

Management Services for Performance Verification in Broadband Multi-Service Networks

Panos Georgatsos¹, David Griffin²

Abstract

This paper presents a practical management system for performance monitoring and network performance verification to support the larger goals of performance management systems. We show how the rich and powerful features of OSI systems management can be used in a hierarchical manner in the TMN, to achieve sophisticated performance monitoring and performance verification without imposing a large communications overhead in the TMN and hence in the managed network itself.

1. Introduction

The intelligence provided by network management has been widely recognized as an important aspect of telecommunications networks. The coexistence of different services, with potentially widely differing requirements on performance and quality of service, on the same broadband networking infrastructure, imposes the need for effective management schemes. ATM technology itself provides several degrees of freedom for multiplexing traffic that if not managed properly can prove disastrous in terms of network provisioning and performance.

The ITU-T introduced the Telecommunications Management Network (TMN) [9], as a means of provisioning the required management intelligence and have distinguished between the management and control planes in the operation of communications networks [6], [7]. The TMN relies on the OSI systems management concepts and functions, developed by ISO and ITU-T (cf. X.700 series recommendations), to model the network and service resources at various levels of abstraction and the communication of management information. Utilizing these concepts, TMN suggests a hierarchical management architecture enabling separation of concerns and encapsulation of lower level functionality. TMN therefore, implies a hierarchical distribution of management intelligence and it can be regarded as a separate network, logically distinct from the network being managed.

The majority of the management systems deployed today are concerned with network configuration and network monitoring. There is a trend [1]-[5] to increase the intelligence of management functions to encapsulate human management intelligence in decision making management components to move towards the automation of the monitoring, decision making and configuration management loop.

Performance management has been identified as one of the major management areas by ISO and ITU-T. The aim of the performance management functions is to guarantee that the network meets the required performance targets of the range of services that it supports. The on-going work in performance management related areas (routing, VPC

¹ ALPHA Systems S.A., Athens, Greece, Tel: +30.1.48 26 014-16, Email: panos@alpha.ath.forthnet.gr

² FORTH-ICS, Heraklion, Greece, Tel: +30 81 39 17 22, Email: david@ics.forth.gr

bandwidth management, etc.) justifies the need for efficient network performance management systems.

The actual efficiency of a performance management system cannot be established unless reliable measures on network performance are provided. Furthermore, network performance assessment is required for identifying undesirable trends in network performance and triggering the appropriate management functions so that the necessary corrective actions to be taken. In [2] the need for network performance verification capabilities has been emerged as an essential part of a management architecture for VPC and routing management.

Taking into account the multi-service environment, network performance should be evaluated and assessed on the basis of measures relating to the individual performance characteristics of the bearer connection types (or classes of service, CoSs) supported by the network. Although there are practical implementations of network monitoring and performance evaluation systems, the problem of network performance assessment from a network management perspective and taking into account the different performance characteristics of the multi-class environment, has not been fully addressed.

Within the above framework and recognizing the emerging need for enhanced automated network management systems, this paper concentrates on the problem of network performance assessment within the overall context of network management.

A TMN approach is followed. The paper defines an appropriate management service, the Performance Verification management service, for network performance assessment. The paper describes the Performance Verification management service in the context of the performance management functional area and proposes a specific approach and algorithms to network performance verification. The management service is decomposed and mapped to the TMN architectural framework.

By taking advantage of the rich and powerful features of OSI systems management, the paper shows how the calculation of the required measures can be made for all CoSs over all source destination pairs, introducing minimum overhead into the management system. Specifically, by pushing the monitoring functionality down to the NEs, polling from the network management layer is avoided; the Performance Verification management service residing in the network management layer receives only the emitted threshold crossing notifications indicating unacceptable network performance. Requirements on supporting monitoring objects are identified as well as enhancements on the 'classical' OSI metric and summarization monitoring objects [16], [17].

The paper is organized as follows: Section 2 introduces the Performance Verification management service, describing its objectives and its interactions with other management services/components and the network. Section 3 describes the main functional aspects of the Performance Verification management service and outlines a specific approach for network performance verification. Section 4 presents a TMN compliant management architecture for the Performance Verification management service. Section 5 elaborates on system design aspects, presenting efficient means for realizing the specified management functionality on the identified architecture. Finally, section 6 presents the conclusions and highlights aspects of future work.

