THE BELL SYSTEM TECHNICAL JOURNAL
Vol. 62, No. 7, September 1983
Printed in U.S.A.

Total Network Data System:

Trunking Systems

By P. V. BEZDEK* and J. P. COLLINS*

(Manuscript received August 3, 1982)

The Total Network Data System/Trunking (TNDS/TK) systems are those modules of TNDS that support the engineering and administration of the message trunk network. TNDS/TK consists of the Trunk Forecasting System (TFS), the Trunk Servicing System (TSS), and Common Update/Trunking (CU/TK). Using representative trunk group loads and switching system growth data as input, TFS forecasts trunk needs for five future years. TSS fine tunes the first year of the forecast by showing where to rearrange the network to meet current demand. And CU/TK supports TSS and TFS with a record base that contains a description of the network and user-stated parameters. We describe TNDS/TK from the standpoint of its environment, functions, system internals, and future direction. We also present a high-level view of its algorithms. For more detail, the reader is referred to the "Theoretical and Engineering Foundations" article and its references in this issue.

I. THE CIRCUIT ADMINISTRATION CENTER

1.1 General

The Total Network Data System/Trunking (TNDS/TK) systems support the activities of people in the Circuit Administration Centers (CACs). These are work centers in the Bell Operating Companies (BOCs) that are responsible for engineering and administering the

* Bell Laboratories.

message trunk network. Each BOC has at least one CAC; large companies may have several CACs with responsibilities divided geographically.

The CAC functions consist of (1) forecasting the required size and placement of trunk groups over five future years (trunk forecasting), (2) determining how best to modify the existing network to satisfy current demand (trunk servicing), and (3) initiating and monitoring the work orders that result in rerouting traffic from one trunk group onto another (route administration). TNDS/TK directly supports trunk forecasting and trunk servicing. However, it does not directly support route administration, although it does provide a major input to routing by defining, in forecasting, where reroutes are planned.

For a description of trunk servicing and trunk forecasting that is more comprehensive than that given below, the reader is referred to the "Environment and Objectives" article in this issue.

1.2 Trunk servicing functions

By sending Traffic Measurement Requests (TMRs) to TNDS/EQ, the trunk servicer indicates needs for the collection of traffic data on each trunk group throughout the year and frequently during the expected busy season. Traffic measurements that do not satisfy validation tests or represent normal customer behavior must be excluded from use. The servicer must initiate action to correct problems in the data collection process so that subsequent weeks' data can be collected successfully. The accepted measurements are used to estimate each trunk group's offered loads, current service levels, and the number of trunks required to provide objective service.

The servicer issues and tracks Message Trunk Orders (MTOs) to add or disconnect trunks to maintain objective levels of service and utilization. Representative busy-season loads, called "base" loads, must also be selected for each trunk group to support the forecasting process. Annually, Trunk Administration Measurement Plan (TAMP) reports are produced, which describe the adequacy of trunk-group data collection and compare the actual trunk network with a theoretical, low-cost network that produces the desired level of service and utilization.

1.3 Trunk forecasting functions

Trunk forecasting determines the future size and location of trunk groups as a major input to the construction program. Although the generation of the forecast is mechanized by TNDS/TK, the forecaster must manually determine and monitor much of its input. The activities this involves are described below.

The forecaster must define central office growth, typically in terms of main stations and traffic volume per main station. This person must specify the presence or absence of trunk groups, based on switch plans, homing arrangements, switch capacities, costs of facilities, tariff changes, and marketing demand forecast changes. Also based on these sources, the forecaster must indicate where the major shifts in traffic load will occur and select the appropriate load projection and sizing algorithms. The forecaster must perform these functions for scheduled forecasts and, often on short notice, for unscheduled ones in support of new switch plans or other construction program constraints. Finally, the forecaster must make input corrections that will reconcile future forecasts with current load variations that servicing reacts to with unplanned MTOs.

II. A FUNCTIONAL DESCRIPTION OF TNDS/TK

2.1 Overview

Figure 1 shows that TNDS/TK is a batch system made up of three component systems. The Trunk Servicing System (TSS) supports the servicing functions described above. Except for that part of the system related to annual TAMP, TSS is usually run once a week. The Trunk Forecasting System (TFS) performs most of the calculations involved in producing a General Trunk Forecast. It consists of several modules that may operate independently, though all are run between two and four or more times per year. Supporting TSS and TFS is Common Update/Trunking (CU/TK), which maintains a record base that stores a network description and associated parameters.

Separate installations of TNDS/TK are present in each BOC. Depending on the size of the company, an installation may process data for as few as 4000 to more than 100,000 trunk groups.

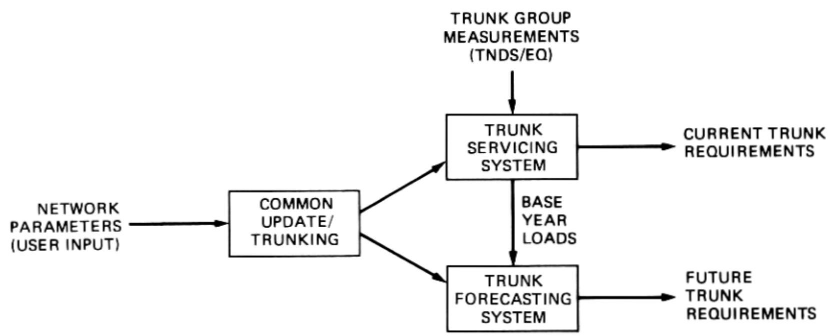

Fig. 1—TNDS/TK components.

2.2 *Common Update/Trunking (CU/TK)*

2.2.1 *General*

CU/TK is the interface between the user (the servicer or forecaster) and the rest of the system. Its record base stores a description of the network, engineering and administrative parameter values, some intermediate calculations, revisions to calculations, and some traffic loads. CU/TK provides 27 reports that describe its content. It also validates the input and advises the user of actual and potential errors.

