An Efficient Low-Cost Traffic Management System for Large Switching Networks*

Ana Luiza Barros Diniz[1] *Sérgio Campos*[1] *José Marcos Nogueira*[1]
Christiano Mata Machado[2] *Luiz Geraldo Souza*[2]

1: Departamento de Ciência da Computação
Universidade Federal de Minas Gerais
{`analuiza,scampos,jmarcos`}`@dcc.ufmg.br`

2: Telemar Minas
{`chris,lgeraldo`}`@telemig.com.br`

Abstract

Traffic management is a fundamental task in a telecommunications company. It uncovers actual traffic capacity in plant switches and how this capacity is being used. This work describes the design and implementation of a traffic management system that has been developed in a partnership by Telemar-MG and Departamento de Ciência da Computação-Universidade Federal de Minas Gerais. This system collects, transmits and stores traffic data, as well as generates reports and data graphics of switch traffic, trunk and signaling trunk traffic. To collect data and generate alarms we have used the SIS system. From the data collected it is possible to generate graphics that show switch traffic. This system allows the analysis of switch behavior and the identification of problems that may be occurring or of problems that may potentially occur, before they manifest themselves.

1 Introduction

Managing a large telecom network is a vital but extremely difficult task. The Telemar company, for example, has on the order of hundreds of large and medium scale phone switches. It is essential to know at any time which switches are functioning properly and which ones have problems. If a problem is not detected soon, the company loses revenue, since thousands of calls may be left unanswered.

Problems can be categorized into two classes: in the first case, failures take the switch off line. These failures are easily detected, and in this case technicians can be sent to fix the problem immediately. However, it is often the case that more subtle problems may occur, making *part* of the switch to stop working. Some calls can be completed, but not as many as expected. Performance suffers, with only a fraction of the requested calls being

---

*This work has been supported in part by Telemar Minas (former Telemig)

completed. Another cause for performance problems is high traffic. Too many connection attempts create congestion and can cause a drop in performance. These problems are much harder to detect, since the effects are not apparent because switch does not stop.

Using a traffic management system we can solve this problem. The traffic management system collects data from phone switches relating to traffic that goes through the switch. This data is used to identify changes in traffic pattern that may indicate potential problems. In this work we present the design and implementation of a traffic management system that has been developed for Telemar-MG, SGT. This system is part of SIS — The Integrated Supervision System developed by the Computer Science Department of the UFMG university in partnership with Telemar-MG to supervise its telecommunication network. The SGT system allows the detection of traffic loss in switches and trunks *before* they become critical. Correcting these problems becomes significantly cheaper in this case. The system can also be used to redimension the network. Knowing the traffic pattern in a certain trunk or switch enables the company to estimate its ideal capacity and therefore plan expansions more accurately. Besides switches and trunks, the SGT system also manages switch processors, common devices and signalling trunks.

There exist commercial traffic management systems available such as [1, 2]. However, these systems tend to be very expensive. Moreover, they usually put strong restrictions on the user, such as a limited number of licenses, high maintenance costs, the necessity to perform changes only through the original vendor. The SGT system, on the other hand, has been developed at a low cost and in a very short period of time, due to the use of the existing SIS infrastructure. There are no restrictions on the number of users. In fact, the access to the system is done using the world wide web. It is possible to access all reports from the system using any browser. Finally, the source code for the SGT system is part of the SIS system. This makes it straightforward to maintain the system efficiently.

Another important advantage of SGT is that it uses an unified data model. Traffic data from all types of switches are stored in the same format. This simplifies handling the data, but also simplifies management. Each switch vendor generates traffic reports in a different format, and frequently using values that are not directly comparable with those of other vendors. SGT translates all types of reports into an unified model, which makes it possible to manage the network as a homogeneous plant.

Because of these features, we believe that the SGT system provides an efficient and economic alternative to telecom traffic management. It can handle a large number of elements efficiently, enabling its use on companies with large plants. It has been inexpensive to develop and it is also inexpensive to maintain since the source code is included.

