The Bell System Technical Journal
Vol. 55, No. 10, December 1976
Printed in U.S.A.

Last-Trunk-Usage Measurements in Step-by-Step Switching Systems

By B. R. LaCAVA, W. D. MILLER, and B. YAGED

(Manuscript received July 22, 1976)

A measurement approach is described which provides the capability of detecting traffic congestion in the graded multiples of a step-by-step switching system. The basic idea is to measure the load carried on the last trunk of each graded multiple and, using the techniques described, to determine the congestion level in each grading. The method, which requires only one lead per grading, works because Bell System grading patterns, which differ in structure depending on size, all have a common last-choice trunk and because the load carried on the last-choice trunk increases as the congestion in a grading increases. The basic idea and the interpretation of last-trunk-usage data in the presence of day-to-day variations in the offered load for configurations with and without rotary out-trunk switches are described in detail. Last-trunk-usage measurements are compared with other possible traffic measurements on gradings in terms of effectiveness in detecting various service impairments. The statistical accuracy of load and blocking estimates from last-trunk-usage measurements is also discussed. The result is a simple, effective measurement technique with the combined advantage of rapid detection of gradings with service problems and relative ease of implementation.

I. INTRODUCTION

Because of physical limitations, the trunks that connect the successive switching stages in a step-by-step switching system are arranged in sets of partial-access patterns called graded multiples (also called gradings, graded subgroups, or subgroups). The Bell System uses standard patterns for the gradings; a representative configuration is shown in Fig. 1. A call arriving hunts across the 10 trunks marked as heavy lines to find an idle trunk; if these 10 are busy, the call is blocked with a reorder tone returned to the customer. Details of the traffic flow through the system are given in Ref. 1. A long-standing problem in step-by-step offices has been the difficulty in detecting traffic congestion in the graded multiples, since detailed measurements are not obtained.

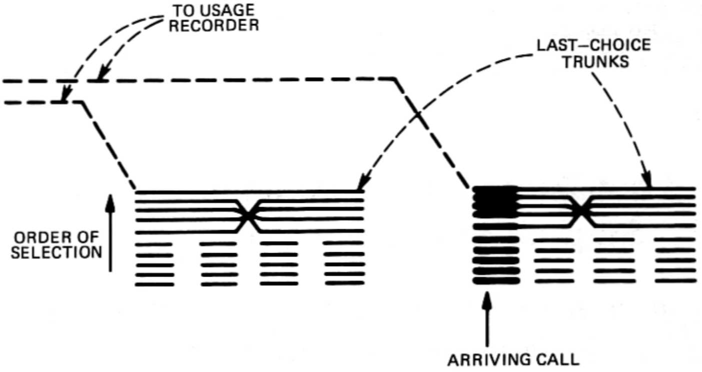

Fig. 1—Typical step-by-step subgroups showing limited trunk access for an arriving call.

In this paper, a new measurement approach is described which provides an economical and effective means for identifying specific graded multiples with substandard blocking levels. The basic idea is to capitalize upon the fact that, although the Bell System patterns differ in detailed structure depending upon the number of trunks, they all have a common last-choice trunk. Thus, the traffic performance, i.e., fraction of calls blocked, of a graded multiple can be monitored by measuring the load carried on the last-choice trunk (see Fig. 1). As the congestion in a grading increases, the load carried on the last trunk also increases. By properly interpreting the last-trunk usage (LTU), it is possible to infer the congestion level of each grading. Since only one connection per grading is required, the method makes efficient use of measurement equipment. The data can be collected by a variety of available usage-measuring equipment. Notice that the wiring can be validated by making the last trunk busy during nonpeak hours and checking for the correct usage measurement. Also, it is possible to detect (and thus avoid possible data misinterpretation) whenever subgroup performance cannot be determined due to the last trunk having been taken out-of-service, since its usage will appear as 36 CCS (hundred call seconds) in each hour.

Notice that this measurement approach is an application of a more general concept wherein the traffic performance of a set of servers is evaluated by measuring the behavior of one carefully selected element of the set. It is possible that this general concept may have broader applicability and should be considered in designing new measurement techniques.

Section II provides an overview of the measurement procedure and describes the model used to develop the last-trunk-usage procedures.

Section III describes the additional considerations required when the subgroups access rotary out-trunk switches (ROTS). Section IV discusses the use of LTU and other possible traffic measurements in detecting congestion problems—general overloads, imbalances between subgroups, and imbalances within subgroups. The LTU measurement procedure does not detect equipment irregularities. Section V relates the statistical accuracy of blocking and offered-load estimates to inherent variations in LTU data. Finally, Section VI discusses the application of the method.

II. OVERVIEW OF LAST-TRUNK-USAGE PROCEDURES

The LTU monitoring procedures were developed using a computer simulation model of step-by-step (S × S) graded multiples.[1] The model has the following properties: Poisson arrivals, inherent load-

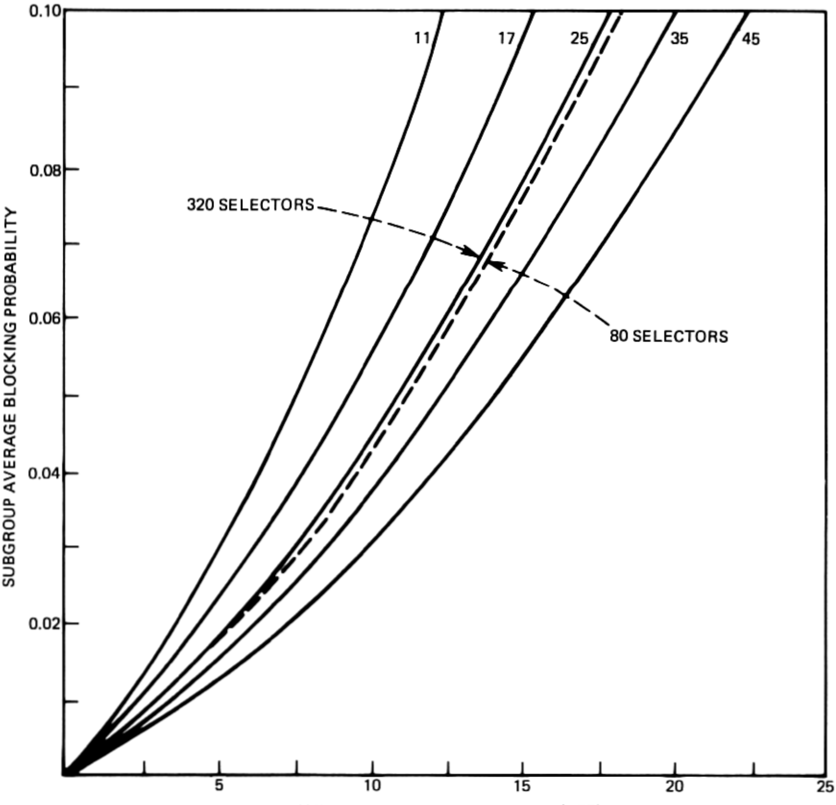

Fig. 2—Subgroup average blocking probability versus average last-trunk usage for 11-, 17-, 25-, 35-, and 45-trunk gradings with low day-to-day variation in offered load.