2.2.2 *Record base content*

CU/TK data fall in two broad categories: those that apply to the company as a whole and those that apply to a portion of it, possibly to a single switch or trunk group. Normally, the first type is created and maintained by a person with company-level responsibility. The second is the obligation of the individual servicers and forecasters.

The company-level information is global in nature and serves several purposes. It may specify company policy or provide a key processing control. A single company input may also obviate numerous identical inputs by individual users. Among other things, company-level inputs define Bell System Common Language name change associations, report formatting and distribution criteria, engineering options, and the forecast years for TFS.

The servicer/forecaster-level information is more detailed and network dependent. *Traffic Unit Records* define the nodes in the network. They give the start and end dates of switching systems or NXX's that terminate trunk groups or for which growth data are to be stored. These records also allow forecasters to override the company-level specification of minimum trunk group size, described later. *Circuit Group Records* define the links in the network. They specify the life spans of trunk groups, load projection and sizing algorithms, trunk servicing criteria and defaults, base year loads, and alternate routes. Figure 2 shows a sample Circuit Group Record (TU500). *Growth Records* contain the main station and traffic forecast data that apply to central offices and are used in projecting trunk group loads. *Traffic Transfer Records* contain a description of the network impacted by area transfers, rehomes and reroutes, and the loads involved. In addition, there are inputs to revise intermediate TFS calculations, request nonautomatic reports to analyze the forecast, and specify that TSS ignore certain weeks of nonrepresentative traffic data.

2.2.3 *Input validation*

CU/TK verifies that the inputs contain the proper character set (field validation) and that all data on a record are consistent (intrarecord analysis). Consistency *between* records is verified downstream

```
TOTAL NETWORK DATA SYSTEM                                 PROCESS DATE 07 23 81
TRUNKING SYSTEMS                       TNDS               BISP REPORT     TU500
ACCEPTANCE TEST NETWORK                                   RESP CODE        UNF4
SERVICING                           CIRCUIT GROUP
                                       RECORD
                                                              12-31 74      ON
AN000227
            TROY NY 03 MGO   ALBY NY SS CGO   M  PH 55 IE                  I

PROJ CODE       13    PNT CGSN                    ALT RTE AL1    AN000632 100
GRTH CD A       00      PCT 2 WAY                         AL2    AN000178 100
        2       00    REQD MONTH    01                    AL3
CONV CODE        7    DSTN-CTGY                      COMP TNDM
TRMT CODE             MDLR ENG CD                  A CL FINAL    AN000632
CONF CODE             MIN PRI CCS                 ALT RTE ZL1
PROB BLKG             STIM-FACTOR                         ZL2
NBR SUB GPS           CNST PGM CGY                        ZL3
NBR ACSG EQ                    USE                   COMP TNDM
D/D VAR                  OWSHP A Z                Z CL FINAL
W-FACTOR              FCST RMKS CD                T AREA A
ECCS            09    PSEUDO A RTE                T AREA Z
CCS/TRK                      PCT A                OTHER CGID    00237
TS CONV CD            PSEUDO Z RTE                TF-RESP CD    UNF4
TF INFO MSGK1 N   16.6 I    .0 O    .0

SVC OPT 1             DSGN-PK FCTR                REPORT TIME   E
        3             CKTS  IN SVC  0015          RPTS EXC-IND
PRI-HR          10            PEND                HIST CGSN
SEC HR          11       TRND DOWN                SVC RMKS CD
AVG HT         170    TDAS UPDATE                 TS RESP CD    UNF4
TS INFO

* * * * * * * * * * * * * * * * BASE VALUES * * * * * * * * * * * * * * *

BASING PRD     BASE   CNTL   BASE   TOTAL      OFL    PRIMARY      BASE
END DATE       TRKS   BASE   NAME   OFRD CCS   CCS    OFRD CCS   MO  EXT MO

 09 30 80        15    AL2    HBH      267                        3
                              AL1      267                        3
                              AL2      320       8               11

* * * * * * * * * * * * * * * * * CHANGES * * * * * * * * * * * * * * * *

 S119    ALT RTE AL1    AN121781                  PCT OFL AL1    100
         ALT RTE AL2    AN121484                  PCT OFL AL2    100
         A CL FINAL     AN121781
```

Fig. 2—Circuit Group Record (TU500).

in TSS and TFS. CU/TK rejects inputs that fail its analysis and notifies the user with a Record Update Errors Report (TU001) and several statistical summaries.

2.3 The Trunk Servicing System (TSS)

The following sections outline the major functions of TSS. Representative processes within these functions are described in some detail, illustrating the types of processing that TSS performs.

2.3.1 Cycle setup

Although a TSS cycle can only process traffic measurements taken during a single week, the data received from the Traffic Data Administration System (TDAS) can pertain to several weeks. (The different technologies employed in data collection, ranging from photographic to electronic, produce varied delays before the measurements are made available to TSS.) The TSS cycle-setup function therefore separates the measurements that pertain to a user input "study week" from those that pertain to other weeks. The study-week measurements are passed to the measurement validation function; the others are held for later runs of the system. Measurements that pertain to those weeks before the study week are reported to the users so that back-dated cycles can be run as appropriate.

The CU/TK circuit group records contain (at one time) a time-varying description of the trunk network. The cycle setup function also extracts a static description of the trunk network, appropriate for the study week to be processed.

2.3.2 Measurement validation

The systems that collect traffic measurements and deliver them to TSS do not eliminate machine or human error. TSS must attempt to keep erroneous measurements away from its load-estimation processing. Although certain measurement errors generate absurd data, other errors can create apparently normal data or unlikely but theoretically possible data. The TSS validation function screens out measurements statistically likely to be in error. Rejected and missing measurements are reported to the users, so that problems in the measurement collection systems can be identified and corrected.