2 Traffic Management

Grade of Service Computation

The behavior of a switching network can be understood using the model described below. This model has been proposed by A. K. Erlang [3]. In our analysis, each entity on the telephone network (e.g. switches or trunks) is seen as a *connection server*, and requests for connections are seen as clients for the server. For example, a trunk connects two switches and has the capacity to transmit n calls simultaneously. A new call is considered a request for the trunk server. Besides the capacity of the server, two other parameters define the behavior of the trunk: the arrival rate of requests λ and the completion rate μ. We assume that requests for connection arrive at rate λ following a Poisson arrival process. If there are less than n ongoing calls when a new request arrives, it is served. If the trunk is at its maximum capacity, the new request is denied service. From these characteristics we can determine the probability that a new request will be blocked P_b, called *blocking probability*. P_b is computed from the formula below, called *Erlang-B* formula, where $\rho = \lambda/\mu$:

$$P_b = \frac{\rho^N/N!}{\sum_{l=0}^{N} \rho^l/l!}$$

The value ρ is the average number of ongoing calls in the server, i.e., the traffic on the trunk at any instant. This value is measured in *erlangs*: one erlang corresponds to the traffic generated by one call with the duration of one hour. The derivation of this formula can be found in [3].

The blocking probability is the most important result in traffic analysis. P_b corresponds to the fraction of calls that are lost due to congestion. It is imperative to minimize P_b, because each lost call reduces the company revenue. The blocking probability can also be seen in another way. The higher the probability, the worst service is being provided to the client. Because of this, P_b is frequently referred to as *grade of service*, GOS. A traffic management system must determine the GOS of each entity and warn the manager whenever this value becomes critical.

Representative Values for Traffic

However, the computation of the GOS is not enough to manage traffic effectively, because it refers to the traffic on an instant. It is necessary to analyze the behavior during an *interval* of time. But in order to do this, we must collect traffic data periodically and summarize this data. Data is usually collected and summarized every 15 minutes. With the 15 minute data it is possible to determine the time of the day that has the highest traffic, called *busy hour*.

The busy hour is defined as the period of 60 minutes during the day that had the highest traffic. The traffic during the busy hour is considered the representative value of traffic for the day. The concept of representative value is then extended to weeks, months and years. Consequently we can determine a representative value of traffic for each of these periods. The main objective of traffic management is to compute representative values that characterize the behavior of the switch or trunk.

3 Traffic Management in the Real World

Traffic management is a simple problem in its essence, but extremely complex in its implementation. The large amounts of data generated by the switches make the implementation of the system very difficult. Large switches can have hundreds of trunks and serve thousands of calls simultaneously. But the traffic management system has to collect data from each switch (and every entity in the switch) every 15 minutes. The need to collect and condense this data makes an accurate analysis difficult.

Traditionally, a simplified model is used to analyze traffic. But this model reduces significantly the accuracy of the analysis. It is assumed that the busy hour is fixed and traffic is analyzed only during this period. For example, the busy hour is assumed to be between 9AM and 11AM. This method makes it much simpler to analyze the network, but at a high cost in precision. Several events that can modify the behavior of the network can easily be ignored if these events occur outside of the interval of interest. An example is the growth of the Internet. The Internet has changed significantly the traffic pattern on the switching network, generating heavy traffic during traditionally low traffic times such as evenings and nights. A traditional analysis cannot identify these changes of behavior.

Other factors make it difficult to implement a traffic management system. One example is that frequently switches are located in remote areas. Another one is the heterogeneity of the plant. A network such as Telemar-MG's is composed of switches from several manufacturers, each providing different models and versions of each switch. Because there is no standard for accessing telephone switches, each model uses a different interface. In many cases the interface assumes that data will be seen on a screen or printed directly, and therefore have no concise and clear syntax. Moreover, each switch provides traffic data in a different way. A global analysis of the network becomes extremely complex because it is necessary to understand how each switch works. It is often the case that the values provided by one switch are not provided by another one, forcing the system to implement a common model and translators from each type of switch

4 The SGT Traffic Management System

The SGT system has been designed to overcome the problems described above and provide an efficient low-cost traffic management tool for a heterogeneous network. To achieve this end we have developed SGT completely integrated with the SIS system. Traffic data is collected from plant switches and stored in a database. This data can then be analyzed to determine the behavior of each element, looking for events that can cause service degradation. In addition, the system generates alarms automatically when predefined conditions occur, warning the manager immediately when an abnormal situation exists [4].