balancing, blocked calls cleared, and a negative exponential distribution for holding times. Retrials are not included in the model. A representative set of grading patterns was selected for detailed study. For each grading, the simulation was initially used to determine the equilibrium blocking and last-trunk usage over an appropriate range of offered loads.

Figure 2 shows the relationship between average blocking and average LTU for several gradings. In this paper, "average blocking" means "average call congestion." The use of average blocking and average LTU, rather than their equilibrium values, is appropriate since it is assumed that the offered loads will vary on different days according to the *low* day-to-day variation model;[2] measurements during a time-consistent hour must be averaged over several days to provide stable results (and reduce the possible impact of a single long-holding-

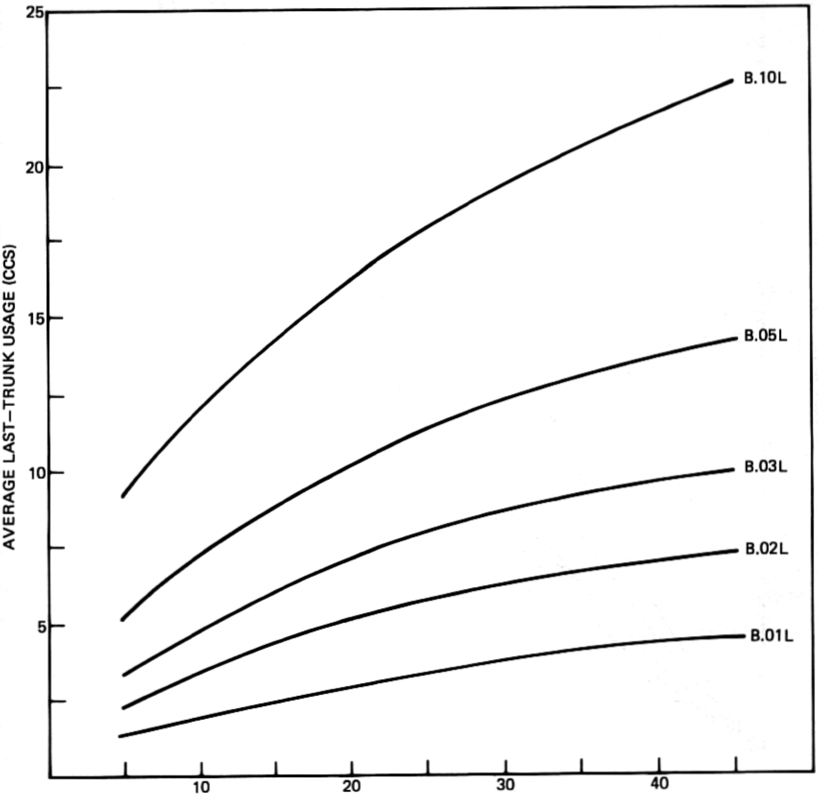

Fig. 3—Average last-trunk usage versus the number of trunks in the graded multiple for blocking at B.01L, B.02L, B.03L, B.05L, and B.10L.

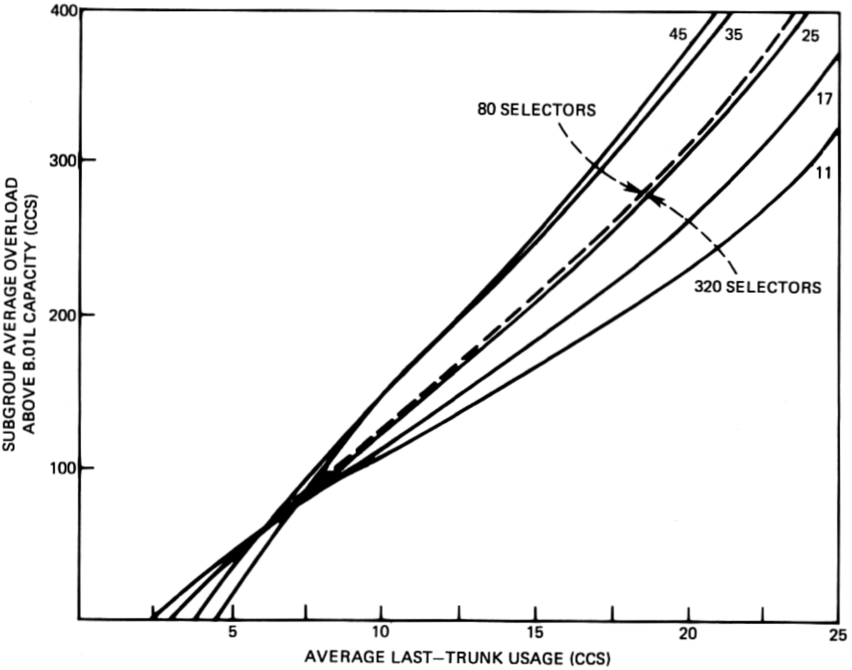

Fig. 4—Subgroup average overload versus average last-trunk usage for different graded multiple configurations (non-ROTS accessing) with low day-to-day variation in offered load.

time call). Figure 2 shows that last-trunk-usage data provide a good indication of subgroup performance, or blocking level. The LTU-blocking relation is somewhat dependent on the number of trunks in the grading, but for a given number of trunks, is generally not very sensitive to the number of accessing selectors. Because of this insensitivity, selector dependence was not studied in detail. Note that Fig. 3 uses the same data as Fig. 2, but relates average LTU and number of trunks in the grading for fixed average blocking levels.