2. The Performance Verification management service

This section describes the scope of the Performance Verification management service in the framework of ATM network management, specifying its objectives and its relationship with other management services/components, the network and the TMN users.

The managed environment is assumed to be a public ATM network supporting a wide range of services made up of network bearer service classes. The term CoS denotes a network bearer service class. Different CoSs are distinguished according to their bandwidth requirements and performance targets. The view taken is that there is a range of network bearer services of different quality (in terms of performance targets) and costs to support the AAL services. This view is in accordance with the views of the ATM Forum [11], where they explicitly recommend the augmentation of the AAL service classes with a range of quality service classes.

Within the above environment, the objectives of the Performance Verification management service are:

- to ensure that the network meets the performance targets of the different CoSs supported by the network,
- to warn performance management related components when network performance has dropped below the acceptance levels per CoS, so that corrective actions can be taken,
- to analyze customer complaints with respect to the quality of the network services they use.

The scope for network performance verification is necessitated from the fact that the network supports multiple CoSs of guaranteed performance. In the case of a network providing connections without a guaranteed performance, the scope for network performance verification and performance management in general is limited to performance monitoring.

Although the Performance Verification management service can be regarded as a management service in its own right, it could well be taken as a management service component of other management services in the performance or other management functional areas (e.g. billing in cases where the network accounting policy allows for compensations when the quality offered by the network is outside the specified levels). In [2] the above management service has been proposed as a management service component of an overall management system architecture for VPC and routing management for networks supporting guaranteed performance connection classes. In particular, this management service was responsible for verifying network performance and for emitting quality of service alarms in case of undesirable trends in network performance.

The Performance Verification management service is beneficial to the network operation in an indirect manner in the sense that it is beneficial to the management system that runs on the top of the network. Specifically, the Performance Verification management service quantifies the performance of the network management system,

providing therefore an indisputable measure of the efficiency of the management system.

The Performance Verification management service should not be confused with the performance monitoring services of the network. The performance monitoring capabilities of a network are responsible for calculating and supplying measurements concerning various network entities like nodes, VPCs, VCCs. The Performance Verification activities are making use of the network monitoring capabilities, requesting the measurement of the statistics required for network performance evaluation and verification. In the same sense other network management components - in the performance or other management functional areas (e.g. accounting) - are making use of the network monitoring capabilities for the purpose of applying their intelligence. Therefore, the functionalities of the Performance Verification and Performance Monitoring management services are quite distinct.

The Performance Verification management service belongs to the performance management functional area. Figure 1 shows the relationship of the Performance Verification management service with other performance management service, management functional areas and the TMN users.

The boundaries of the management responsibility of the Performance Verification management service are shown in Figure 2, which depicts the interactions between the management and control planes from the point of view of the Performance Verification management service.