A few sample validations are described below. These validation tests are performed for each hour in which the appropriate measurements are available. The measurements discussed here are defined in the "Theoretical and Engineering Foundations" article earlier in this issue.

The peg count (*PC*) and overflow (*OFL*) measurements taken at one end of a trunk group are compared. If

$$OFL \geq PC > 0,$$

then both measurements are rejected. It is not valid for *OFL* to exceed *PC*. Although they can theoretically be equal and positive, most cases of such measurements are caused by measurement error.

When usage (U) is measured at both ends of a trunk group, the measurements need not agree, because of the sampling technique used. Based on a few assumptions about holding times and random calling patterns, each measurement should be an approximation to the true usage on the group, and hence to the other measurement. Specifically, if

$$(U_1 - U_2)^2 > (200 \text{ seconds})(U_1 + U_2),$$

TSS determines that the two measurements are not sufficiently close for both to be acceptable. It rejects the lesser measurement because typical problems in collecting usage data produce undervalued measurements. If, however, the two measurements are sufficiently close, TSS uses their average as its estimate of true usage.

2.3.3 *Load estimation*

TSS can estimate the average load offered to a trunk group from any of five measurements, or from two particular combinations of these measurements. The specific calculation used for a trunk group depends on trunk group type and the availability of measurements after the validation function. In most cases, several parameters that describe the traffic on a group and the group's performance are calculated, rather than average offered load alone. These parameters are determined from up to five hours' data, representing a fixed clock hour for five days of a business week. Thus, up to 24 sets of estimates, one per measured clock hour, are produced for the business days of the study week.

If, for example, usage, peg count, and overflow measurements are available from the same hour for the appropriate hours, TSS estimates the offered load (a) for each such hour as

$$a = \frac{PC - OFL \times R}{PC - OFL} U,$$

where R is based on trunk group type and accounts for customer retrial of blocked calls. This equation is based on the assumption that calls blocked by the trunk group and not retried would have the same average holding time as completed calls. These estimates of offered load are averaged over five days to produce study week hourly average offered loads. The variance among the daily offered loads in each clock hour is computed and stored, as well as estimates of blocking, average holding time, and peakedness.

If only usage measurements are available for a trunk group because of equipment limitations or lost or rejected data, TSS executes a more time-consuming algorithm that produces fewer, less reliable results. First it averages the usage measurements in each clock hour over the measured days. Then it uses standard numerical-analysis convergence techniques around a program that calculates expected usage, given offered load, to find an offered load that corresponds to the average measured usage. Blocking is computed as a by-product of this process. But average holding time and peakedness cannot be computed in this case. In fact, a user estimate of peakedness (or, if necessary, a TSS default value) is needed as an input to the process.

2.3.4 *Study period formation*

To increase the reliability of its traffic estimates, TSS must produce averaged parameter values that cover study periods of up to four weeks, again for each measured clock hour. The values from the current study week are averaged with values from preceding study weeks, weighted by the number of measured days for each hour in each week. As a result, the measured days are effectively averaged with all equal weights. The function outputs up to 24 sets of averaged estimates, one per measured clock hour.

Many groups are not measured every week. If a group is measured during the study week, TSS normally uses prior weeks' data (at most eight weeks older) to form averages of four measured weeks. Servicers can optionally submit Administrative Period Control inputs to CU/TK, inhibiting TSS from using past data, in cases where the traffic offered to the group is changing significantly (caused, for example, by seasonal variations in demand). If a group is not measured during the given study week, but during one or more of the preceding three weeks, TSS estimates "current" study period loads equal to the previous study period loads, for use in selecting busy hours.

2.3.5 Busy hour selection

From these hourly study-period load estimates, TSS must select a particular hour's load for which to size each trunk group. The network would obviously satisfy service objectives in all hours if each group were sized to satisfy its own service objective in its own busy hour (where "busy" is suitably defined). But because traffic can overflow from one trunk group to another, and unused capacity in one group can relieve a nearby overloaded group, a network engineered in this way would be more costly than necessary. For economic reasons, the hour for which a trunk group should be sized, the administrative hour, depends on loads offered to that group and the surrounding network.

The TSS process for selecting busy hours, Significant Hour Engineering, involves clusters of trunk groups as well as individual groups. A cluster is a final trunk group, together with all the surrounding high-usage groups that (1) share one (fixed) endpoint with the final trunk group, and (2) overflow traffic directly or indirectly to the final trunk group based on the alternate-routing logic of the switch at the shared endpoint.

In cases where all the groups in a nontrivial cluster (i.e., more than one trunk group) have study period load estimates available, TSS computes a cluster load for each hour by summing the carried load on the high-usage groups and the offered load on the final group, in corresponding hours. The hour with the greatest average cluster load per measured trunk is identified as the cluster's busy hour and the "control hour" of the final group. And, the hour with greatest measured blocking on the final group is termed the final group's "service busy hour." These two selected hours may coincide.

Selecting busy hours for high-usage groups is, from a logical view, recursive. Each high-usage group has a set of "significant" related groups defined in the CU/TK circuit group records, consisting of (1) the groups in its alternate route(s) and (2) its cluster final(s) (the plurals here are applicable to two-way high-usage groups only). The set of control hours of the significant groups is computed and then identified as the set of significant hours for the high-usage group itself. Then the significant hour with the greatest load on the high-usage group becomes the group's control hour.

For example, in Fig. 3, there are three clusters: groups AB and AD, BC and BD, and CD. Based on greatest average cluster load per measured trunk, the control hours for groups AB, BC, and CD might be hours 7, 8, and 12, respectively. The significant hours for group BD are now known to be 8 and 12. If the load on BD in hour 8 exceeds that in hour 12, then the control hour for BD is hour 8. The significant hours for group AD are now known to be 7 and 8. If the load on AD in hour eight exceeds that in hour seven, then the control hour for AD is hour eight. An even greater load on AD in hour 9 or 12 would not be considered.

