Evaluating PNNI Performance

Sílvia Farraposo[†], Edmundo Monteiro[‡]
silvia@estg.iplei.pt, edmundo@dei.uc.pt

Laboratory of Communications and Telematics
Centre for Informatics and Systems of the University of Coimbra
Polo II, Pinhal de Marrocos, 3030 Coimbra, Portugal

Responsible Author

Escola Superior de Tecnologia e Gestão de Leiria
Morro do Lena, Alto-Vieiro Apartado 3063
2401-951 Leiria - Portugal
Telf: +351 244820332 Fax: +351 244820310

Keywords

PNNI routing, PNNI signalling, performance evaluation, testbeds.

Abstract

The purpose of this paper is to present some tests and the obtained results to evaluate PNNI performance. The experimentation arisen from the RiaXo project, who is intended to characterise PNNI at WAN environment. In this way, after a brief presentation of the RiaXo project, the scenarios defined and the experiments performed to obtain this goal are also presented.

The analysis of the results, derived from experimentation, can supply directives for an improved testbed configuration to implement at RiaXo environment and form a range of values useful to comparisons LAN versus WAN.

1. Introduction

Eventhough the PNNI 1.0 specification [1] dates from the beginning of 1996, only at the end of 1998 the first commercial versions of this protocol, with all the functionality implemented, were introduced in the market. This fact as induced several studies about ATM technology to include a PNNI evaluation item, as is the case of the RiaXo project [2], proposed by the Portuguese *Fundação para a Computação Científica Nacional* (FCCN) and Universities in collaboration with Portugal Telecom (the public telecommunication operator).

[†] MSc Student at Laboratory of Communication and Telematics, Centre for Informatics and Systems of the University of Coimbra. Responsible author.

[‡] Auxiliary Teacher at Informatics Department of University of Coimbra.

RiaXo project is composed by several workpackages, one of them devoted to dynamic routing at ATM networks, intended to evaluate PNNI behaviour. Nevertheless, due to PNNI being recent there is not enough quantitative information about it and it was necessary to gather some. In this sense, the main concern of this work was to define PNNI performance evaluation in a LAN environment, in order to gather values to compare with WAN experiments.

To evaluate PNNI performance, several testbeds were defined, each one characterised by several factors, like the number of nodes, the number of links, Peers Groups creation and hierarchy routing. Particularly, to study PNNI signalling, it was necessary to inject ATM traffic into each scenario: UBR and CBR traffic.

The content of this paper is distributed as follow: section 2 describes aspects concerning the PNNI routing performance evaluation: objectives, scenarios used and results obtained, section 3 presents PNNI signalling performance evaluation, including the scenarios defined and the items under observation. Finally, section 4 regards to final conclusions and future work.

2. PNNI Routing Performance

Evaluate PNNI routing performance can be very difficult due to its complexity. One possible approach is presented here, and its main concern is the process of convergence necessary to obtain a stable routing structure and external occurrences capable of perturbing it.

2.1 Objectives

The concern of this study about PNNI routing was to qualify and to quantify the process of network convergence. So, the tests were defined to permit:

- Evaluate the time needed to all nodes, to have the same information in their databases, at different scenarios.
- Evaluate the impact at the routing structure, when a fault occurs. The values presented were measured in function of the time the routing structure needs to become again synchronised.

2.2 Defined scenarios

The structures defined were the result of some restrictions in the equipment used and the purpose of the study. In this sense, the scenarios were defined to offer a platform with variable parameters, like the number of nodes, the number of links, the number of peers groups and the number of levels of the hierarchy. Table 1 presents the scenarios defined.

Each one was the result of connecting three types of switches: the Core Builder 7000™ from 3Com™, the ForeRunner LE155™ from Fore™ and the ForeRunner ASX-200BX™ from Fore™. All of them had UTP and OC3/STM1 155 Mbps interfaces and PNNI 1.0 version. Particularly, the Core Builder PNNI version presented some limitations about routing hierarchy and communication between Peers Groups.