Several activities within a telecommunications company can be performed with the assistance of the SGT system. Some areas where SGT can be used are: failure management, configuration, planning and maintenance. It helps failure management and maintenance by generating alarms when the GOS increases, indicating problems. To assist in the configuration of the plant, the SGT provides information about congested or idle switches, generating a more accurate picture of the capacity of the plant. By providing information about all switches and trunks on the plant, SGT creates a global picture of the network. By also storing this information over time it also creates a historic view to the network which makes it possible to use past history of traffic to plan the evolution of the system.

4.1 The Information Model

The information model describes each entity managed by the system, the relationship between them, the most important data collected from each one and how this data is stored. The entities modeled are described below. Figure 1 shows the internal structure of a switch and how each each entity relates to one another.

The switches are the main element of the network. A switch has a subscribers module, a switching matrix, common devices, interconnection trunks for switched traffic and signalling traffic and one or more processors. The switching matrix of a switch is responsible to connect the various traffic sources. Subscribers connect to a main switch through a subscribers module. Internal trunks connect the subscribers module to the switching matrix and or the common devices.

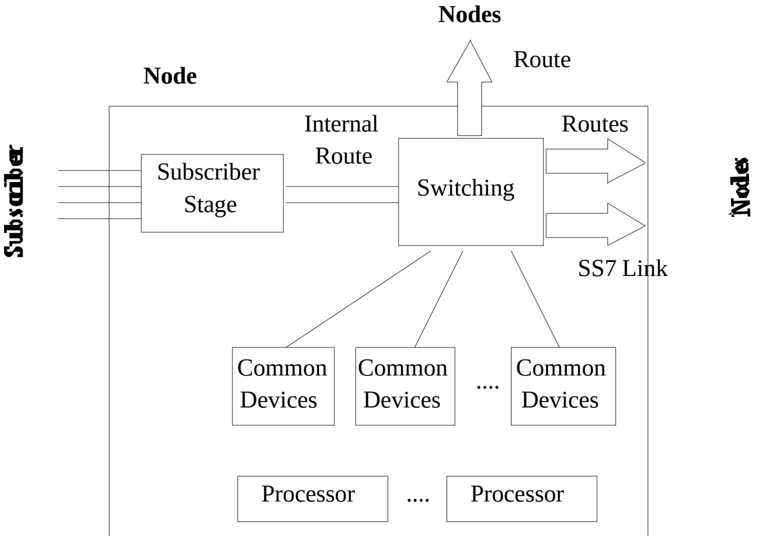

Figure 1: Structure of a switch

There are six types of traffic that must be considered inside a switch, as illustrated in figure 2. The *originated* traffic is the traffic that relates to subscribers connected directly to the switch. The *internal* traffic corresponds to calls between subscribers connected to the same switch. Internal traffic does not leave the switch. *Inbound* traffic refers to traffic arriving from other switches. *Transit* traffic is part of the inbound traffic that is directed towards other switches, i.e., to a subscriber not on the switch. The *outbound* traffic is traffic that is destined to other switches. The *terminating* traffic is the addition of the inbound traffic directed to subscribers in the switch and of the internal traffic. The total traffic is the sum of originated and inbound traffic. It can also be computed as the sum of the outbound and terminating traffic.

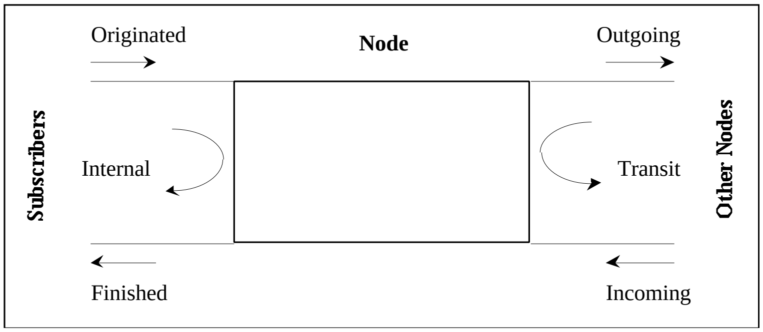

Figure 2: Types of traffic in a switch

Trunks are sets of lines that connect two switches. Trunks are used to transmit voice, user data and signalling data. Signalling data can be transmitted using independent physical channels or on the same line as the voice data.

