (April 1978).
3. P. J. Kiviat, R. Villanueva, and H. M. Markowitz, *SIMSCRIPT II.5 Programming Language*, Los Angeles: CACI, 1973.
4. B. L. Marsh, "The Telephone Plant—An Overview," presented at the 1976 International Conference on Communications, Philadelphia, 1976.
5. N. L. Johnson and F. C. Leone, *Statistics and Experimental Design*, New York: John Wiley, 1964.
6. W. L. G. Koontz, "An Approach to Modeling Operating Cost in the Loop Network," B.S.T.J., *57*, No. 4 (April 1978).
7. H. T. Freedman, T. R. Harms and W. L. G. Koontz, "A 'Human Simulation' Validation of a Telephone Loop Network Simulation Model," Proceedings of the Winter Simulation Conference, San Diego, December 1979.