Configuration	Description	Scenario
#1	Peer Group with two nodes.	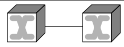
#2	Peer Group with three nodes.	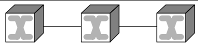
#3	Peer Group with four nodes.	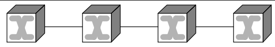
#4	Peer Group with two nodes. There are two alternative paths to reach each endpoint.	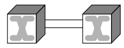
#5	Peer Group with three nodes. There are two alternative paths to reach each endpoint.	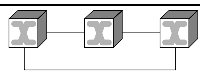
#6	Peer Group with four nodes. There are two alternative paths to reach each endpoint.	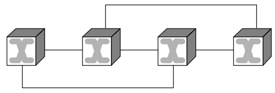
#7	Two nodes distributed into two Peers Groups.	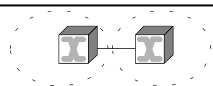
#8	Three nodes distributed into two Peers Groups.	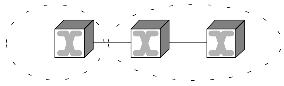
#9	Four nodes distributed into two Peers Groups.	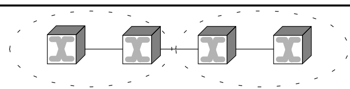
#10	Routing hierarchy with two levels: one physical with three nodes distributed along two Peers Groups and one logical with two nodes belonging to the same Peer Group.	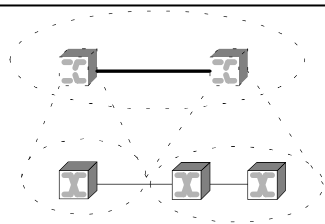

Table 1 Defined scenarios for PNNI routing performance evaluation.

2.3 Evaluating PNNI routing experimentation

When a physical link is established between two switches, a convergence process is initiated. This process stops when both have the same view of the network and each one starts to exchange data periodically. This occurrence was chosen to mark the end of each test.

Routing domain convergence

The stability of a network is essential to PNNI because it depends on it to select the paths for calls end-to-end, whenever they occur. This condition permits to any node to have the same information about the network and resources on it, and to respond equally to the same call demand.

This concern was an argument to perform some experiments to identify factors responsible for changing the stability of a network, and so its performance.

Figure 1 presents the mean times a network needs to converge, in each one of the scenarios presented above. Particularly, Figure 1 presents how the number of nodes, the number of links and the complexity of the routing structure may affect the time needed by the network to reach stability.

The figure show that the number of links and nodes at the routing structure (configurations #1 to #6) do not affect the mean time needed to network convergence. The variations shown in the graphic, less than 1 second, may be due to external factors, like processing time at ATM switches.

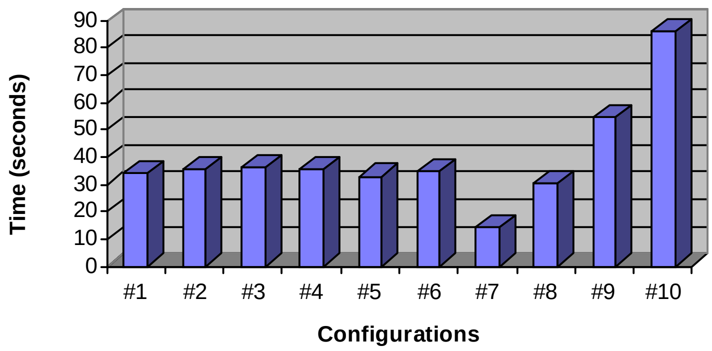

Figure 1 Mean time a network needs to converge.

The values associated to more complex routing structures (configurations #7 to #10) shows that the time necessary to stabilise the network depend on it. In this way, a topology with two nodes, like configuration #7, only need approximately 18 seconds to become stable, meanwhile an hierarchic structure need approximately 85 seconds to become operational.

Changes in the routing structure

A critical factor in performance routing is the way elements in a network react to unpredictable occurrences, like link and node failures. In this sense, in the experiments performed changes were applied to link and node states, and the time needed by the network to recover its stability was evaluated. Only scenarios #6 and #10 presented above were used, since they are the more similar to a real network structure. Table 2 presents the faults introduced during experimentation. All of them permit the continuity of communication, after their occurrence.

Configuration	Description
#6a	Node failures at structure #6.
#6b	Link failures at structure #6.
#10a	Link failure between Peers Groups.
#10b	Peer Group Leader failure in one of the Peers Groups. The Peer Group Leader does not connect to the neighbour group.
#10c	Logical node failure.

Table 2 Failure configurations to evaluate the stability of the routing domain.

Figure 2 show that the occurrence of a node failure is more critical than a link failure, since more time is spent in recovering network stability. Also, from the figure above it is clear that a node or link failure affects more a flat routing structure than a hierarchic routing structure. Particularly, when the routing topology has two or more levels, it seems to need less time to recover from a fault situation, which proves its robustness.