A switch has one or more processors. There are two types of processors, common and signalling processors. Common processors are responsible for processing calls, and signalling processors are responsible for dealing with signalling information only.

In the SGT system we have created an unified information model, where data from all types of switches are mapped into the same data model. An unified model solves the problem of dealing with switches from different manufacturers. Reports from each type of switch are translated into our unified model. In this way the SGT system only needs to deal with one type of traffic data, simplifying the implementation and use of the system.

An additional problem is that not all switches generate reports containing the same information. There are traffic values that cannot be obtained from certain types of switches. Because of this our unified model allows the possibility that some traffic parameters may be left unspecified. Consequently, reports may vary from one type of switch to another. This is, however, unavoidable because of the characteristics of each switch type.

4.2 Traffic Computation

a switch has traffic counters that store, among other parameters, the amount of traffic on the switch and the number of trunks available at any time. SGT reads these counters every 15 minutes. Each reading is called a measurement. All measurements collect data for all entities in the system. SGT stores temporarily all data from the last four measurements. From this data we determine the time and traffic of the busy hour. In the SGT database there are two busy hours, one for daytime and another for nighttime.

The most important information derived from this data are the traffic and the grade of service. At the end of each day the representative traffic value for the day is computed. The representative values for the week, month and year are computed similarly.

4.3 SGT General Architecture

The SGT system collects, stores and analyzes traffic data. The architecture of the system can be seen in figure 3. Data is collected from each switch every fifteen minutes, processed and stored on disk. The resulting data is stored in the unified model, regardless of the type of the switch that generates it. Each file containing traffic data is then sent to the database, where the information is permanently stored. Traffic reports can be generated as text or graphics, showing the traffic pattern for that entity for the period. For example, a report can show hourly traffic for a trunk for a given day. Besides the hourly traffic, the report also shows the GOS for each hour and the busy hour of the day (see in section 4.6).

The most important components of the SGT are:

- **Configuration**. This subsystem is responsible for maintaining the static data of the system, such as switch location or the number of trunks connecting two switches.
- **Alarm Generation**. Analyzes the traffic values and GOS measured at each 15 minute interval and, if needed, generates alarms.
- **Report Generations**. Recovers database information and generates traffic reports.

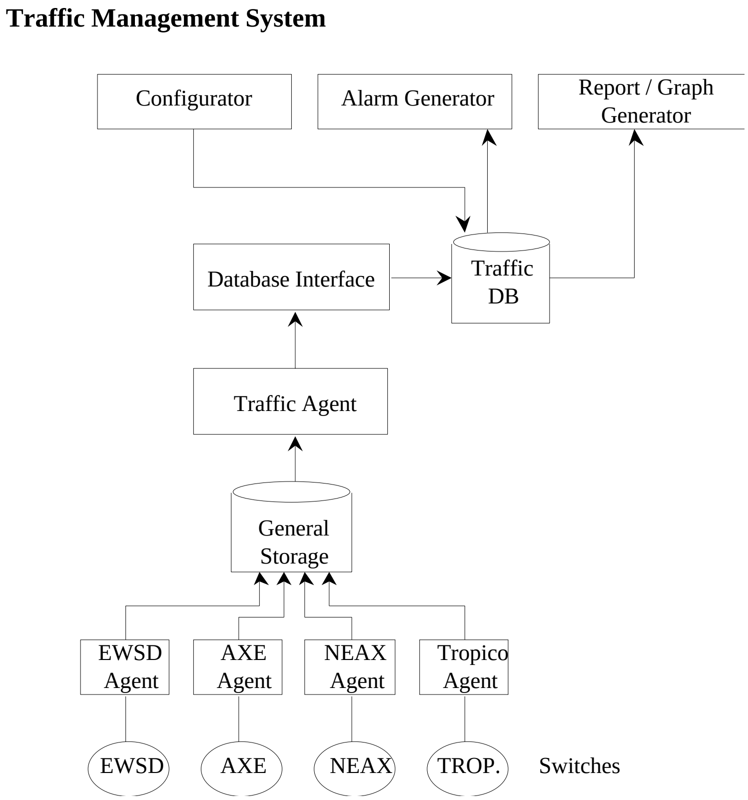

Figure 3: SGT Architecture

Figure 4 shows how data flows in the SGT system. The main modules are the data collecting agent, the traffic agent, the database and its interface.

