

Traffic Monitoring in a Network of PBXs

M. Sestak and A.R. Kaye

Dept. of Systems and Computer Engineering
Carleton University, Ottawa, Ont., K1S 5B6, Canada
(M. Sestak now with Mitel Corporation, Ottawa.)

Abstract

The paper deals with monitoring traffic flow in a network of PBXs of arbitrary size and configuration. We propose two alternative monitoring architectures which differ primarily in the mode of generating call detail records (CDRs). In Architecture A the CDRs are generated in a manner similar to that now employed for supporting billing functions. In Architecture B a single CDR is generated by making full use of common-channel signalling. We show that Architecture B results in greatly reduced requirements for storage space and processing at the monitoring centre and in the transmission bandwidth required to deliver CDRs to the centre. We present algorithms which can be used with either architecture to provide both real-time and retrospective traffic monitoring.

1 Introduction

In recent years Private Branch Exchange (PBX) networks have grown greatly in size and in the complexity of features offered and they now make use of digital trunks and common-channel signalling and support both voice and data requirements. The need to provide for and to simplify the monitoring and management of PBX networks has increased along with this growth as it has in the field of computer-communication networks. However, facilities for monitoring have not kept pace with the needs and each PBX in a network is still typically managed as an isolated entity.

This paper is concerned with the monitoring of traffic flow in a network of PBXs at a central monitoring station and is based on [1]. We propose two alternative monitoring architectures for the generation, collection, storage and processing of monitoring data which would support a monitoring station with the capability to answer queries regarding:

- The call path associated with a specific device or collection of devices.

- Call paths for all the traffic in a network.
- Call paths for traffic originating from or terminating on selected PBXs.
- Call paths for all traffic affecting selected PBXs.
- Call paths of traffic occurring on links between a selected pair of PBXs.
- Call paths for the voice or data subsets of the above categories.

These capabilities will be supported both in real time and on a retrospective basis.

Raw data for monitoring purposes is generated in the familiar form of call detail records (CDRs). The two alternative architectures differ mainly in the approach to the generation of call detail records. The first approach, used in what we shall refer to as Architecture A, is a slight modification of the current approach in use in the PBX industry; the modifications are required to permit the real-time aspects of our monitoring system. The second, Architecture B, is a more radical departure from current practice which we propose in order to take advantage of the widespread use of CCS in modern PBX networks and to reduce the resource requirements associated with the monitoring functions.

The CDR architecture in common use today is primarily intended to serve the billing function and was designed at a time before CCS was in common use. We shall show that the current form of CDRs can be used, with our Architecture A, for the monitoring functions described above but that the resultant monitoring resource requirements are significantly greater than they need be. We define these resource requirements in terms of the disk space and processing capability required at the monitoring centre together with the bandwidth required to transmit the CDR data from the network to the monitoring centre. We then proceed to show that these resource requirements can be significantly reduced for a very wide range

402A.2.1

of PBX networks by the use of our Architecture B. Neither of our Architectures imply any impairment of the ability to carry out the normal billing functions for which CDRs are used.

Our approach is first of all to develop a traffic model which can be applied to any specific network topology. We then apply this model to two generic topologies which serve to establish bounds on the results that would be expected for actual networks. These are used to compare the computing resources required to support the network monitoring capabilities using the two different CDR architectures.

2 Traffic model

In order to develop quantitative comparisons of the resource requirements associated with monitoring we must first develop a model of the traffic carried in a PBX network. This involves three elements: call types, network topology and the volume and distribution of calls of different types.

2.1 Call types

A specific PBX may handle six different types of call which are of interest in this paper:

- Intra-PBX call - originating and terminating within the PBX.
- Incoming external call - originating outside the network and terminating on the PBX;
- Outgoing external call - originating on the PBX and terminating outside the network;
- Incoming network call - originating on the PBX and terminating on another;
- Outgoing network call - originating on the PBX and terminating on another;
- Tandem call - passing through the PBX en route between two others.

The CDRs which are generated will vary in content and, in the case of Architecture A only, in number according to the type of call being recorded.

2.2 Network topologies

The number of tandem calls as a percentage of total calls in the network, and the amount of CDR data generated, will depend on the topology of the network. The resource requirements for monitoring are directly related to the number of CDRs generated and their size. In our model for projecting the rate of generation and size of CDRs, we represent the effect of topology by the average number of inter-PBX hops traversed

by a call from source to destination.