2.3.6 Trunk group sizing

TSS computes the number of trunks required to provide objective service on final groups, based on average offered load, peakedness, and day-to-day variation. Required-trunks values are calculated for the traffic in the group's control hour and in the group's service busy hour. TSS chooses the larger required-trunks value as the value for the current study period, and the associated hour is called the group's administrative hour.

High-usage groups are sized for economic, rather than service-based, criteria. Using control-hour traffic data, TSS determines the number of high-usage trunks required to minimize the combined cost of the direct and alternate routes.

The Circuit Group Servicing Record report shown in Fig. 4 is produced after the trunk group sizing process. The bottom part of the report page shows a rough graph of trunks in service and study-period trunks required against time, for the past 65 weeks. The top left gives

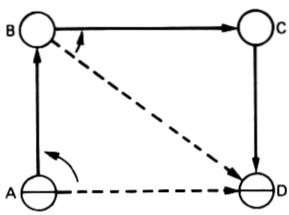

Fig. 3—Sample network for busy hour selection. Solid lines indicate final trunk groups, and dashed lines high-usage trunk groups. Curved arrows show the alternate routing used for calls overflowing a high-usage group.

```
TOTAL NETWORK DATA SYSTEM        TNDS - TRUNK SERVICING SYSTEM                  PROCESS DATE 01 11 81
TRUNKING SYSTEMS                                                                BISP REPORT    TB668
ACCEPTANCE TEST NETWORK          NETWORK CIRCUIT GROUP SERVICING RECORD           RESP CODE    UMF4
SERVICING
CYCLE DATE  12 08 80
                                                                               CONV TS-  CONF REPORT DESIGN
                                                     PRI SEC  SVC-OPT ECCS  HT  CODE CONV CODE TIME     PF
AN000227   00237    TROYNY03MGO ALBYNYSSCGO M- PM 55 IE  10  11  0 3      09  170  7              EDT    1.00

 LAST STDY-WK           CIRCUITS               SEL     CIRCUIT GROUP ADMINISTRATIVE HOUR      *  SIGNIFICANT HOURS
  OF STDY-PRD  I-SVC REQD FCST TND-DN PND   MEAS HOUR OFRD-CCS OFL-CCS OFRD-PF VAL-DAYS  *   LEG HR   CGSM
 ------------  ----- ---- ---- ------ ---   ---- ---- -------- ------- ------- --------  *  ---- --  --------
 80 OCT 06       15   10   16               UPCO  11    251        0     1.01     18     *   AL1 19 AN000632
    OCT 13       15   10   16               UPCO  11    263        1     1.02     17     *   AL2 11 AN000178
    OCT 20       15   10   16               UPCO  11    260        1     1.02     17     *   AL3
    OCT 27       15   10   16               UPCO  11    252        1     1.06     18     *   ZL1
    NOV 03       15   10   16               UPCO  11    250        1     1.05     17     *   ZL2
    NOV 10       15   10   16               UPCO  11    242        0     1.04     17     *   ZL3
    NOV 17       15   10   16               UPCO  11    246        0     1.04     18     * A-CLST 19 AN000632
    NOV 24       15   10   16               UPCO  15    239        0     1.00     17     * Z-CLST
    DEC 01       15   10   16               UPCO  15    245        1     1.10     18     *   CBH 14
    DEC 08       15   11   16               UPCO  11    265        0     1.00     19     *   SBH
 * = STUDY WEEK DELETE APPLIED.---- # = CKTS REQD = I-SVC - 1.----ADM PRD CNTL =        ----+ = EXCEEDS SUBGROUP CCS LIMIT.--
--------------------------------------------------------------------------------------------------------------

                              79             80
 LAST                             19NOV     28JAN     07APR     16JUN     25AUG     03NOV
 STDY-WK                     15OCT  .  24DEC  .  03MAR  .  12MAY  .  21JUL  .  29SEP  .  08DEC          REMARKS
 (MONDAY)         CIRCUITS       .    1    .    2    .    3    .    4    .    5    .    6    .
 OF STDY-PRD                 12345678901234567890123456789012345678901234567890123456789012345
                          25
                          24
                          23
                          22
                          21
                          20
                          19
                          18
                          17
                          16
                          15 SSSSSSSSSS SSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSSS
                          14
                          13 RRRRRR  RR RRR
                          12       RR      RR                        RR
                          11                 RRRRRRRRRRRR R   R RRRRR  RRRRRRRRRRR           R
                          10                             R RRR R                   RRRRRRRRRR
                           9                                                      R
                           8
                           7
                           6
                           5
  NOTE:                    4
  S = CKTS IN SVC          3                                                                     FORECASTING
  R = CKTS RQRD            2                                                                     TROYNY03MGO
  E = CKTS SVC=RQRD        1                                                                     PAGE      2
```

Fig. 4—Circuit Group Servicing Record.

more detail on administrative hour load conditions for the past ten study periods. The top right shows the significant hours identified for the current study period.

2.3.7 Trunk group banding

Based on observed blocking on final groups, and comparison of trunks in service with trunks required on high-usage groups, TSS assigns a "service band" to each measured group for each study period. The five bands identify groups that are overprovided (underloaded), possibly overprovided, close to objective, possibly underprovided (overloaded), and underprovided. The thresholds that separate these bands are adjusted for the statistical reliability of the data available for a trunk group. For example, small trunk groups with a few days' data available would be expected to have wide variations in their TSS-calculated load parameters. For such groups, the limits of the "close to objective" band are wider than for other groups.

In each study period groups judged to be overloaded are reported to the servicers. However, groups may appear underloaded simply because they have (currently) spare capacity needed for an upcoming busy season. Therefore, TSS only displays underloaded groups for corrective action if such groups have been underloaded for their entire banded history (up to 65 weeks).