The Data Collecting Agent

Data collection is performed by specific agents, developed for each switch type separately. The switches monitored are always digital. The agent orders the generation of traffic reports directly on the switch. From each report we extract data such as traffic in the last 15 minutes and the number of lines being used to service this traffic. The information extracted is stored in local files on the XDR (External Data Representation) format and then transferred to a global repository.

The Traffic Agent

The traffic agent is responsible for reading the files on the repository and transmitting them to the database interface. This agent also receives a list of alarms that have been generated by the database interface. The traffic agent then generates the alarms.

Database Interface

The database interface receives data from the traffic agent and updates the database accordingly. It is also responsible for computing the GOS as described in section 2. At each measurement the traffic in each entity is analyzed, and if needed, a list of alarms is created to be sent to the traffic agent.

Database

We use a relational database. The entities and their attributes are listed in [4] and the database specification can be found in [5].

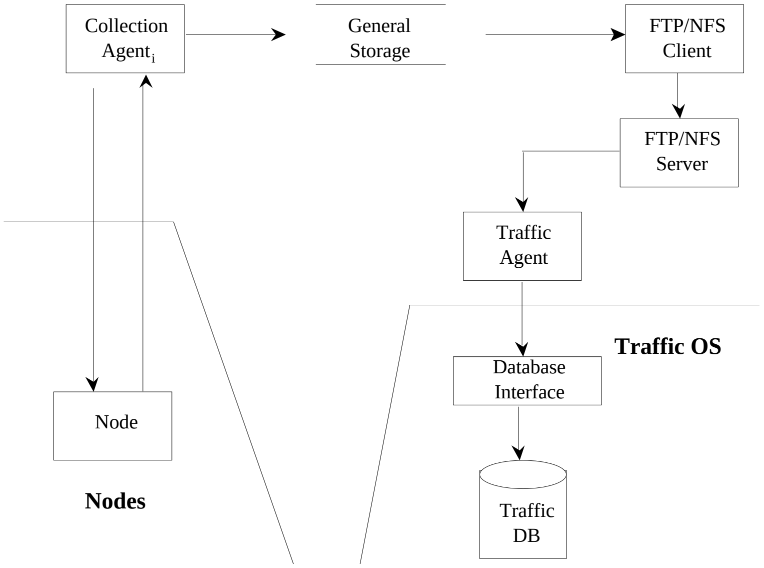

Figure 4: Data flux in SGT

4.4 Reliability of the Data

It is always possible that during the operation of the system failures occur during data collection. This compromises system operation because data must be summarized continuously. One common occurrence is the absence of one measurement. The SGT system implements an interpolation algorithm to compensate for missing measurements. This algorithm tries to ensure that data is always consistent in the database.

The system always knows which measurement should come next. When the next one actually arrives the system checks to see if it is the one expected. If the new measurement is older than expected, it is ignored. If it is newer than the one expected, then one or more measurements have been lost. SGT then averages traffic on the previous three measurements to compensate for the missing one. Up to three consecutive lost measurements can be interpolated. If more then three are lost an alarm is generated. Using this algorithm, some measurements may be lost without compromising data consistency.

4.5 Integration with the SIS System

The Integrated Supervision System (SIS) is a distributed network management platform [6]. The SIS system is based on the client-server paradigm. Several agents collect data from various types of switch. SIS provides reliable communication between switches, agents and central databases. The SIS architecture can be seen in figure 5.

This architecture in conjunction with the ability to communicate with virtually all types of switches makes SIS an ideal host for the SGT system. Therefore, we have developed SGT using as many SIS functions as possible. Data is collected and transmitted by SIS to the SGT system. Alarm generation has been also simplified, since SIS manages failures as its main function, and alarms can be treated as a special type of failure. The connection of SGT and SIS can be seen in figure 6, where UC modules correspond to central SIS management units. They can also be subdivided in principal, regional and secondary units, respectively UCP, UCR and UCS.