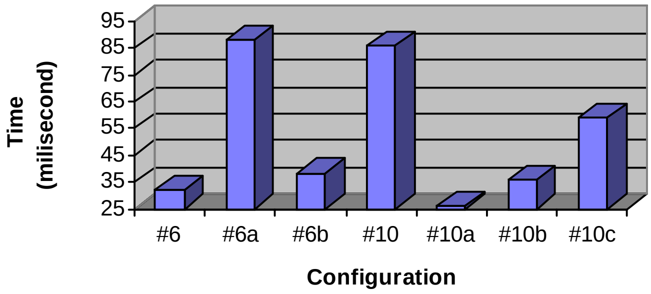

Figure 2 Recovering time after a fault occurrence.

3. Evaluating PNNI signalling performance

Since signalling data does not transport user information, it must be minimised in a routing environment, letting resources free to user information. So, it is important to know what to expect from a routing structure when end-to-end calls are issued. The results presented here focus on call setup and call release.

3.1 Objectives

The experiments performed under this topic were defined to qualify and quantify the setting up and releasing of user calls. Also under evaluation were some factors that could affect setting up and releasing times. The following items made up the experimentation performed:

- Evaluate the relationship between the necessary time to setup a call and the Quality of Service (QoS) premises involved.
- Evaluate the relationship between the necessary time to setup a call and routing structure.
- Evaluate the relationship between the necessary time to release a call and the QoS premises involved.
- Evaluate the relationship between the necessary time to release a call and routing structure.

3.2 Defined scenarios

When evaluating PNNI signalling, the scenarios defined were intended to be like real routing structures: one with several nodes and redundant links inserted in a Peer Group, and another with the nodes distributed over two Peers Groups and a logical level to permit communication between them. Table 3 presents these scenarios.

Configuration	Description	Scenario
#1	The four nodes belong to the same Peer Group. There are always two routes to each node.	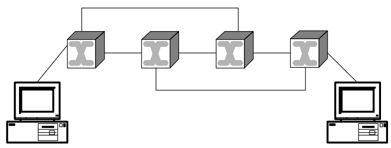
#2	The nodes are distributed over two Peers Groups and a logical level was created to permit communication between them.	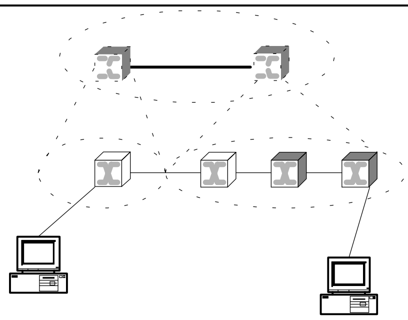

Table 3 Defined scenarios for PNNI signalling performance evaluation.

3.3 PPNI signalling experimentation

The following experimentation was oriented to respond some questions, like how call setup and call release are affected by QoS requirements, routing structure and available resources.

3.3.1 Call setup

At the present experimentation, PNNI signalling messages were generated by the use of two applications [3], responsible to establish end-to-end connections with two different types of contract: Unspecified Bit Rate (UBR) and Constant Bit Rate (CBR). UBR traffic does not have any QoS requirements, while CBR traffic need the attribution of a value to Cell Delay Variation (CDV) and Peak Cell Rate (PCR) parameters.

Type of traffic

The experiments realised under this point were intended to evaluate how the time spent with PNNI signalling is affected by the type of traffic during a call setup. To do so, UBR and CBR end-to-end calls were requested over configurations #1 and #2 presented above.

Figures 3 and 4 presents the average time needed to establish an UBR and CBR dynamic connection (exchange of Setup, Call Proceeding and Connect messages between switches), respectively. Both figures show the values obtained with different data rates.

The observation of both figures suggests that the time associated with PNNI signalling is independent of the data rate requested but dependent of the type of traffic. The constancy of values show up one of the principal advantages of using outband signalling, where a special channel is used to transmit exclusively signalling data.