The difference in offered load from that required for B.01L, a frequently used objective in S × S systems, is plotted against average LTU in Fig. 4 for the same grading sizes as in Figs. 2 and 3. (The notation B.XXXV means that "Erlang-B" assumptions were used, that the blocking probability for no day-to-day variation or the average blocking probability for results with day-to-day variation is O.XXX, and that V indicates the amount of day-to-day variation with V blank for no day-to-day variation and L, M, or H for low, medium, or high day-to-day variation, respectively.) The curves in Figs. 2 and 3 show that the subgroup overload increases in an approximately linear fashion with LTU, up to a blocking of roughly B.10L.

Again, results are not very sensitive to the number of selectors. However, errors in offered load or subgroup overload estimates may arise if certain service-affecting problems exist, but are unknown; these will be discussed in Section IV.

As mentioned, the results indicate that curves of blocking vs LTU, constant blocking curves for LTU vs the number of trunks, and subgroup overload vs LTU are not strongly dependent upon the number of selectors that access the grading. Consequently, the general problem of studying all 155 standard combinations of selectors and subgroup trunk patterns found in the Bell System was simplified by ignoring selector dependence and considering only the different grading configurations. Note that the relationship of LTU to offered load, in contrast to the overload in CCS, is somewhat dependent upon the number of selectors accessing the graded multiples; this occurs because the capacities of the gradings are dependent on the number of accessing selectors.[1]

III. EFFECTS OF ROTS ON LTU PROCEDURES

Rotary out-trunk switches (ROTS) are used in step-by-step systems to concentrate traffic from several graded multiples, hereafter called the access subgroups, onto a single outgoing trunk group, whose occupancy is higher than those of the access subgroups. The model of traffic flow in a graded multiple used in Section II is not appropriate in the presence of ROTS for the reasons described. ROTS access subgroups are generally engineered at one-tenth the blocking desired for the entire ROTS system to ensure adequate access to the outgoing trunks. A representative local ROTS configuration is shown in Fig. 5 for m access subgroups (10-trunk, full-access subgroups in this case) going to n ROTS subgroups, each consisting of about 20 to 30 switches. The trunks from the access subgroups are spread across the ROTS subgroups according to specified patterns to distribute traffic across the ROTS trunks and to provide access from a given selector to the several ROTS subgroups. If the number of outgoing trunks per ROTS subgroup is less than 21 (or 22 in toll groups), the outgoing trunks from each ROTS subgroup may be multipled to adjacent subgroups.

In local ROTS groups, approximately 40 percent of the trunks in the ROTS access subgroups are wired directly to the outgoing trunks, bypassing the ROTS switches. Thus, certain outgoing trunks can be accessed either directly from trunks in the ROTS access subgroups as well as indirectly through a ROTS switch. If the latter occurs, the directly linked trunk in the ROTS access subgroup is made busy to prevent a second call from seizing the outgoing trunk. Subsequent calls arriving at the access subgroup that find a directly linked trunk busy will use trunks higher in the hunting order (closer to the last trunk).

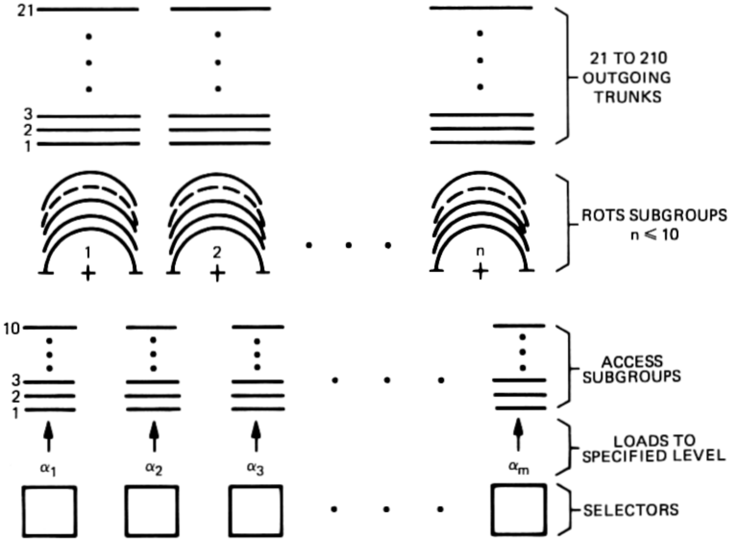

Fig. 5—Representative ROTS group configuration.

In addition, if all outgoing trunks in a ROTS subgroup are busy, all idle ROTS switches in that subgroup are made busy. This has the combined effect of forcing traffic higher in the access gradings and, if such a switch is connected to the last trunk of an access subgroup, of generating busy-back (non-call carrying) usage on that trunk. Any call searching for an idle trunk, which progresses to the last-choice trunk and finds the last-choice trunk busy from either type of usage, is blocked.

Consequently, the access gradings must carry the actual offered load plus the induced busy-back load. Since any calls that are blocked in the ROTS system receive reorder tone from the selectors of the access subgroups, the blocking observed on the access subgroups equals the system blocking, even though the access subgroups are usually engineered at one-tenth the desired system blocking. In addition, the intricate access arrangement results in significant interaction between the traffic parcels offered to the different access subgroups. The result is that the performance of a subgroup is influenced by its own offered load, the congestion level on the outgoing trunks, and the performance levels of the other access subgroups.

Therefore, it was not clear whether the relationships between LTU and blocking shown in Section II applied in the presence of ROTS. To determine this, a computer simulation of ROTS configurations modeled by Neal was studied.[3] The traffic assumptions of the graded multiple simulation model described in Section II are applicable here.