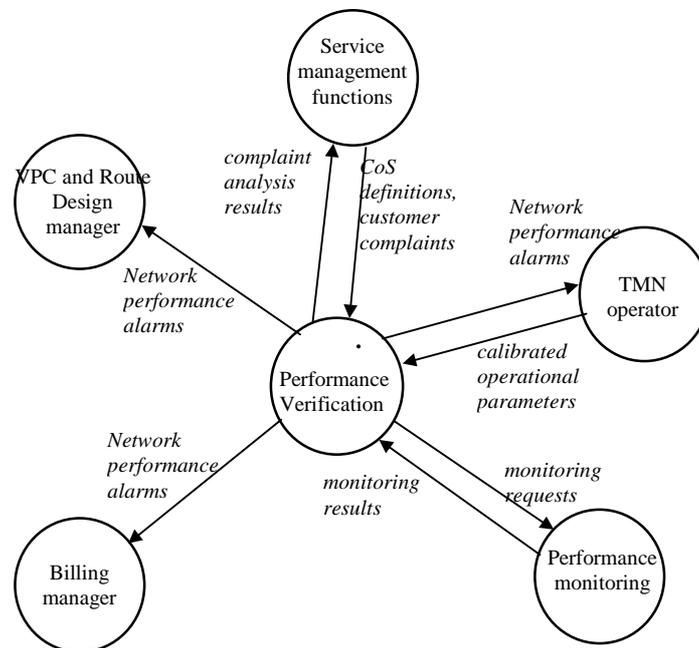


Figure 1: Enterprise view of the Performance Verification management service

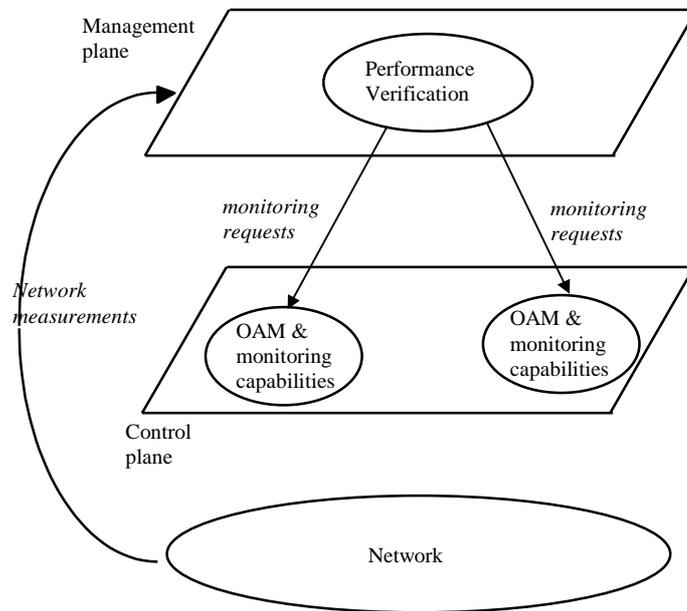


Figure 2: Management and Control plane interactions.

3. Performance Verification functional aspects and approach

This section highlights the main aspects of the functionality involved in the Performance Verification management service and proposes a specific approach for achieving its objectives. As implied by its objectives, the functionality of Performance Verification management service includes

- evaluation and verification of network performance with the purpose to notify other management services of network unacceptable performance and to analyze customer complaints with regard network performance.

Evaluation and verification of network performance is done on the basis of measurements concerning the performance of the CoSs that the network supports. It is assumed that the following parameters define the performance of the CoSs:

- rejection ratio (blocking probability),
- cell delay,
- cell loss ratio,
- jitter.

The above parameters are those which are directly influenced by the activities of the VPC and route design managers [2]. Other performance parameters may be associated with network CoSs, such as set-up delay; but because the role of Performance Verification is to feedback performance analysis results to the performance management system, to influence its future behaviour, these parameters are outside the scope of Performance Verification.

These parameters have been widely accepted as meaningful connection performance parameters and they are in accordance with the performance parameters defined by ATM Forum [11]. The jitter refers to the variance in the cell delay within a connection; it is equivalent to the 2-pt CDV (two point cell delay variation) defined by ATM Forum [11]. It is assumed that for each CoS there is an upper bound (performance target) defining the range of acceptable values. Network performance is within acceptable levels if these bounds are preserved within certain confidence limits.

Measurements of rejection ratio statistics could be done directly from appropriate raw data (e.g. counters of connection requests and rejections) available at the management interfaces of the network elements. The measurement of cell related statistics (delay, loss ratio, jitter) requires the existence of measurement instruments at the source and destination end-points or the existence of OAM (operation and maintenance) capabilities at access switches including appropriate management interfaces.

Based on measurements on the above parameters, Performance Verification evaluates network performance and by comparing it with the maximum allowable values (performance targets) verifies whether network performance is within acceptable levels or not. Evaluation and verification should be done in the scope of the total population of CoSs and source-destination pairs. Therefore the measurements should cover all (or a representative portion of) the source-destination pairs and supported CoSs, to increase the validity of the produced results. On the other hand, the verification process should introduce minimum overhead to the management system for collecting the required network information and it should not be sensitive to transitory situations.

For a given CoS, network performance evaluation and verification could be done in two modes:

- Per source-destination (s-d) pair, or
- Network-wide (statistically averaged over all possible s-d pairs by sampling)

Network-wide performance estimates could be taken either exhaustively (over all possible s-d pairs) or following sampling techniques in cases where the number of all s-d pairs is huge. In the latter case, sampling per strata should be pursued, to guarantee homogeneity of the population strata.