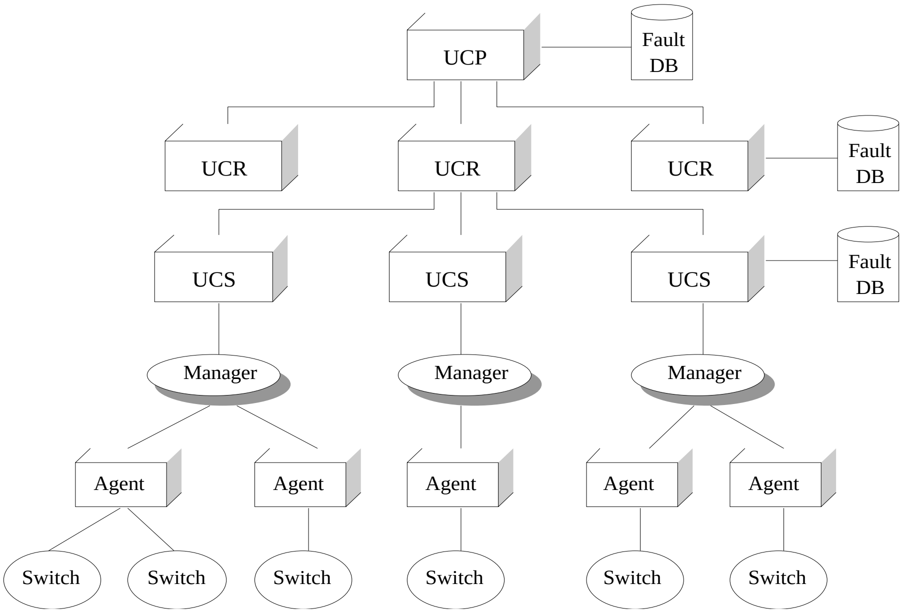

Figure 5: Structure of SIS

4.6 Traffic Reports

One of the difficulties of implementing a traffic management system is the analysis of the data collected. The volume of data is very large, making it necessary the use of tools to assist in this analysis. The report generation subsystem performs this function in SGT. Graphic and text reports generated by this subsystem can be accessed using the web. Graphics can be seen by year, month, week, day and hour. The yearly report contains the representative values for the last twelve months of the year. Moving through these graphics is straightforward. By clicking on a day, for instance, brings up the hourly graphic for that day. The reports show traffic values, GOS and busy hour for all periods. The GOS is shown as a color scale that indicates idleness, normal occupation or congestion.

Figure 7 shows a typical traffic report, with hourly traffics for 24 hours. These values are real traffic values. Each bar shows the information for one hour. The height of the bar indicates the capacity of the switch. The dark gray bar shows the actual traffic, while the light gray shows the blocked capacity, i.e., how much of the switch was not working at that time. The white portion of the bar shows the active capacity of the switch. When a large light gray area exists this means that the switch may have problems, because it blocks a large part of the traffic it was supposed to handle. The small bar below the graphic indicates the day and night busy hours. In this figure, the day busy hour occurred at about 11AM and the night one around 6PM.

5 Results

Currently, the SGT system is monitoring several large Telemar switches. Initially, EWSD switches (from Siemens) have been monitored, with agents collecting data using the FTAM protocol. Other agents that have been or will be incorporated to SGT include agents for switches Tropico-RA, AXE (from Ericsson) and NEAX (from NEC). The graphic and text reports are also available over the web. These reports can be visualized on a browser