For the different grading patterns and different ROTS configurations,

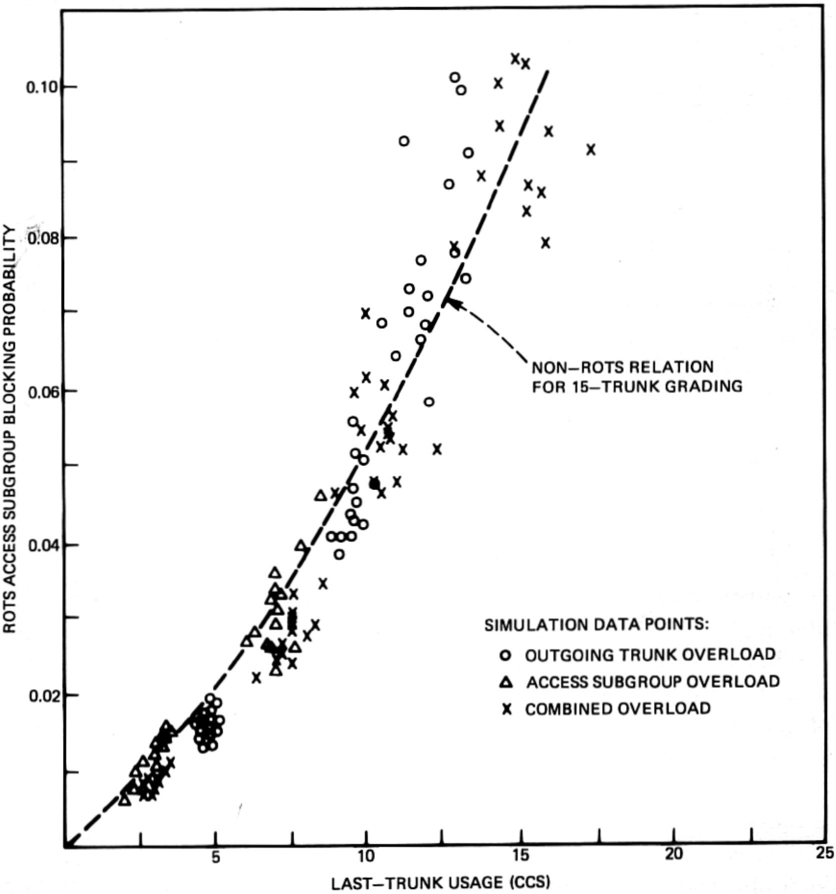

Fig. 6—ROTS access subgroup blocking versus last-trunk usage for ROTS group of 15 subgroups of 40/15 with up to 135 outgoing trunks and no day-to-day variation in offered load–general overloads.

it was found from simulation that approximately the same relationships between LTU and blocking apply as when the graded multiples do not access ROTS, giving acceptable estimates up to ten-percent blocking. Figure 6 shows a curve which relates LTU and blocking for non-ROTS subgroups, as well as data points from simulation studies of a configuration with 15 access subgroups of 15 trunks each and up to 135 outgoing trunks; the offered loads to the subgroups were equal (or balanced). By varying the offered load and the number of outgoing trunks, arrangements were investigated in which the access subgroups were under-engineered, as well as situations where the outgoing trunk group was under-provided. The data points clustered about the non-ROTS curve, for blocking levels up to about B.10–. Significant differences between the data points and the curve occur above ten-percent blocking because of interactions caused by busy-back usage.

These ROTS, non-ROTS comparisons are based upon *no* day-to-day variations since substantially more computer time would be required to study *low* day-to-day variations explicitly. However, the relative comparisons still hold for *low* day-to-day variations, and, thus, Fig. 2 can be used to estimate the average blocking of subgroups which access ROTS.

The same non-ROTS curve is shown in Fig. 7, but the data points represent ROTS simulation results where the offered loads to the subgroups are unequal (imbalanced). In this case, the data points were more clustered about the non-ROTS curve than in Fig. 6. These effects were confirmed for additional ROTS configurations of 11-, 21-, 25-,

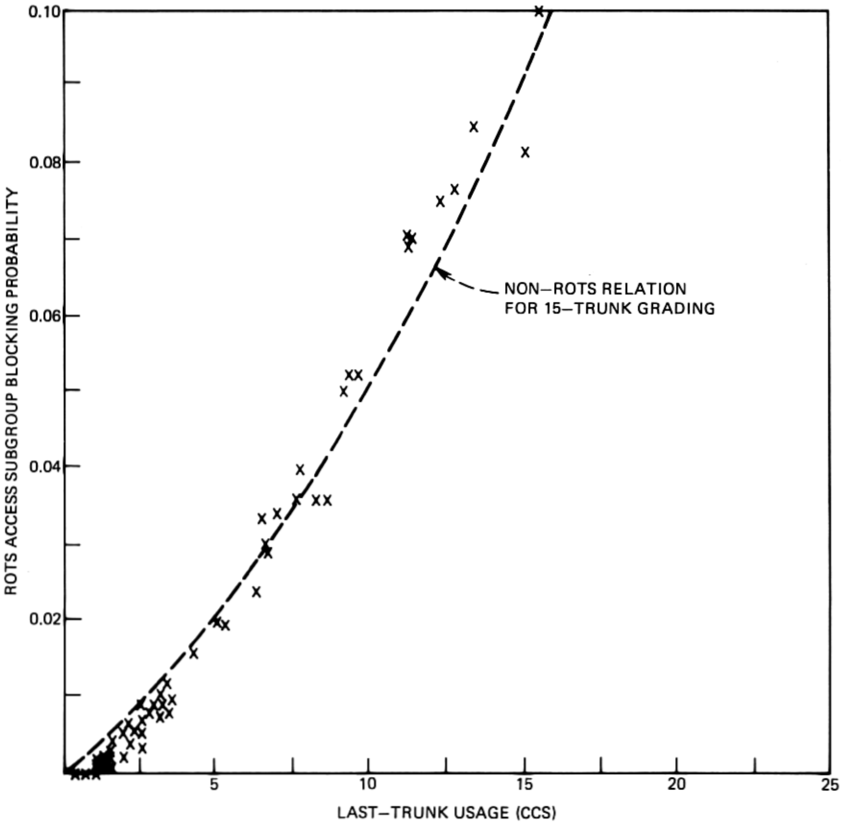

Fig. 7—ROTS access subgroup blocking versus last-trunk usage for ROTS group of 15 subgroups of 40/15 with 105 outgoing trunks and no day-to-day variation in offered load–imbalanced loads in access subgroups.

35-, and 45-trunk subgroups with between 34 and 210 outgoing trunks, demonstrating that the LTU-blocking relation of Section II could apply as well to ROTS configurations under a wide variety of conditions. An important feature is that the LTU measurement monitors the blocking of the access subgroups, which is, in fact, *equal to* the blocking for the entire configuration, the quantity of interest from a service standpoint. The impact of the scatter in ROTS data points on blocking estimates is discussed in Section V.