The proposed approach to network performance verification adopts the first mode, since it is considered more useful to the management services requiring network performance deterioration alarms. Indeed, the routing management systems construct and manage routes of guaranteed quality for all source-destination pairs for a given CoS; therefore, as a means for quantifying routing management efficiency, network performance estimates per s-d pair are required. This view was adopted in [2], where a hierarchical management architecture for VPC and routing management was proposed. Network-wide performance estimates may hide unacceptable performance situations for some s-d pairs. Moreover, network-wide performance estimates inevitably require the use of polling for retrieving the necessary sample values which subsequently creates a communications overhead in the management system. As it will be shown in the following sections, network performance verification per CoS and per s-d pair, paradoxically enough, although all s-d pairs are involved, does not require as

great a management overhead as in the network-wide case; in fact it creates minimal management overhead.

For a specific s-d pair and a given CoS the essence of the proposed approach is to calculate the following probabilities as a means of ensuring and therefore verifying that the network performance is within acceptable levels:

$$\text{Prob}^{(i)}[\mathbf{R}_{sd} \leq {}^{(i)}T_r] \geq A \quad (1)$$

$$\text{Prob}^{(i)}[\mathbf{R}_{sd} > {}^{(i)}T_r] < 1 - A \quad (1a)$$

where:

${}^{(i)}\mathbf{R}_{sd}$ denotes a measurement of a CoS related performance parameter R (rejection ratio, cell delay, cell loss ratio, jitter) for CoS i and source destination pair s-d. For rejection ratio, ${}^{(i)}\mathbf{R}_{sd}$ is calculated as an instantaneous value or as a moving average estimate. For the cell related performance parameters (cell delay, cell loss ratio, jitter), ${}^{(i)}\mathbf{R}_{sd}$ is calculated as the arithmetic average over a specific number of connections.

${}^{(i)}T_r$: is the maximum allowable value of performance parameter R for CoS i.

A: is the confidence level.

Formula (1) says that it is almost certain (with a confidence level A) that the network performance with regard the performance parameter R for connections of CoS i between the source-destination pair s-d will be within the acceptable levels (performance target) specified for that CoS. Violation of condition (1) or (1a) means that the network performance with regard the performance parameter R for CoS i between the source-destination pair s-d has fallen below the acceptable level. In such cases, network unacceptable performance alarms should be triggered in order the necessary corrective actions to be taken by the appropriate management services.

The calculation of the previous probability can be done continuously, or during specific verification periods, with dynamic duration and inter-period times, depending on network state. For a given source access node, the above probability measure could be calculated either exhaustively (over all possible s-d pairs) or for a sample of destination nodes following sampling techniques -in cases where the number of the s-d pairs for all destination nodes is huge. In the latter case, sampling per strata should be pursued, to guarantee homogeneity of the population strata.

Note that by taking the Probability of the event $ER = [{}^{(i)}\mathbf{R}_{sd} \geq {}^{(i)}T_r]$, transitory fluctuations in network performance are not taken into account. The measurement of the probability is taken at a given window, within the verification interval; then the probability is approximated by the frequency (ratio) of the times of occurrences of the event ER over all the observations that were made within this window. The choice of the length of the probability window should be done taking into account the nature of the measurements (instantaneous values or moving average type of values or number of connections over which the cell related performance parameters are calculated) for the performance parameter S. These are left as design options.

Table 1 summarizes the proposed approach outlining the main steps of an algorithm fulfilling the objectives of the Performance Verification management service. The details of the Performance Verification algorithm is not in the scope of the paper, since the paper deals more with architectural and design issues.

<p>Retrieve performance targets ${}^{(i)}T_r$ for all CoS related performance parameters R.</p> <p>For all access source nodes</p> <p>For all CoSs, determine destination nodes over which network performance will be verified</p> <p>For each CoS related performance parameter, initiate appropriate monitoring activities for the selected s-d pairs for monitoring the probability of unacceptable performance (see (1a)).</p> <p>Collect threshold crossing notifications.</p> <p>If one received, wait for a specific time period and collect any other notification that may come within this time period. At the end of the time period send the collected notifications (network unacceptable performance alarms) to the interested management services.</p> <p>Schedule the next verification interval.</p>

Table 1. Performance Verification functionality.

It is worth noting that the Performance Verification management service operating with the previously described functionality creates *minimal overhead* into the management system. As it will be shown in section 5, all the required monitoring activities may be pushed down at the QA (Q3 adaptation) level providing the management interface of the network elements. Only the notifications resulting from threshold crossings are forwarded to the management system. The number of notifications depends apart from the performance of the network, on the sensitivity of the measures. By regulating appropriately the measurement characteristics of the connection related performance parameters and the probability window, the tradeoff between the validity of the measurements and their sensitivity can be managed.