2.3 Volume and distribution of calls

There appear to be no published data which would tell us the volume and distribution of calls by type in a network of PBXs. Consequently we have developed a model which projects such data as a basis for our subsequent work. We start with Table I which shows data used within the PBX industry to project heavy traffic call volumes for a single PBX based on its configuration [2].

In order to project traffic statistics for an N-PBX network we apply the following rules which modify the values shown in Table I for each PBX in the network:

- All PBXs on the network have identical configurations.
- Of the intra-PBX traffic shown in Table I, 70% remains internal while 30% becomes outgoing network traffic and terminates on other PBXs with a uniform distribution.
- The incoming external calls of Table I, arriving at one PBX, terminate on any PBX in the network with a uniform distribution. Thus the incoming external traffic which does not terminate on its PBX of entry becomes outgoing network traffic on that PBX.
- Outgoing external traffic is not routed to other PBXs before it exits the network. Thus it remains the same as shown in Table I for each PBX.

These rules are based on assumptions which we believe, on the basis of experience, to be reasonable but which, in any case, could be varied in order to test the generality of our results. We have carried out some sensitivity testing which shows that varying the principal assumption, the percentage of intra-PBX traffic which remains intra-PBX in a network, does not affect the validity of our comparative results.

The only further information we need in order to determine the total number of calls of all types which are handled by each PBX is the average number of inter-PBX hops traversed by a network call; this is based on the number of tandem calls occurring in the network. In any network this value will depend on the degree of connectivity in the network and the traffic pattern. In carrying out our comparisons, we use two topologies which, while they do not represent realistic situations in networks of more than a few PBXs, nevertheless represent the extremes of possible situations. These are the fully connected network, in which there are no tandem calls, and the linearly connected network in which tandem calls are maximized for a given number of PBXs.

As an example of the results of applying the rules for extending Table I to a network of PBXs, Table II shows the various types of traffic for a four-PBX network in which the PBXs have 1000 lines, 122 trunks and the average number of hops traversed by a network call is 1.5.

3 Monitoring architectures

By a monitoring architecture we refer to the basic architecture of the complete monitoring system. This includes information as to the time and place of generation of the CDRs for each type of call, the data content and mode of constructing a CDR and the transmission, processing and storage of monitoring data. In both our architectures, the CDRs are transmitted to a monitoring centre where they are processed and the resultant data is stored for accessing when requests for views of the network traffic are made by a network operator.

3.1 Architecture A

In PBXs today CDRs are generated at the end of a call for the purpose of producing billing reports. They contain information regarding both establishment and termination of the call. Each PBX involved in a call generates a CDR at call termination which records its own involvement in the call. The result is that a single network call traversing two intermediate PBXs, for instance, generates a total of four CDRs from the four PBXs that handle it. This produces a requirement for processing at the network billing centre to correlate the four CDRs which, of course may be received at different times interspersed with CDRs recording other calls.

Our proposed Architecture A follows the same approach as just described except that, in order to meet the requirement of real time monitoring, it is necessary to generate CDRs at the beginning of a call as well as at termination. Although this doubles the number of CDRs generated, it does not double the amount of data to be transmitted since most of the data in the call establishment CDR need not be sent again at call termination. The reason for our use of Architecture A as a basis of comparison is that it is probably the simplest modification that could be made to the current CDR architecture in order to add real-time monitoring capability. If only historical monitoring of traffic and congestion were required, that is not including current calls, the current CDR architecture could be used without change.

For Architecture A, a possible CDR format, the one we have used in our work, is shown in Figure 1.

3.2 Architecture B

Architecture B differs from Architecture A principally in the fact that only a single CDR is generated at call establishment and at call termination regardless of how many PBXs are involved. This concept takes advantage of the availability of common-channel signalling by causing the successive PBXs on the call path to generate call-path information progressively during the signalling process. Once the signalling is complete and the call is established, a complete CDR is transmitted from the terminating PBX to the monitoring centre. This single CDR contains all the information that would be included in the multiple CDRs of Architecture A. At call termination, a single CDR is generated by the terminating PBX. A possible format for the CDRs, again the one we have used, is shown in Figure 2.

4 Storage, processing and Transmission

This section describes the RAM and disk data structures used at the monitoring centre to record all CDR data and to support both real-time and retrospective queries. The same data structures are used for both Architectures A and B but there are differences in the amount of storage space required. To illustrate the structures we use the example of a network call shown in Figure 3, in which subscriber #1200 dials subscriber #7300 and is connected using trunk 2 on PBX1, trunks 62 and 91 on PBX2, trunks 18 and 311 on PBX3 and trunk 147 on PBX4.