It is possible to measure the actual last-trunk call usage, excluding busy-back usage. This procedure, examined via simulation, did not give as good agreement to the non-ROTS LTU-blocking relations for general overloads, in that this type of overload affects the call usage on the last trunk for a given blocking level. The majority of points in the scatter diagram of Fig. 8 for access subgroups of 15 trunks fall

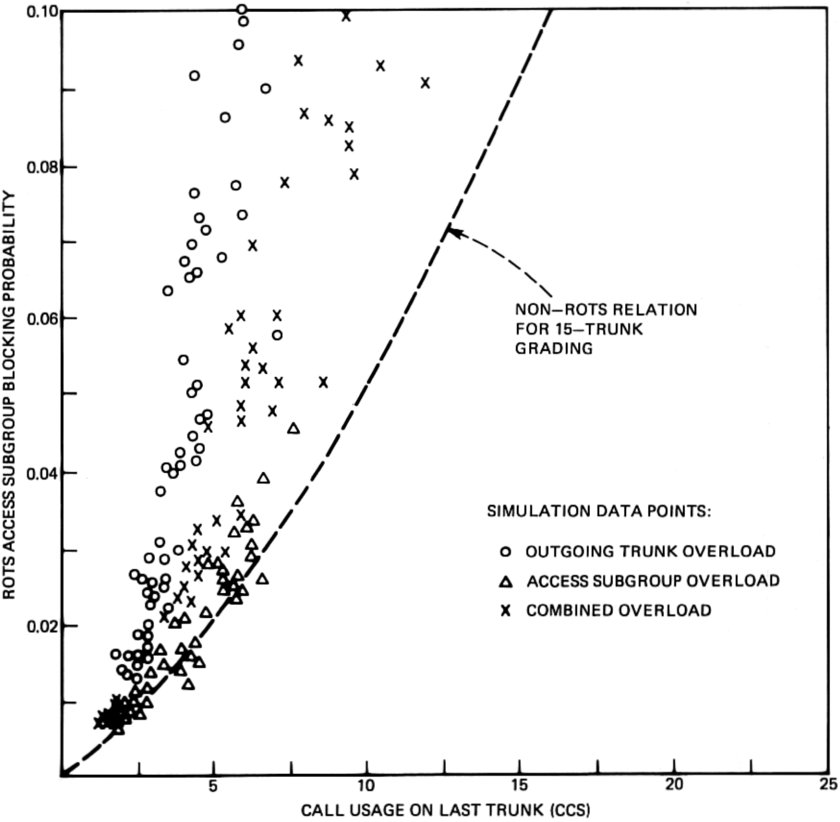

Fig. 8—ROTS access subgroup blocking versus call usage on last trunk (excludes busy-back usage) for ROTS group of 15 subgroups of 40/15 with up to 135 outgoing trunks and no day-to-day variation in offered load-general overloads.

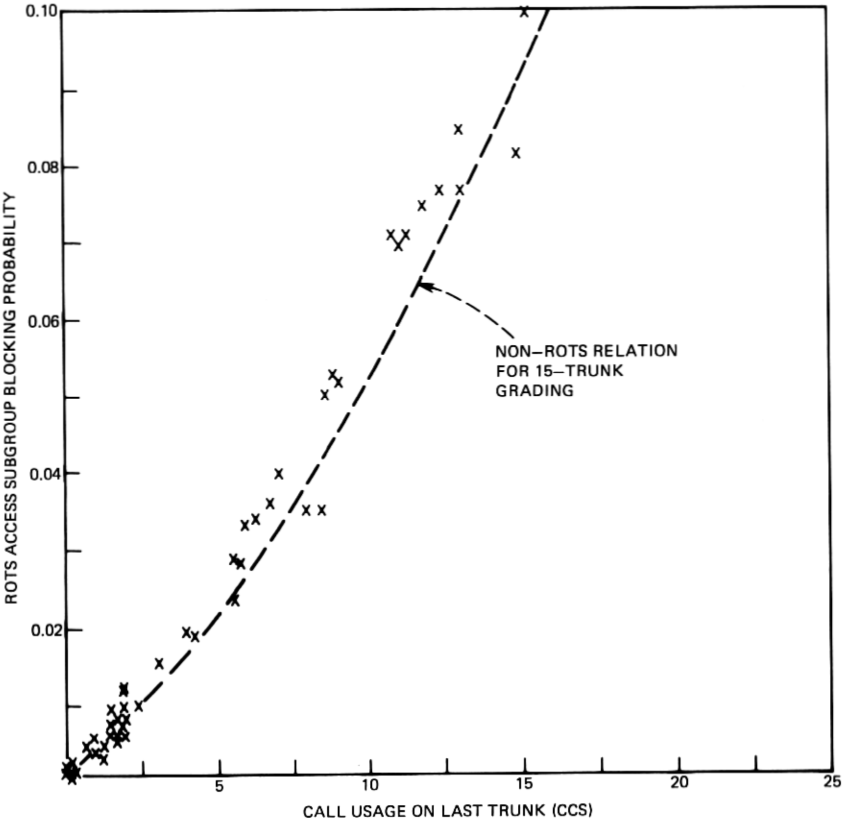

Fig. 9—ROTS access subgroup blocking versus call usage on last trunk (excludes busy-back usage) for ROTS group of 15 subgroups of 40/15 with 105 outgoing trunks and no day-to-day variation in offered load—imbalanced loads in access subgroups.

to the left of the non-ROTS relation for a 15-trunk subgroup. There was less disagreement for the case of load imbalances in accessing subgroups. Figure 9 shows that the scatter diagram for 15-trunk subgroups follows the corresponding non-ROTS relation, but that the simulation points generally fall to the left of the curve, implying a small bias. Again, similar results were observed for both general overloads and imbalances in simulating other ROTS configurations.

Last-trunk-busy (LTB) registers on graded subgroups are wired to exclude made-busy counts. Thus, the results in Figs. 8 and 9 also apply to LTB counts in the presence of ROTS, where the abscissa would be measured in LTB registrations rather than call usage on the last trunk. Notice that reconnecting the LTB register to include busy-back effects complicates data interpretation, since it appears difficult to estimate the correct mean holding time when busy-back usage occurs.