The reports of overloaded or underloaded trunk groups from this function advise the users of significant service or utilization problems. Servicers must then determine what actions are appropriate, if any, or if only a transient condition in loads or measurements has occurred. The Circuit Group Servicing Record, as well as reports of weekly average load data, daily measurements, and measurement validation results, are available to support the servicers in this task.

2.3.8 Base selection

Just as TSS must select an appropriate current load from all the hourly loads in a study period to do trunk group sizing, so must TFS have an appropriate future load from all the hourly loads in a future forecast year to forecast future trunk requirements. The TSS base selection function executes a process similar to busy hour selection, starting with the hourly loads in a user-specified span of study periods. It outputs the most significant trunk group loads and hours in the specified base period, so that TFS can later estimate the loads in corresponding hours of future years.

Servicers can input Extended History Delete transactions to CU/TK to exclude a specified study period's data for a specified trunk group from the base selection function. This feature allows the user to prevent data that represents unusual traffic, or errors in measurement, from affecting the trunk forecast. The user can also replace any program-generated base load with a human-generated load before TFS is run, after reviewing the Base Selection Details Report.

2.3.9 Network optimization

After base selection has determined the significant trunk group loads from a twelve-month historical period, the TSS network optimization function determines a low-cost trunk network that should have been in place to handle these loads. Unlike the trunk group sizing function, which calculates the required size of each group for its actual, measured load, the optimization function estimates the theoretical load on overflow-receiving groups, to adjust for proposed trunk changes in the surrounding network, and sizes these groups accordingly. Optimization can therefore suggest widespread, related changes in trunk-group sizes but it does not suggest adding or removing trunk groups.

2.3.10 Trunk Administrative Measurement Plan (TAMP) reporting

In support of TAMP, the AT&T plan to index the performance of the trunk administration process, TSS produces reports annually from base selection and network optimization outputs. Each final group is assigned to one of the five bands previously described, according to its blocking in its annual busy hour. In addition, every group is assigned to a band, based on a comparison of trunks in service (in the annual busy hour) and optimized required trunks. The distribution of trunk groups among these bands is reported, by administrative responsibility and by trunk group type (end office, operator connecting, tandem, tandem connecting, and auxiliary services). Also, the TAMP reports list and summarize the number of days' data available to compute each group's busy hour load.

2.4 The Trunk Forecasting System (TFS)

The principal output of TFS is the General Trunk Forecast, a document that is a major input to the BOC construction program. A company-created forecast period table controls the span and complexity of the forecast. Through the table, the company defines the five forecast years, the first four of which must be consecutive and are normally chosen to be in the immediate future. For planning purposes, the company may choose a more distant fifth year, perhaps 20 years ahead.

The company also uses the table to define "base year subdivisions" and "forecast periods." These are partitions of the base (current) year and forecast years that tell TFS how much detail to create. Up to 20 forecast periods (normally quarters for each forecast year) and four base year subdivisions may be specified. TFS assumes conditions that exist on the last day of a forecast period exist throughout the period, so the importance of defining more than one forecast period per year is clear. Single annual periods would result in an improperly sized network if major events such as area transfers did not coincide with busy seasons, as frequently occurs. Although a company may opt for numerous forecast periods, it does so at the expense of computer run time which is proportional to the number of periods chosen.

The sections below describe the major functions of TFS, including user interactions, reports, and a high-level view of the calculations. Following these is a description of the sequencing of operations considering both TFS and CU/TK.

2.4.1 Growth factor computation

To forecast traffic load on trunk groups, TFS requires knowledge of the growth in load of the switches that the trunk groups terminate on. For each such switch, the forecaster enters growth data through CU/TK. The data are usually in the form of main stations (MS) and hundred call-seconds per main station (CCS/MS), as of a given date. The data may represent all customers served by the switch, or a subset of the customers according to NXX or class of service, e.g. coin, residence, etc. The data are entered normally by year to cover the base year and all forecast years.

TFS uses these data to compute growth factors in the form of ratios of future load to present load. First, TFS assures that data are available for each base year subdivision and forecast period end date; when the data have not been expressed to coincide precisely with the needed dates, TFS derives what it needs by linearly interpolating from surrounding data. Next, the data are aggregated to the appropriate level for forecasting trunk group loads. For example, if the forecaster input data by NXX but required growth factors for only the total switch (comprising several NXX's), then the NXX data would be summed for the entire switch.

Products of the like-dated individual growth data terms are then formed, e.g., MS and CCS/MS are multiplied to form CCS. Finally, by simple division of a future product by a base year product, growth factors are produced. For each switch, one growth factor is computed for each forecast period relative to each base period (as defined by the base year subdivisions). The forecaster has the option of bypassing this process by manually stating growth factors, if computed ones are inappropriate.

With calculations complete, the system outputs a Growth Factor Report. At this point, errors in the input data are most evident. To avoid the need to rerun the process with new inputs, the system allows the forecaster to substitute new growth factors for the incorrect ones through CU/TK.

2.4.2 Network disassembly

TSS base selection provides TFS with trunk group loads that represent busy hour and busy season conditions during the base year. TFS must convert these loads to a form appropriate for the projection function. For the most part, this amounts to removing overflow traffic from the loads of groups that are alternate routes for other (high-usage) trunk groups. This "disassembly" is necessary, since the overflow is a function of a previous network structure and must not be projected with the (network-independent) first-route traffic.

For each high-usage group, TSS base selection identifies the load offered to the group and its overflow. Using the alternate route information on Circuit Group Records, TFS associates the overflow with the appropriate alternate route groups. (When the overflow is fragmented over multiple alternate routes, TFS allocates a portion of the overflow to each route according to user-specified percentages.) TFS then subtracts the overflow values from the offered loads of the receiving group.

Because of differences in data collection schedules and measurement errors, a group's received overflow may appear to exceed its offered load. In such a case, TFS will assume a 0 value for the first route load. If the forecaster anticipates this and concludes that a 0 value is an understatement, the forecaster may input a "minimum primary CCS" that will override any lower computed value.