4. Functional architecture

In this section we place the Performance Verification Management Service in the context of an architectural framework that will allow design and implementation of the management functionality introduced in the previous sections. We have adopted the TMN approach [9] recommended by the ITU-T. By following the methodology of Recommendation M.3020 [10], Management Services are decomposed into Management Service Components (MSCs) which are in turn decomposed into Management Functional Components (MFCs). Expanding this methodology, the derived MFCs are mapped to the hierarchical layers of the TMN [9] and to the TMN function blocks of the TMN functional architecture.

The Performance Verification Management Service is decomposed into:

- a *performance verification* MSC, performing the functionality of the Performance Verification management service as outlined in the previous section. This MSC is responsible for controlling the required network monitoring activities, analyzing the measured statistics against performance targets, and for analyzing and validating customer QoS complaints against measured performance. Following the functional analysis presented in the previous section, this MSC is further decomposed into the following MFCs:
 - a *performance analysis* MFC which is responsible for: determining which performance parameters in which network resources and on which

connections should be monitored in order to estimate network performance; comparing measured performance against target performance to identify whether the network is not meeting its performance obligations; and raising appropriate alarms to trigger a re-configuration of the network to resolve performance degradation.

- a *customer complaints analysis* MFC which interfaces to the customer complaints functions in the service level; compares customer complaints on quality degradation to the measured performance in the network; analyses whether the customer complaints are justified; initiates additional performance monitoring probes in the case where complaints arise on specific portions of the network which previously were not being monitored explicitly.
- a *performance management control* MFC which is responsible for: interfacing to the network performance monitoring functions (see below); controlling the initiation and termination of monitoring activities such as logs and alarm thresholds; specifying the attributes of monitoring activities. These activities are performed at the request of the *performance analysis* and the *customer complaints analysis* MFCs.
- a *performance monitoring* MSC which is responsible for retrieving and reporting on the necessary data from the network elements to support performance monitoring. Performance monitoring is performed on two aspects of the network: the network resources and the network CoSs supporting customer calls. The performance monitoring MFC is involved in the collection of raw data as presented by the network elements and also for transforming this data by summarization, averaging, and statistical analysis to more comprehensive forms (e.g. rates, probabilities, averages) according to the requirements of the *performance verification* MSC. The performance monitoring MSC is decomposed into:
 - a *current load model* MFC which is responsible for collecting, storing and reporting on data concerning the traffic on specific network resources (e.g. nodes, links, VPCs).
 - a *connection performance model* MFC which is responsible for collecting, storing and reporting on data concerning network CoSs. This may often be achieved by interacting with the OAM capabilities of network switches, or by means of network monitoring tools.
- a *connection-type model* MSC which stores the performance targets associated with each CoS the network supports. It provides the performance verification MSC with the cell loss, delay, delay jitter, connection rejection ratio, etc. targets. It is decomposed into a single MFC:
 - a *connection-type model* MFC
- a *network model* MSC which stores information about the available network resources and their configuration, including network topology and resource capacity. This MSC is used by the performance verification MSC to identify

where appropriate monitoring probes should be placed. It is decomposed into a single MFC:

- a *network model* MFC

The *performance verification* MSC is specific to the Performance Verification management service, its functionality and algorithms have been derived in the previous section. The *performance monitoring* MSC and the derived *connection performance model* MFC may also be regarded specific to the Performance Verification Management Service; but the *current load model* MFC is more generic and may be used by other management services [2], [3]. The *connection-type model* and the *network model* MSCs are not specific to the Performance Verification Management Service, they are supporting MSCs of a more general use in management systems and may be used by a number of management services. They even may be regarded as MSs in their own right. Note that elsewhere the *network model* MFC may be a sub-component of a configuration management MS.

Figure 3 shows the allocation of MFCs to OSFs and their allocation to the TMN architectural layers. The ICF and MF function blocks [9] have been included to guarantee access from the network layer to the network element layer information models, and to enhance the network element specific information model, respectively. The actual choice of whether the physical counterparts of the QAF, MF or ICF are implemented as appropriate in a NE, QA, MF or a Network Element Level OS is left open according to the capabilities of the underlying network elements. However, it is assumed that a management component, for each network element (or a set of network elements) supplying a Q3 interface will exist below the network management layer.