While the average LTU-blocking relations of Figs. 2 and 3 provide acceptable average blocking estimates for ROTS configurations, the subgroup average overload-LTU relation of Fig. 4 does not always apply. Not only is it difficult to accurately estimate the magnitude of the overload, but whether the overload occurs in the access subgroups, the outgoing trunk group, or both is not evident. Once the LTU procedure detects blockages, additional measurements must be taken to determine the cause of the blockage. If the outgoing trunk group is adequately engineered, as determined from total carried load or all-trunks-busy measurements, the access subgroup is limiting; in this case, Fig. 4 can be used to estimate the subgroup overload. If the outgoing trunk-group capacity is limiting and additional trunks are added, additional LTU measurements should be taken to see if the access subgroup capacity is still sufficient after augmenting the outgoing trunk group.

In summary, the basic LTU-blocking relations apply to ROTS configurations, although some moderate additional estimation uncertainty arises from the scatter shown in Figs. 6 and 7. Thus, the LTU procedure appears to be an acceptable performance-measurement tool, even though it alone cannot determine the source of the congestion problem. An additional measurement on the outgoing trunks is needed to decide what corrective action is required for a congested ROTS configuration.

IV. STEP-BY-STEP SERVICE-AFFECTING PROBLEMS

4.1 Overview

LTU procedures can be used to detect several service-affecting problems, namely: general overloads, load imbalances, and out-of-service trunks. In each case, the problem is described and how LTU measurements are effective in detecting the problem is shown; the ability of carried load or last-trunk-busy measurements to detect the problem is also discussed. First, a general understanding of the principles used in engineering S × S subgroups is needed.

Consider a typical stage in a S × S train (Fig. 10), where the incoming traffic load is directed to n groups of selectors. Each selector group has access to a set of trunks on each level. In practice, the selector grouping for different levels may not be identical. Two basic assumptions in S × S engineering are: (*i*) incoming traffic is distributed equally between selector groups (i.e., $L_1 = L_2 = \cdots = L_n$) and (*ii*) each selector group offers the same proportion of the incoming traffic to a given level j (i.e., $a_{1j} = a_{2j} = \cdots = a_{nj}$). As a result, all subgroups on a given level generally have the same number of trunks. Based on these assumptions, we need only measure the total carried

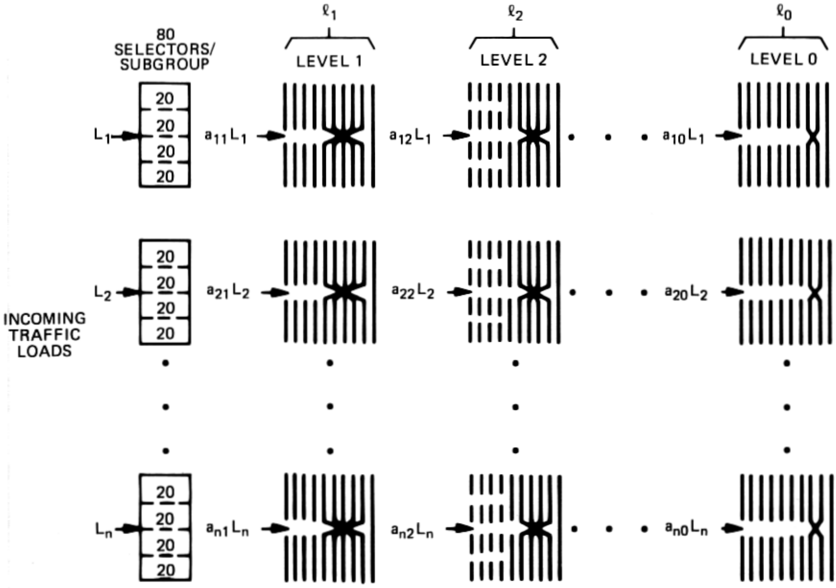

Fig. 10—Step-by-step provisioning and monitoring problems: aggregate measurements of total carried loads ℓ_i.

load on all trunks on a level (e.g., l_1) to estimate the carried load for each subgroup (e.g., l_1/n) and to determine the number of trunks required to meet a stated service objective.

4.2 General overloads

A general overload occurs when all subgroups on a level have excessive blocking, as indicated by comparing the LTU values with a threshold determined from Fig. 2. The amount of overload on each subgroup can be determined from Fig. 4 if the subgroups do not access ROTS. If the subgroups access ROTS, then additional measurements, such as outgoing trunk carried load or all-trunks-busy counts on the ROTS subgroups, are required to determine if the blockage arises in the access subgroups or the outgoing trunks.

Carried-load measurements on subgroups within a S × S train are not commonly made because of the excessive number of measurement leads required, but are usually made on interoffice trunks. Such load measurements should detect a general overload in the non-ROTS case and in the ROTS case only if the problem is insufficient outgoing trunks. Last-trunk-busy registers, although not generally available, should indicate overloads, assuming that the correct mean holding-time is used and that the last trunk is in service. With ROTS, LTB counts are limited by the effect described in Section III.

4.3 Load imbalances

Load imbalances arise when the incoming loads per selector vary over a wide range (e.g., trunks from No. 4A crossbar or S × S offices) and the trunks are not uniformly distributed across selector groups, or trunks from different offices, which exhibit different calling patterns to different levels, are not well distributed. In either of these cases, the traffic offered to a level is not the same for each selector group (e.g., $a_{1j}L_1 \neq a_{2j}L_2 \neq \cdots \neq a_{nj}L_n$), causing some subgroups to be overloaded while others are underloaded. Imbalances lead to higher blocking to selected incoming trunks as well as an increase in the average congestion across all trunks.

Last-trunk-usage procedures effectively detect blockages arising from load imbalances between subgroups, since the congestion level of each subgroup is individually monitored and recorded. This results in the specific identification of congested subgroups for corrective action.