2.4.3 Projection

With first route loads and growth factors as input, TFS estimates future load. The first step is to compute projection ratios. These are created by substituting the growth factors into a formula that the forecaster has selected for each group. For example, if the AT&T recommended formula $(A + Z)/2$ were selected, the projection ratio would reflect the average growth of the originating end (A-end) and terminating end (Z-end) of the group, as determined by A's and Z's growth factors.

Of the total set of growth factors available for a specific A and Z, the ones used in the formula are those developed for the base year subdivision and forecast period containing the measurement period (busy month) of the base load. So if a trunk group had a base load with a May busy month, and the base year and forecast years were calendar years with a quarterly base and quarterly forecast periods, then the growth factors used would be those that apply to the second quarter of each forecast year relative to the second quarter of the base year.

The multiplication of the base loads by the projection ratios completes projection. In a similar manner, the loads stated on traffic transfer records are projected.

The system allows several options that the forecaster can exercise in advance of a projection run. Instead of relying on growth factors, the forecaster may state manually derived projection ratios or specify an annual percent for compounding. The forecaster may supply a stimulation factor or select a projection formula that accounts for community-of-interest considerations or regional growth. For a specific A or Z, the forecaster may request the use of growth factors that apply to only a part of the switch. Using "pseudo" trunk groups, described later, the forecaster may project separately the individual traffic items on a group. Finally, if all the system-provided methods are inappropriate, the forecaster may state future loads that are externally derived.

2.4.4 Sizing

Disassembly, projection, and sizing are run as a single process; the forecaster has the opportunity to review all calculations only after required trunks are computed. The same blocking objectives and economic criteria are used in TFS sizing as in its TSS counterpart. In general, TFS computes one trunks-required value per group per forecast year based on the load in the forecast period that contains the group's busy month. For office(s) affected by major event(s) such as cutovers, however, the forecaster may specify that required trunks be computed as of the date of the event(s). This will result in more than one trunks-required computation per year (except where busy months and event dates fall in the same forecast period).

As a first (logical) step, the process adds together (1) projected first route loads, (2) projected traffic transfer CCS values (+ or −), and (3) user-stated load adjustments (+ or −), if any. Transfer loads are needed since projected trunk group loads alone may not compensate for the effects of reroutes and new or deleted groups that occur after the base year. Next, the process sizes those groups that receive no overflow (only-route and primary high-usage groups). Overflow from the primary high-usage groups is then computed.

The rest of the procedure is the inverse of disassembly. Overflows are added to the offered loads on the alternate routes. When a group receives the overflows from all expected sources, the process computes the peakedness of its load and its required trunks. If the group is high usage, overflow and variance are computed. The previous steps are then repeated. The result is that the network is sized iteratively, bottom-up through that portion of the five-level network hierarchy included in the company's database. The process will add to the computed size of a group a trunk adjustment (+ or −), if any is stated; for high-usage groups this is performed before overflow is computed.

To facilitate more economic purchasing of trunk equipment, TFS will convert the above values to modules of 12 or 24 trunks. The forecaster, however, must request this action on a group-by-group basis. Which of four modular sizing procedures is invoked depends on whether the group is high usage, one way or two way, or uses digital facilities or digital terminal equipment. Modular sizing takes place before computing overflow from high-usage groups.

A major function imbedded in sizing is the determination of where new trunk groups are warranted. Although TFS mechanizes the test for new groups, the forecaster must identify candidates in advance by creating what are called "pseudo" Circuit Group Records in CU/TK, complete with base loads. To test fully all possibilities, the forecaster should input large numbers, perhaps thousands, of pseudo groups. But the practical limitations of deriving base loads and specifying all the necessary parameters reduce the number created to a small subset of those possible.

TFS projects the pseudo trunk group base load as if the pseudo group were an actual group. The sizing process calculates required trunks for it and determines whether a minimum-trunk-group-size threshold is exceeded. If so, the pseudo group is retained as a planned group (one that will appear in some future year of the forecast). If not, the pseudo group's loads are distributed onto existing groups specified by the forecaster. This is necessary since the pseudo group's loads represent actual traffic that is not accounted for elsewhere and would otherwise be lost.

Apart from planning new groups, pseudo groups give the forecaster more flexibility in projecting traffic. The forecaster may remove the load of one or more traffic items from an existing group's base load, define pseudo groups for them, and project them by different methods. If the forecaster indicates that the pseudo groups were created solely for separate projection, sizing will suppress the minimum-size test and add the pseudo groups' projected traffic back to their existing route(s).

The principal user output from sizing is the Preliminary General Trunk Forecast (PGTF). In addition to trunks, it displays several of the major intermediate calculations: base and future offered CCS, projection ratios, and peakedness. For each group, the report indicates, among other things, whether it was a pseudo, is modularly engineered, received overflow, or includes transfer loads. At this point, the forecaster may revise only future offered CCS and required trunks, as none of the other items appear on the final TFS reports. Since the analysis may be complicated, the forecaster may request up to three additional reports. The Forecast Detail Report, Detail Traffic Transfer Summary, and Overflow Summary supplement the PGTF with the information their titles imply.

With revisions complete, TFS produces its final reports. The General Trunk Forecast (GTF) may be issued in several forms, combining the newly computed trunk values with absolute differences or percent differences from previous forecasts. Figure 5 shows a GTF. The Office Trunk Studies display load and trunk values in a way useful for central office equipment engineering, and the Construction Program Worksheets provide a starting point for preparing AT&T construction program reports.

TFS also produces a computer tape of forecast results that is input to another Bell Laboratories product, the Facility and Equipment Planning System (FEPS), and BOC-developed programs. FEPS is a module of the Trunks Integrated Record Keeping System (TIRKS).