The functionality of the Performance Verification management service has been placed at the network management layer following the directives implied by the decomposition of the logical TMN architecture. Performance verification is concerned with collecting performance data from a number of network elements, collating that information and comparing performance data obtained from more than one network element. For these reasons, it needs to have a global view of the network, and therefore must exist at the Network Management Layer. However, the performance monitoring functionality is distributed over the network management and element management layers. This distribution allows frequent data collection for a single network element to be carried out close to the source of the data.

A hierarchical system architecture is an important consideration in fulfilling the objectives of the Performance Verification management system. By virtue of the proposed hierarchical system architecture, the management overhead for acquiring the required network statistics may be minimized, as monitoring activities can be delegated down the hierarchy as close as possible to the network elements themselves where the raw performance data is generated.

Based on the identified architecture, and by making use of the rich facilities of OSI management, the following section proposes an efficient means for the design and implementation of the functionality of the Performance Verification management service.

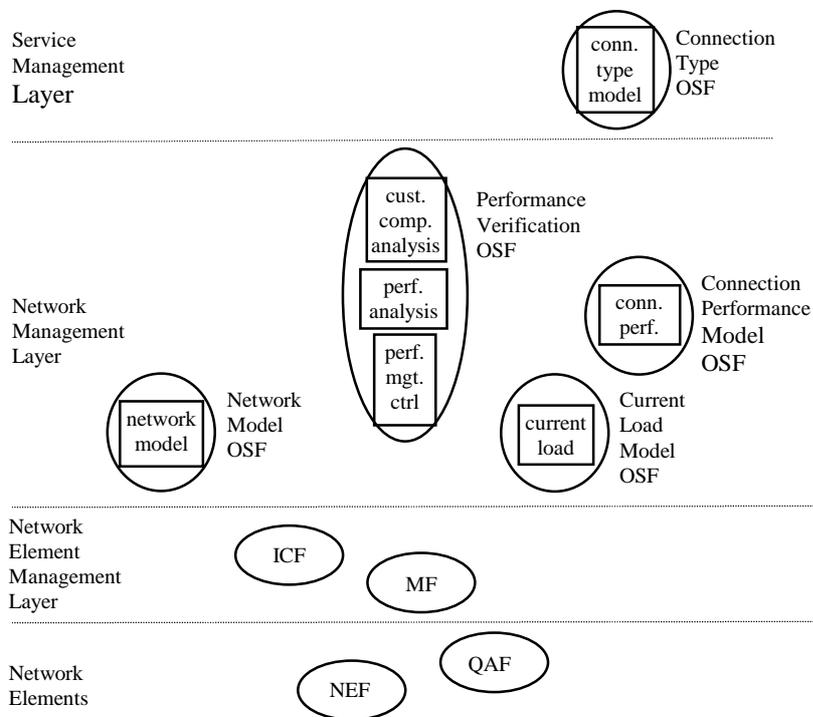


Figure 3: Performance Verification TMN functional architecture.

5. Design aspects

By following the TMN approach which uses OSI systems management concepts, performance monitoring is achieved by virtue of the OSI systems management functions (SMFs). In particular, event reporting [12], alarm reporting [13], log control [14], test management [15], workload monitoring [16] and measurement summarization [17] SMFs are used. These are standard facilities of OSI systems management, pertinent to any Q3 interface, and demonstrate the advantage of adopting the TMN approach for implementing the Performance Verification management service. It is believed that these facilities provide a rich and powerful set of generic management tools, which - when properly used - may considerably reduce the cost of developing and of operating the management system.

The hierarchical TMN architecture and the OSI SMFs imply and furthermore facilitate distribution of the required monitoring activities over the network management, network element management and network element (management interface) layers. This distribution enables frequent data collection for a single network element to be carried out close to the source of the data. The results of monitoring activities in the network elements are then forwarded to the higher management layers either on request or at exception, at threshold crossing instances. This design consideration ensures that the management communications overhead is as small as possible as the bulk of the data in the management plane will be transferred locally.

Furthermore, by adopting the per s-d verification approach (see section 3), rather than the network-wide one, decomposition at the functional level is achieved, in the sense that individual network element performance measures do not need to be further summarized, in higher management layers; network performance is assessed individually on a s-d basis. By extending the measurement summarization functions to include probability calculations, as described in section 3, the calculation of the required performance measures (see (1), (1a) in section 3) can indeed take place at the network element (management interface) level, incurring no polling cost to the management system whatsoever.