In some cases, particularly when the larger gradings are used, balance within a grading may become a problem. Measuring the LTU is also effective here. This follows both from the flatness of the B.01L curve and from the significant spread between the curves in Fig. 3. To illustrate, consider a 45-trunk grading. From Fig. 3, we interpret an average LTU of 4.5 CCS as about one-percent average blocking under the assumption that the offered load is balanced across the individual groups of first-choice trunks (legs) of the graded multiple. Suppose the maximum imbalance occurs, i.e., all traffic is offered to only one leg and the remaining receive no calls. In this case, only ten of the 45 trunks carry any traffic; the observed average blocking would be about 2.5 percent (average LTU of 4.5 CCS on ten trunks). Thus, if the LTU is controlled so that the indicated average blocking for a 45-trunk subgroup with balanced loads is about one percent, no parcel of traffic will see more than 2.5-percent average blocking. The difference between this value and the average blocking estimate, smaller when less extreme types of imbalances occur and when the subgroup has fewer than 45 trunks, is important only when a specific group of incoming trunks is focused on an overloaded leg. If incoming trunk groups are spread across the legs, each group of customers would experience about the average blocking (across legs) for the subgroup; this average blocking is accurately estimated by the LTU procedure. Thus, the LTU procedure provides reasonable control over imbalances among the legs, although it cannot assure maximum equipment utilization. However, the estimates of subgroup average overload shown in Fig. 4 always exceed the true value in the presence of imbalances within a subgroup.

Carried-load measurements are commonly obtained on interoffice trunk groups, but are generally aggregated across all subgroups on a level for the non-ROTS case or are obtained on the outgoing side of the ROTS switches when ROTS are used. In both cases, these composite carried-load measurements cannot detect imbalances, since the carried load actually decreases as the imbalance increases. As before, LTB registers are capable of detecting imbalances although they are not usually available in the field.

4.4 Out-of-service trunks

Out-of-service trunks that lead to an increase in subgroup congestion can be detected by the LTU procedure. An out-of-service trunk that does not lead to undesirable blocking levels (if the subgroup's offered load is sufficiently below engineered capacity so that loss of a trunk has no service impact) would not be detected by the LTU procedure. The shallow slope of the B.01L curve of Fig. 3 results in the average blocking estimate being insensitive to a modest number of made-busy trunks. Thus, the LTU procedure can detect such blockage, but not indicate the cause of the blockage (i.e., overload, imbalance within the subgroup, or made-busy trunks). With out-of-service trunks, the overload estimate of Fig. 4 always exceeds the true overload. Measurements of total-carried-load may not indicate a service problem caused by out-of-service trunks, particularly if the load measurements are aggregated across subgroups.

4.5 Equipment malfunctions

Malfunctions leading to abnormally short holding times (the "killer trunk" phenomenon) cannot be located with the LTU procedure.

V. UNCERTAINTIES IN BLOCKING AND LOAD ESTIMATES

The uncertainty in average blocking estimates is important in deciding when corrective action is required. The uncertainty in these estimates arises from three effects: variations in single-hour LTU measurements, day-to-day variations of offered load, and approximations used in the LTU procedure. Estimates of blocking and offered load from single-hour LTU data cannot be used, because they have a large coefficient of variation (CV); this is defined as the standard deviation to mean ratio of a random variable and indicates the "spread" of values of the variable about its average. Consequently, LTU data must be averaged over several hours to provide reliable estimates. This section shows that last-trunk-usage procedures using time-consistent busy-hour measurements can detect moderate to severe problems with acceptable confidence in five days whereas at least 20 days are needed for provisioning studies.

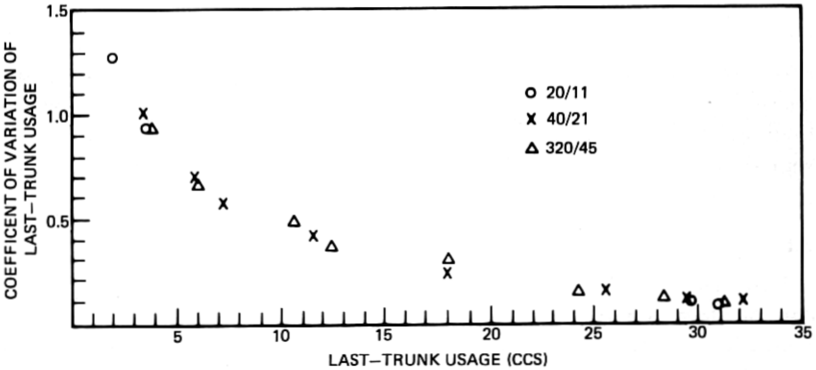

Fig. 11—Simulation points of coefficient of variation of single-hour last-trunk usage for several graded multiple configurations with no day-to-day variation in offered load.

For low subgroup blocking (B.01_), the cv for a single-hour LTU measurement is approximately one. The cv decreases to roughly 0.1 for subgroup blocking above B.10_. Simulation results for subgroups with 11, 21, and 45 trunks indicate that the cv of an LTU measurement is not strongly dependent upon grading size or pattern for Poisson traffic with *no* day-to-day variations (Fig. 11). The effects of day-to-day variation in the offered load and discrete sampling of the last-trunk usage on estimates of blocking and load were studied using procedures formulated by Neal and Kuczura,[4] Hill and Neal,[5] and Hill.[6] The standard deviation of the average blocking estimate for all subgroups (with blocking near B.01L) is approximately 0.01 for LTU values averaged over five hours (typically the same hour on five consecutive days). Thus, if five measurements give a blocking estimate greater than B.02L, we can be 84-percent confident that the true blocking is greater than B.01L. This result assumes a statistical model where the error is normally distributed with mean zero. Hence, LTU measurements during the busy hour enable relatively quick (i.e., five days) detection of subgroups with moderate to severe service-affecting problems. If five measurements give a blocking estimate less than B.02L, but greater than B.01L, the adjustment of performance levels to a B.01L objective should only be based on longer study intervals (e.g., 20 days) and a demonstrated practical need for such adjustments.

Five days of data are not sufficient to accurately determine the number of trunks required to meet a given blocking objective. With five days of data, the standard deviation in the average offered load estimate is about 40 CCS for a 21-trunk subgroup; thus, we can be

68-percent confident that the true offered load will be within 40 CCS of the estimated offered load, which corresponds to a possible deviation of two to three trunks from the number required to precisely meet the objective. The 68-percent confidence interval can be reduced to within one trunk by taking at least 20 days of data. It would seem that this is the minimum number of days needed for any provisioning study.

ROTS subgroups are subject to approximately the same level of statistical variations as non-ROTS subgroups. However, the scatter of Fig. 6 indicates that there will be some uncertainty introduced by using non-ROTS LTU-blocking relations for ROTS subgroups. Fortunately, the effect of possible approximation bias or additional variance is negligible when compared to the estimation error arising from other sources (measurement variation and day-to-day variations) during a five-day study.