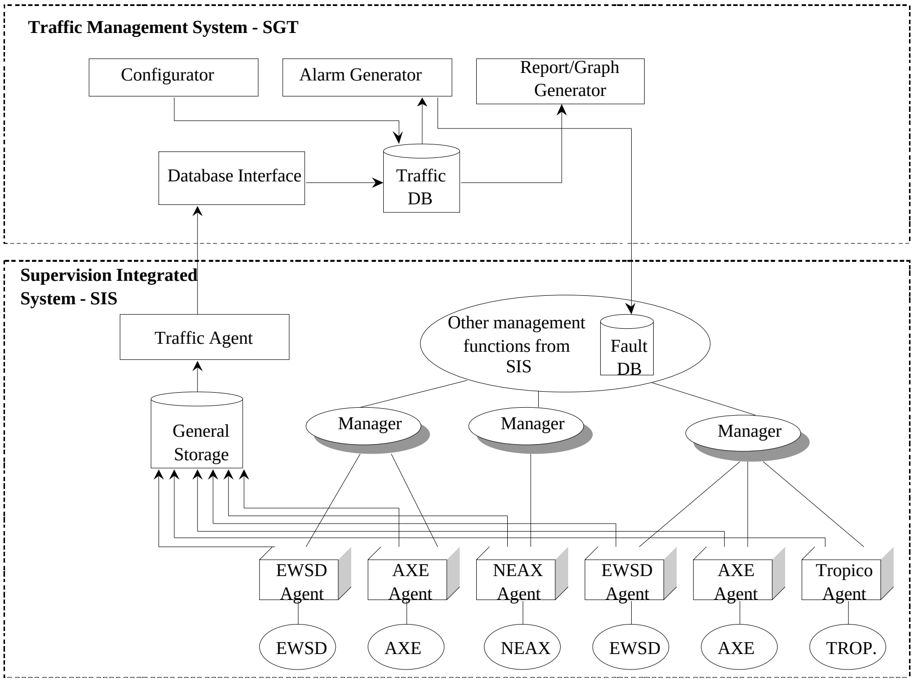

Figure 6: Integration SGT-SIS

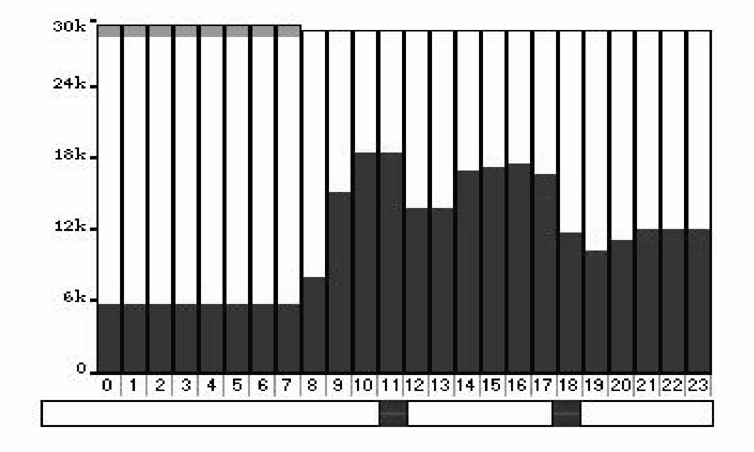

Figure 7: Typical switch traffic

or using the SIS subsystem SisNave.

The reports generated enable a comprehensive analysis of the performance of switches. The information produced show clearly the amount of traffic that is actually carried by a switch compared with the amount of traffic that *can* be carried. The result of this comparison can show if a switch has problems in many ways. For example, its active capacity may have become smaller due to significant blocking of lines which would not have been noticed unless congestion occurred, and in this case it would have been late to avoid problems. Another example is that traffic may be increasing consistently over time, indicating that an upgrade will be necessary. Without SGT this may have been noticed only when congestion was critical. With SGT it is possible to have this information on a glance, since a critical situation is indicated by a different color (meaning a worst GOS).

An an illustration of how SGT can be used, we have used it to identify how Internet traffic can affect switches that have Internet providers as subscribers. The graphics shown here are all data taken from actual Telemar switches under normal operation. Normally, telephone traffic is significantly higher during business hours, as seen in figure 7. However, a switch hosting Internet providers has a different pattern. Besides the usual business hour traffic, it has significant traffic at night, as shown in figure 8. On Saturdays, when prices are lower, we can observe a more pronounced effect, with the highest traffic at night, and the day busy hour late in the evening. This behavior is illustrated in figure 9. Notice that SGT not only shows that traffic changes significantly under different conditions, but it is also able to *quantify* these changes, making it possible to not only identify potential problems, but also to determine how much is needed to correct it. For example, switches that host Internet providers have to be dimensioned differently from other switches.