2.4.5 Sequencing with CU/TK

Figure 6 shows the relationship between CU/TK and TFS in a complete forecast run. The diagram refers to two TFS functions that have not been described to this point. Growth Factor Analysis (GFA)

TOTAL NETWORK DATA SYSTEM
THUNKING SYSTEMS
ACCEPTANCE TEST NETWORK
SERVICING

TNDS - TRUNK FORECASTING

ORIGINATING GENERAL TRUNK FORECAST

FORECAST RUN DATE 03 31 81
PROCESS DATE 10 09 81
BISP REPORT TW530
RESP CODE EF30

GLLD PA GN CG0 1ES CLASS 5

TO	TRUNK TYPE	ORIG UNIT PROB ECCS	RQD MO	BASE TRKS	REQUIRED TRUNKS 81/82	82/83	83/84	84/85	85/86	RMKS CODE	SERIAL NUMBER	OTHER CGID
PHLA PA SL	42T M- 1B 5- DD	11	02				98	101	102	DC	AA074262	00200 GLLD
	42T M- DF 54 PA	010	02	7	7	7	11	11	11	DC	AA084304	00200 GLLD
	42T M- DF 54 TO CR	010	03		2	2	2	2	2	DC	AA092097	00200 GLLD
	42T M- DF 54 SP NC2	W010H	01				19	19	19	DC	AA092380	015C1 1B006D
	42T M- DF 54 SP CN3	W010L	01				60	61	61	DC	AA092382	20701 1B006D
PHLA PA TH	CG0 M- PH 55 IE	07	02	11	12	13				XW	AA055919	00200 00410
	CG0 M- PH 55 IE	12	02				21	21	21	XW	AA055919	00200 00410
PHLA PA WV	CG0 M- PH 55 IE	07	02	6						XW	AA055920	00200 00410
	CGC M- PH 55 IE	14	02		9	10	18	19	19	XW	AA055920	00200 00410
RDPK PA RP	D60 M- DF 55 IE	W010L	11						59	XW	AA094563	00200
	52C M- PH 55 IE	05	02	29	24	26				XW	AA055921	00200 B61M
	52C M- DF 55 IE	W010L	02				58	* 59	59 X	XW	AA055921	00200 B61M
	52F M- PH 55 IE	09	02				8	9	9 X	XW	AA055922	00200 B164M
SPFD PA SF	D60 M- DF 55 IE	W010L	09						87		AA094565	00200 SES
	MG0 M- PH 55 IE	08	02	28	30	32				06	AA055923	00200 B220M
	MG0 M- DF 55 IE	W010L	02				88	88	88 X	06	AA055923	00200 B220M
	MG1 M- PH 55 IE	08	02	17						03	AA055924	00200 B60M
	MG1 M- DF 55 IE	08	02		13	15				03	AA055924	00200 B60M
	MG1 M- DF 55 IE	W010L	02				41	42	42 X	03	AA055924	00200 B60M
WAYN PA LA	42T M- AF 51 DD	W010H	02	77	61	67	145	149	149	DC	AA062071	00200 GLLD
	42T M- 1F 51 DD	25	03	27	63	131	138	140	140	DC	AA067372	00200 GLLD
	42T M- DF 51 DD CR	010	03	2	2	3	4	4	4	DC	AA073149	00200 GLLD
	42T M- DF 51 SP CN3	010	01				2	2	2	DC	AA085604	20701 1B006D
	42T M- DF 51 SP NC2	010	01				3	3	3	DC	AA085805	015C1 1B006D
WAYN PA WY	CG0 M- PH 55 IE	14	02				4	4	4	XC	AA082120	00200 00406
WCHS PA WC	CG0 M- PH 51 TE	14	02				11	11	9	XC	AA086392	00200 00406
WLMG DE WL	41T M- DF 53 SP CN3	W010L	07	26	22	22	22 X			XC	AA055886	20701 D779SP
	41T M- PH 53 DD	11	02	6	7	8	8 X			XC	AA055929	00200 A75640
	41T M- DF 53 SP NC2	W010H	07	10	11	11	11 X			XC	AA057168	015C1 D779SP
	TOTAL INTRA BLDG TRKS			192	181	191						
	TOTAL INTER BLDG TRKS			611	648	756	1566	1569	1513			

\# = MANUALLY REVISED

GLLD PA GN CG0

Fig. 5—General Trunk Forecast.

and Circuit Group Analysis (CGA) extend the validations of CU/TK to include interrecord error checks. The philosophy behind this sequence, referring to the numbered steps in the diagram, is as follows:

1. Computer runs of CU/TK and GFA are needed to establish or correct the database for this forecast view. The forecaster receives error reports and, with this step and the ones below, is allowed time to submit corrections.
2. CU/TK is run to accept the corrections.
3. CU/TK is run to allow corrections to erroneous inputs in Step 2, GFA is run to freeze the database, and growth factors are computed.
4. CU/TK and CGA are run to accept growth factor revisions and identify any remaining interrecord errors.
5. Similar to Step 2.
6. Similar to Step 3, concluding with all remaining calculations.
7. CU/TK accepts requests for additional analysis reports, issued by TFS.
8. CU/TK accepts the forecast revisions. These are merged with the preliminary results and final outputs are produced.