Specifically, when the Performance Verification OSF requests a particular performance measure, a monitoring activity is created in either the Current Load Model OSF or the Connection Performance Model OSF. In turn, these later OSFs delegate element level monitoring activities to the network element (management interface) layer (in an element level OSF, MF, QAF or NEF, according to the actual implementation or the type of elements to be monitored). The monitoring activities and data retrieval between the Performance Verification OSF and the Current Load Model/Connection Performance Model OSFs is achieved by the creation of monitoring activities in the form of metric and summarization objects. By creating thresholds to identify degraded performance conditions, and by using event and alarm reporting, asynchronous, rather than synchronous polling communications are achieved, reducing both the communications overhead, and the processing overhead required in the Performance Verification OSF. In turn, the interaction between the Current Load Model/Connection Performance Model OSFs and the underlying OSFs/MFs/QAFs/NEFs is achieved by the same mechanisms, allowing asynchronous communications on exception and therefore reducing the communications load between the Network Management Layer and the Network Element Management Layer. Figure 4 illustrates the proposed design approach.

Only if the Network Elements themselves do not support the required SMFs is synchronous polling required between the lowest level management functions and the elements themselves. This means that high load communications inherent in polling mechanisms is limited to local area communications, reducing the load in the rest of the TMN and the underlying managed network. So increasing the capacity for revenue earning traffic.

By distributing the functions required over the Network Management and Network Element Management Layers and by utilizing the powerful framework of OSI management, the architecture of the TMN hierarchical system can be designed to avoid the management communications overhead inherent in centralized systems by pushing management intelligence and frequently used management functions as close as possible to the network elements.

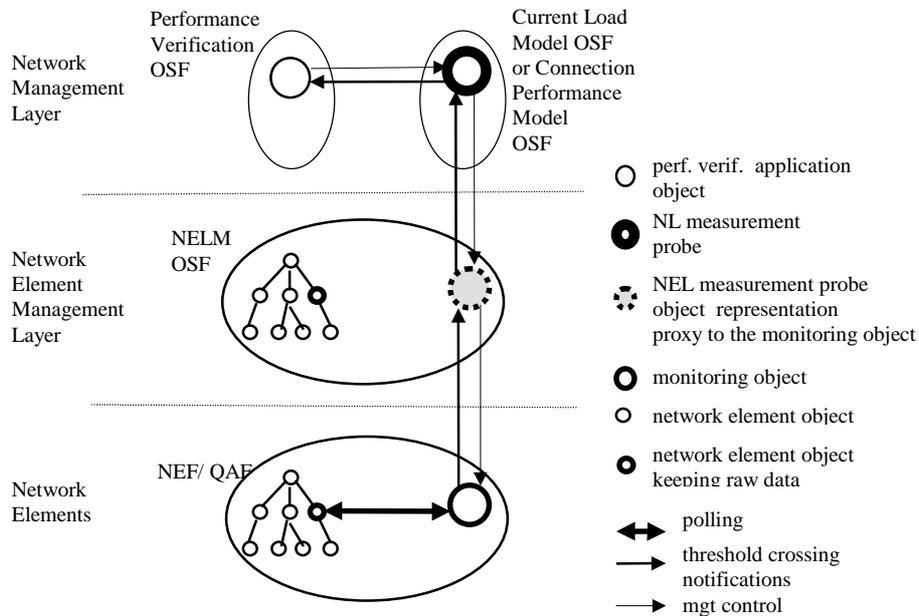


Figure 4: Performance Verification management system design.

6. Conclusions and future work

In this paper we dealt with the issue of network performance verification in broadband multi-service network. Recognizing the emerging need for enhanced automated network management systems, the paper concentrated on the problem of network performance verification within the overall context of network management.

A network management approach (management service definition, algorithms and architecture) was followed based on TMN and OSI management principles. The paper defined an appropriate management service, the Performance Verification management service, for network performance assessment. The Performance Verification management service was described in the context of the performance management area, as well as its interface with management components in other management functional areas, the network and the TMN users.

The introduced Performance Verification management service is beneficial to the network operation in an indirect manner in the sense that it is beneficial to the management system that runs on the top of the network. Specifically, the Performance Verification management service quantifies the performance of the network management system, acting as an indisputable measure of the efficiency of the deployed management systems.