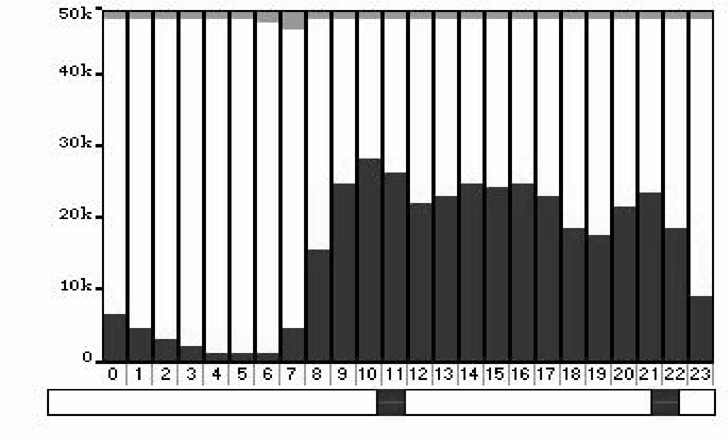

Figure 8: Internet provider traffic — weekday

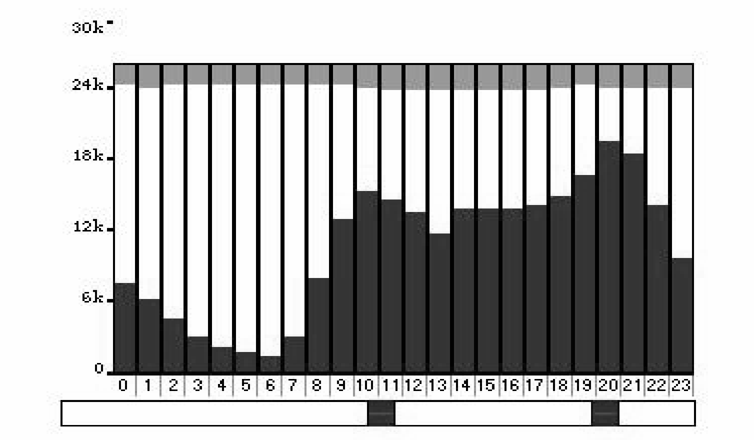

Figure 9: Internet provider traffic — Saturday

6 Conclusion

The SGT traffic management system has been developed to allow a more accurate analysis of switched traffic. As a consequence, the use of SGT can help telecommunication companies to avoid congestion problems, plan upgrades more efficiently, improve its service quality. The final objective is to make the company more profitable by better utilizing its resources.

Several goals have been set for the system. Low cost was an important objective, as well as integration with SIS. Accuracy, high availability, and the ability to collect, store and analyze a very large amount of data. The final result has been an efficient system that allows traffic management to be performed as well as most commercial systems, but at a significantly lower cost.

Acknowledgments

Several other people have participated in the development of SGT during all phases of the project. Among those the authors would like to thank Lina A. Silva, Amilton J. Araújo, Er M. Tassara, André Luiz P. S. Silva, Renato Augusto V. Leite, Sibele V. Oliveira, Geórgia C. Penido e Ângelo M. Menezes.

References

[1] Metrica/npr 3.0. *http://www.metrica.com/index.products.html*, April 1999.

[2] Network planning and traffic capacity management solution. Disponível em *http://www.telcordia.com/solutions/operations/network*, April 1999.

[3] M. Schwartz. *Telecommunication Networks: Protocols, Modeling and Analysis.* Addison Wesley, 1988.

[4] Ana Luiza B. B. Diniz, Sérgio V. Campos, Amilton J. Araújo, André Luiz P. S. Silva, Christiano F. Machado, Er M. Tassara, Lina A. Silva, and Luiz G. P. Souza. Especificação do serviço de administração de tráfego de nós e rotas. RT SIS 5111, UFMG-DCC-ICEx, Belo Horizonte-MG, Novembro 1998. 1ª edição.

[5] Ana Luiza B. B. Diniz, Sérgio V. Campos, Christiano F. Machado, Lina A. Silva, and Luiz G. P. Souza. Especificação do banco de dados do serviço da administração de tráfego de nós e rotas. RT SIS 5112, UFMG-DCC-ICEx, Belo Horizonte-MG, Novembro 1998. 1ª edição.

[6] Meira Dilmar M. Nogueira, José Marcos. The sis project: A distributed platform for the integration of telecommunication management systems. In *Proccedings of IEEE/IFIP 1996 Network Operations and Management Symposium - NOMS'96*, Kyoto, Japan, 15 to 19 of April 1996.