VI. APPLICATIONS

6.1 Types of trunks

The preceding discussions have assumed that the graded subgroups were either connecting intraoffice selector stages or were one-way interoffice trunks. For two-way interoffice trunks, the LTU method applies to all trunk groups with ROTS, since the ROTS isolate the effects of the trunk selection method (whether from common-control or S × S equipment) at the far end from the LTU measurement. The LTU method is not applicable where selector subgroups directly access (instead of going through ROTS) two-way interoffice trunks, unless, at the far end, the trunks are chosen in the same hunting order as at the near end. For trunk groups where the procedure is not applicable, total carried-load measurements should be used, since these are typically small groups that are not susceptible to load imbalance problems because they are full or nearly full-access groups.

6.2 Measurement equipment

There is a wide variety of measurement equipment available for collecting LTU data. However, some Bell System measurement equipment found in small S × S switching entities such as community dial offices (CDOS) may not be ideal for LTU measurements. These electromechanical recorders have a relatively small number of output registers and are configured to report usage only on a grouped basis rather than an individual trunk basis. Last-trunk-usage measurements, requiring only one input lead per graded multiple, would utilize this equipment inefficiently assuming that a sufficient number of output registers were available. For some cases an option to record usage on individual trunks could be added. In general, a more viable alternative would be to use other measurement equipment that is suitable for the collection of LTU data.

6.3 Data collection and analysis

For each level under study at a particular selector stage, last-trunk-usage data should be collected during the time-consistent level-congestion busy hour; i.e., the time-consistent hour in which the total LTU for the level (obtained by summing over all subgroups on the level) is greatest. This hour can be determined by collecting data in several candidate hours during a study period and then choosing the time-consistent hour in which the total LTU for the level is greatest. In some cases, it may be desirable to analyze LTU data outside of the time-consistent level-congestion busy hour to detect focused overloads which may occur in other time periods. Of course, the data should be collected in the busy season to maximize the effectiveness of the measurement.

For the study period, the average busy-hour LTU for each subgroup can be calculated and used with Table I to estimate that subgroup's average blocking. This table, related to Fig. 2, provides a summary of the average last-trunk-usage in CCS at increasing blocking levels (B.005L to B.10L) for bands of subgroup trunk sizes. A last trunk that continually appears as 36 CCS is most likely to be out-of-service.

For a five-day "quick test," any subgroup which exceeds B.02L is a candidate for corrective action. For 20 days of data, subgroups exceeding B.01L should be considered for correction. Lower blocking thresholds may be appropriate outside the busy season. In all cases, additional LTU measurements should be collected after any corrective action to ensure that the blockage is eliminated.

Table I—Average value of last-trunk usage in CCS for different blocking levels *

Number of Trunks in Subgroup	Average Blocking (Percent)					
	0.5	1	2	3	5	10
4 to 9	0.5	1.5	2.5	3.5	5.0	10.0
10 to 14	1.1	2.0	3.5	5.2	7.8	13.0
15 to 19	1.4	2.5	4.6	6.3	9.3	15.5
20 to 24	1.6	3.0	5.3	7.2	10.6	17.2
25 to 29	1.8	3.2	5.7	7.8	11.5	18.2
30 to 34	1.9	3.6	6.2	8.5	12.4	19.5
35 to 39	2.1	3.8	6.7	9.1	13.1	20.5
40 to 45	2.4	4.2	7.3	9.8	14.1	21.7

* Assumes low day-to-day variations.

6.4 Corrective action

When LTU measurements detect congestion, additional investigation is necessary to find the cause of the blockage. Common problems, such as general overloads and load imbalances, were discussed in Section IV. Estimates of subgroup overload (in CCS) for a given average last-trunk-usage may be useful in determining the degree of corrective action for all major problems except general overloads on ROTS groups and maintenance-associated phenomena.

VII. SUMMARY AND CONCLUSIONS

The last-trunk-usage procedures are an application of the general concept of observing a subset of the elements of a traffic system to estimate the performance of the entire system. The procedures are effective because all Bell System standard gradings have a common last-choice trunk and because the load carried on that trunk is directly related to the traffic congestion in the subgroup. This provides a direct measure of subgroup performance, since general overloads and imbalances between subgroups or within subgroups, as well as possible made-busy trunks that have significant impact upon service, cause increases in a subgroup's LTU. In fact, the average load carried on the last-choice trunk is directly related to the service level, and in the range of primary interest (B.01L to B.10L) is not very dependent upon the cause of the degradation. These results apply to both non-ROTS and ROTS applications, although in the latter case additional measurements are required on the outgoing trunks to completely diagnose a problem detected by LTU measurements.

A subgroup's mean blocking is estimated from the average of the last-trunk-usage measurements in the level busy-hour; an estimate of at least B.02L using only five days of data indicates that the subgroup has an average blocking greater than B.01L with 84-percent confidence. In addition, the cost of installing the measurement equipment and test leads is held to a minimum since only one lead is required per subgroup. Hence, the method has the combined advantage of rapid detection of subgroups with service-affecting problems and ease of implementation.

VIII. ACKNOWLEDGMENTS

The authors are indebted to S. R. Neal for many helpful discussions and suggestions throughout the study. We also thank J. S. Hall for programming support and D. W. Hill for consultation.

REFERENCES

1. M. M. Buchner, Jr., and S. R. Neal, "Inherent Load-Balancing in Step-by-Step Switching Systems," B.S.T.J., *50*, No. 1 (January 1971), pp. 135–165.

2. R. I. Wilkinson, "A Study of Load and Service Variations in Toll Alternate Route Systems," Proc. of the Second International Teletraffic Congress, The Hague, July 7–11, 1958, Document No. 29.
3. S. R. Neal, private communication.
4. S. R. Neal, and A. Kuczura, "A Theory of Traffic-Measurement Errors for Loss Systems with Renewal Input," B.S.T.J., *52*, No. 6 (July–August 1973), pp. 967–990.
5. D. W. Hill, and S. R. Neal, "Traffic Capacity of a Probability-Engineered Trunk Group," B.S.T.J., *55*, No. 7 (September 1976), pp. 831–842.
6. D. W. Hill, private communication.