The paper elaborated on the parameters characterizing network performance and on how monitoring of these parameters can be achieved within the following constraints:

- management overhead to be introduced for collecting the required network information
- validity and sensitivity of the required measurements.

The functionality of the Performance Verification management service was analyzed and a specific approach was proposed. A TMN-compliant management system architecture fulfilling the objectives of the Performance Verification management service was presented. By taking advantage of the rich and powerful features of OSI network management, the paper proposed an efficient design approach for implementing the identified functionality. The proposed design approach, although covers all s-d pairs and CoSs introduces minimum overhead into the management system. Specifically, by pushing the monitoring functionality down to the network elements, polling from the network management layer is avoided; Performance Verification, residing in the network management layer, receives only the emitted threshold crossing notifications, indicating unacceptable network performance.

Requirements on supporting monitoring objects to cover the required performance monitoring needs were drawn and enhancements on the 'classical' OSI metric and summarization monitoring objects were identified.

A Performance Verification system has been prototyped following the architecture and design principles presented in this paper.

Future work, includes extensive experimentation for quantifying the performance of the Performance Verification management system design in terms of the overhead introduced to the network management system. Relative comparisons with alternative network performance verification approaches and designs, is another important dimension of future work. Other aspects of future work include further research for enhancing the identified functional components e.g. the customer complaint analysis functionality. Certain aspects of this work are currently being undertaken in the RACE II ICM project.

Acknowledgments

The work described in this paper has been carried out by the authors in the course of the Integrated Communications Management (ICM) project (R2059) in the framework of the RACE II programme. The RACE II programme is partially funded by the Commission of the European Union. The authors wish to acknowledge Peter Baxendale (of University of Durham), Andy Carr (of Cray Communications) Kostas Kassapakis (of ALPHA Systems S.A.), George Mykoniatis (of National Technical University of Athens) and Bruno Rossi (of Monetel/ASCOM) for implementing the proposed management system components.

References

1. K.Geiths, P.Francois, D.Griffin, C.Kaas-Petersen, A.Mann "*Service and traffic management for IBCN*", IBM Systems Journal, Vol.31, No.4, 1992.
2. D.Griffin, P.Georgatsos "*A TMN system for VPC and routing management in ATM networks*", Integrated Network Management IV, Proc. of 4rth intern. symposium on integrated network management, 1995, ed. Adarshpal S. Sethi, Yves Raynaud and Fabienne Faure-Vincent, Chapman & Hall, UK, 1995.

3. P.Georgatsos, D. Griffin, "*Load Balancing in Broadband Multi-Service Networks: A Management Perspective*", Proceedings of the Third Workshop on Performance Modelling and Evaluation of ATM Networks, July 1995
4. M.Wernic, O.Aboul-Magd, H.Gilbert "*Traffic management for B-ISDN Services*", IEEE Network, Sept.1992.
5. G.Woodruff, R.Kositpaiboon "*Multimedia traffic management principles for Guaranteed ATM Network Performance*", IEEE J. Select. Areas of Comm., Vol.8, No.3, July 1992.
6. ITU-T Recommendation I.320 "*ISDN protocol reference model*"
7. ITU-T Recommendation I.321 "*B-ISDN protocol reference model and its application*"
8. ITU-T Recommendation I.150 "*B-ISDN asynchronous transfer mode functional characteristics*"
9. ITU-T Recommendation M.3010 "*Principles of a telecommunications management network*"
10. ITU-T Recommendation M.3020 "*TMN interface specification methodology*"
11. ATM Forum, "ATM User-Network Interface Specification", Version 3, Sept.1990.
12. ITU-T Recommendation X.734, "*Event Report Management Function*"(ISO/IEC 10164-5)
13. ITU-T Recommendation X.733, "*Alarm Reporting Function*" (ISO/IEC 10164-4)
14. ITU-T Recommendation X.735, "*Log Control Function*" (ISO/IEC 10164-6)
15. ITU-T Recommendation X.745, "*Test Management Function*" (ISO/IEC 10164-12)
16. ITU-T Recommendation X.739, "*Workload Monitoring Function*" (ISO/IEC 10164-11)
17. ITU-T Recommendation X.738, "*Measurement Summarization Function*" (ISO/IEC 10164-13)