Rarely, however, does a company execute the full sequence as described. To do so, allowing for forecaster interactions, would require two or three months. Typically, the company will omit Steps 2 and 5 and will merge GFA and CGA into a single step. A company may also

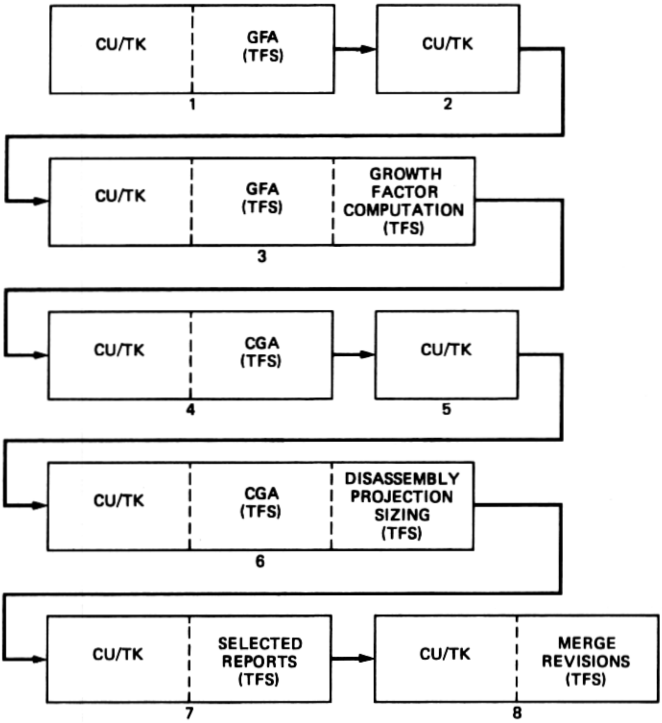

CU/TK – COMMON UPDATE/TRUNKING
CGA – CIRCUIT GROUP ANALYSIS
GFA – GROWTH FACTOR ANALYSIS
TFS – TRUNK FORECASTING SYSTEM

Fig. 6—Sequencing TFS with CU/TK. Arrows show the sequence of operations, not flow of data, and indicate pauses in computer run. During pauses, the forecaster reviews the intermediate output, prepares collections to the output and database, and inputs collections to CU/TK.

overlap the early steps of one forecast run with the concluding steps of the previous run. In some cases, too, only the forecast that pertains to a few switches is of interest. The company may then run TFS in the "Area Forecasting" mode, against only a selected subset of the database, and speed up the process accordingly.

III. SYSTEM INTERNALS

The TNDS trunking systems are designed to run in batch mode on an IBM 370-series computer system or equivalent, supported by the Bell System Standard Operating Environment software. The systems are coded primarily in COBOL, with some programs in FORTRAN or PL/I. Table I describes the size of these systems in several dimensions.

The source code and development documentation for these systems

Table I—Approximate system size

	Lines of Source Code	Compiled Source Parts	Delivered Programs	Data Files
CU/TK	34,000	25	14	15
TSS	77,000	100	58	170
TFS	85,000	110	68	135

are created and maintained using the *UNIX**/Programmer's Workbench (PWB) operating system and text-processing facilities on a minicomputer system. Multiple releases of a system, in simultaneous use at different installations, are maintained conveniently using the Source Code Control System. The Change Management Tracking System monitors the status of Modification Requests through various phases of investigation and resolution.

The programs are converted to executable format, tested, and transmitted to the users' computation centers by use of a telecommunications software system for computer-to-computer data exchange with operating telephone companies (TTRAN).

IV. LOOKING TO THE FUTURE

4.1 With TSS

With a planned trunk demand servicing policy feature, TSS will more effectively separate service problems that require corrective action from transient conditions associated with normal fluctuations in load. This feature will also select the best groups for corrective action within overloaded clusters, subject to user-input facility constraints. A second feature, providing a short-term forecast of loads, will then allow TSS to recommend cost-effective solutions to service problems before they occur.

An extensive revision of the TSS load estimation and study period formation functions will provide more accurate load data and trunk requirements through (1) calculation of loads from individual rather than averaged measurements, (2) inclusion of more days' data in weekend study-period averages, and (3) greater use of calculated parameter values (such as holding time) in preference to user-estimated or theoretical values.

TSS reports will be restructured so that a lesser amount of data is output automatically to the servicers. Servicers will be able to receive more detailed supporting data on request. Currently, TSS is designed to produce all of its reports automatically, thereby providing servicers with an excess of data from which they must extract the information of interest.

* Trademark of Bell Laboratories.

4.2 With TFS

Over the next few years, significant attention will be devoted to improving the user interface for forecasting. There are major implications in this for both TFS and CU/TK. One thrust will be to replace a large portion of the batch operation with on-line capability. Candidates for on-line use include database update, validation and inquiry, report retrieval, revisions to calculations, and the calculations themselves.

Another major enhancement that would be particularly useful in an on-line environment is a "what if" capability. Currently, TFS produces a single forecast with a database that gives a single consistent (though evolving) picture of the network at any given time. A "what if" capability would allow the user to produce several forecasts, each for a different switch and homing configuration. The best forecast, by criteria to be defined, would result in a database update with the preferred configuration.

The companies will also be able to exchange network information in a more mechanized way than is now possible. The exchange will include CU/TK records, intermediate TFS calculations, and forecasted trunks. To do all three will require the companies to coordinate their forecast schedules. The design of this feature will be robust enough to apply to the current network partition or to other partitions required by Dynamic Non-Hierarchical Routing (DNHR) or legislation.

Other proposed enhancements include the Sequential Projection Algorithm (SPA) and a Trunk Implementation Plan (TIP). SPA will replace the current projection method with one based on Kalman filter prediction theory. TIP is a method to convert the demand forecast from TFS into an administered one, considering forecast uncertainty, expense and capital costs, trunks in service, and facility availability. DNHR is also under study as it applies to TFS.

AUTHORS

Paul V. Bezdek, B.S. (Mathematics), B.A. (Physics), 1969, Wayne State University; M.S. (Computer Science), 1976, Steven Institute of Technology; Bell Laboratories, 1972—. Mr. Bezdek is currently Supervisor of the COEES/DOPS Systems Engineering Group. Previously he was responsible for developing functional requirements for the Total Network Data System/Trunking Systems. Member, Phi Beta Kappa.

Julius P. Collins, B.S., M.S. (Mathematics), 1974, Polytechnic Institute of New York; Bell Laboratories, 1974—. Mr. Collins has worked on design, implementation, and testing of several operations support systems for engineering and administration of the public telephone network. He is currently working on an engineering system for corporate telecommunications networks. Member, Pi Mu Epsilon.