On receipt at the monitoring centre, a CDR is processed and its data is used to update a matrix held in RAM as shown in Figure 4. This matrix is essentially a single-PBX call-map.

The call-map matrix is indexed by trunk number and PBX number and each element identifies the entity to which that trunk is connected at that PBX. In Figure 4, the expanded elements refer to three such connections in Figure 3. There is some apparent redundancy where a trunk is connected to another trunk but this allows for ease of answering queries which may refer to either trunk number. Each element requires 3 bytes of storage to allow for: up to six-digit identifiers; 1 bit to differentiate between data and voice calls; and 1 bit to indicate whether the entity referred to is a trunk or an attached device.

Each CDR is also recorded directly on disk, as it is received, as an element of a direct access file. Depending on the size of the network there may be several such active

files, each corresponding to an hour, day or week of activity. Each of these CDR files consists of a sequence of records, one per CDR, in the order in which they were received.

A "snapshot" of the RAM matrix is also stored on disk periodically together with and index to the last CDR stored on disk before the snapshot.

In order to provide rapid tracing of calls through the network it is necessary to identify the remote termination of each trunk on a PBX. This requires a trunk-connection matrix, again indexed by trunk and PBX, whose elements identify the identity of the remote end of the trunk at the remote PBX. The storage requirements for this matrix are not significant in comparison to other storage requirements since it is not dynamic.

The above data structures can be used in combination by query-answering algorithms to reconstruct the state of the network at any specified time, past or present and to provide all of the capabilities listed in the introduction.

5 Query processing

Real-time queries, regarding the current state of the network, can be answered using only the call-map matrix combined with the trunk-connection matrix. Retrospective queries require the reconstruction of the state of the network for the specified time. This is accomplished by retrieving the snapshot matrix stored on disk for the time previous and closest to the specified time. Recovered with the snapshot matrix is the index to the CDR file on disk for the last CDR whose data was included in the snapshot matrix. By combining the snapshot matrix with the CDRs which were recorded on disk after the snapshot and up to the specified time a replica of the RAM call-map matrix at the specified time can be reconstructed and the queries can be answered using it.

6 Comparative costs

We have compared the resource requirements at the monitoring centre for the two alternative architectures in terms of three types of resource: disk storage requirements, transmission bandwidth between the network and the monitoring centre and processing requirements. A typical example of our results is shown in Figure 5 which plots the disk space required per hour against the network size for a particular PBX configuration. The worst and best cases plotted refer to the extremes of possible network topology: respectively a linearly connected and a fully connected network. More realistic network topologies would show results falling between these two extremes.

Over the wide range of network and PBX sizes we have considered, Architecture B requires from about 25% of the storage space of Architecture A in the linearly-connected network to about 81% in the fully-connected network.

Other resource requirements vary in a similar manner to disk space. Transmission bandwidth is a function the volume of CDR data and the requirement consequently varies in the same way as storage requirements. Data-acquisition processing resources are a function of the number of CDRs generated per hour and the volume of data contained in the CDRs. A comparison of CDR generation rates shows that Architecture B generation rates vary from 2% to 75% of Architecture A rates over a wide range of networks.

7 Conclusions

The research presented in this paper has two main results:

- We have proposed two different monitoring architectures, involving different methods of generating call detail records (CDRs), which can be used to construct dynamic data structures which are capable of supporting a wide range of traffic and congestion monitoring capabilities in real-time or in retrospect.
- We have shown that our second proposal, Architecture B, requires significantly less resources to support the monitoring centre in terms of disk storage space, processing power and transmission bandwidth between the network and the monitoring centre.

Architecture A is very similar to the way CDRs are currently produced in typical PBXs; it requires only the production of additional CDRs to support the real-time monitoring capabilities. Architecture B differs in its integration of CDR production with signalling in common-channel signalling networks which are now becoming common in modern PBX networks.

In large networks of PBXs these reductions may be quite significant in terms of cost as is clear from a glance at the absolute volumes of CDR data which must be processed and stored using Architecture A as shown in Figure 5.

References

- 12
1. M. Sestak, Traffic Flow Monitoring in PBX Networks, M.Eng. thesis, Carleton University, Ottawa, Canada, K1S 5B6, May 1990.
2. "Voice Reference Traffic Rates", Internal document QAAM035, Mitel Corporation, October 1985.

Table I: Heavy traffic load for a single PBX

Configuration		Heavy Traffic Calls per Hour			
Lines	Trunks	Intra-PBX	Outgoing External	Incoming External	Total
500	71	765	707	521	1993
1000	122	1638	1420	981	4039
1500	176	2402	2228	1466	6096
2000	230	3159	2971	1954	8084
2500	283	3901	3707	2438	10046
3000	336	4642	4448	2926	12016
3500	388	5378	5189	3414	13981
4000	441	6109	5931	3902	15942

Table II: Traffic load for a 4-PBX network

Heavy Traffic Calls per hour	
Intra-Pbx	1147
Outgoing external	1420
Incoming external	245
Incoming external to outgoing network	736
Outgoing network	492
Incoming network	1228
Additional tandem	615
Total	5883

Bits	Field description
1-8	Record length
9-10	CDR type (Establishment)
11	Call type (Voice or data)
17-32	Date
33-56	Start time
57-96	Calling & called party
97-104	PBX identifier
105-136	Digits dialled (Variable length, max. 26 digits)

a) Call establishment CDR

Bits	Field description
1-2	CDR type (Termination)
3-8	Time to answer (sec.)
9-24	Date
25-48	End time
49-64	Trunk identifier
65-72	PBX identifier

b) Call termination CDR

Figure 1: Architecture A: CDRs

Bits	Field description
1-8	Record length
9-10	CDR type (Establishment)
11	Call type (Voice or data)
17-32	Date
33-56	Start time
57-64	Digits dialled length
65-96	Digits dialled (Variable)
97-	Call path (Variable length)

a) Call establishment CDR

Bits	Field description
1-2	CDR type (Termination)
3-8	Time to answer (sec.)
9-24	Date
25-48	End time
49-64	Trunk identifier
65-72	PBX identifier

b) Call termination CDR

Figure 2: Architecture B: CDRs

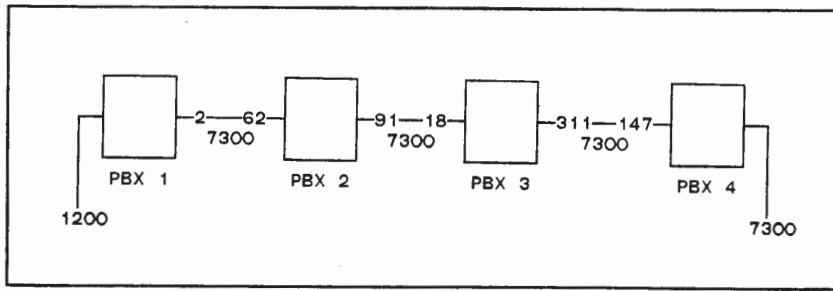


Figure 3: Intra-network call example




Figure 4: Call-map matrix - RAM storage of CDRs

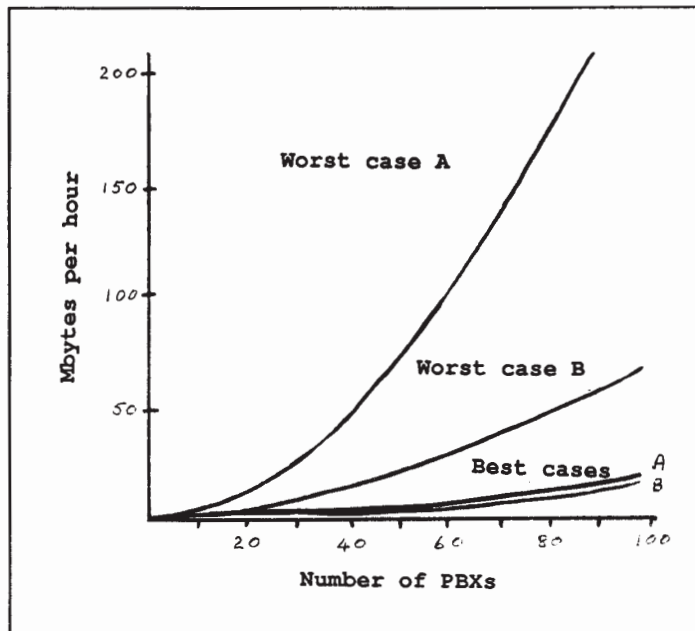


Figure 5: Disk space for a network of PBXs with 2000 lines and 230 trunks