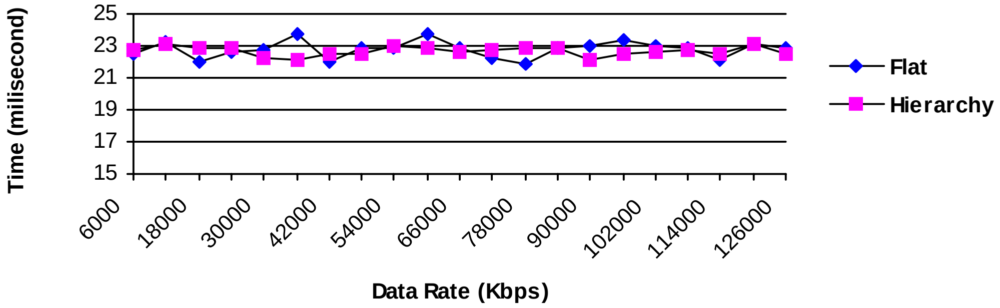

Figure 3 Average time spent with PNNI signalling during an UBR call setup.

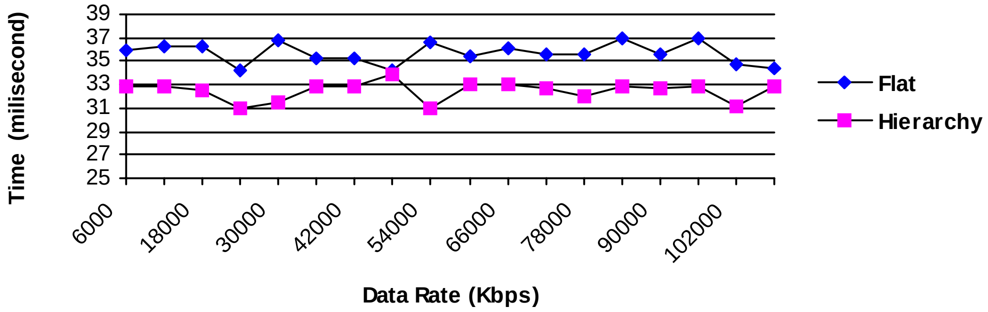

Figure 4 Average time spent with PNNI signalling during a CBR call setup.

The comparison of figures 3 and 4 show that the time spent to establish an UBR call is lesser than the time spent to establish a CBR call. This difference is due to CAC algorithm that must be executed in each node traversed by a CBR call, requiring more computational calculus.

Routing structure

From the two figures presented above, 3 and 4, it is possible to infer about the way time spent in PNNI signalling during call setup varies with the routing topology. In this context, Figure 3 show that structure routing does not affect the establishment of an UBR call. But, when a CBR call is requested, if the structure is flat, it takes about more 10% of time to setup the connection than when the structure has two levels. Although the variation is minimal, it is an aspect to take into account when structuring a network.

Utilisation of resources

Another purpose for this study was to verify if the quantity of connections already established could influence the time required to setting up another call. This is, to verify if signalling requirements are always guarantied.

To get the results only configuration #1 presented in Table 3 was used. Because this scenario has only one Peer Group, it is possible to restrict overhead due to complexity configuration. To establish UBR calls, an average bandwidth utilisation of 65% was considered. Figure 5 present the time spent in a new call setup, when there is already connections established.

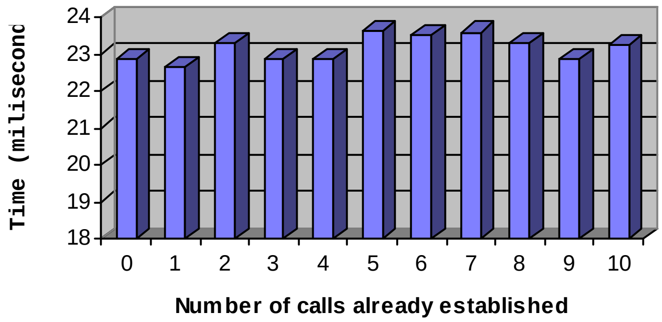

Figure 5 Average time spent with PNNI signalling when establishing a new UBR call.

From the figure above it is possible to see that the establishment of a new UBR connection doe not requires more signalling time even when there is resources allocated to other calls. Another observation was that the network always accepted this type of calls even if the total bandwidth utilisation was near 100%. This observation agrees with ATM Forum definition [1], which says that CAC algorithm cannot reject an UBR call based in available bandwidth, but can reject it based on other motifs, like exceeding a previously established threshold.

Figure 6 presents the results obtained through similar experiments but with CBR traffic. In this case, it was not possible to consider calls with 65% of bandwidth allocation because the network starts to reject new calls when 70% of its bandwidth is occupied. The maximum value permitted depends on Available Cell Rate parameter. So, the values presented refer to CBR call establishment with variable bandwidth requirements.

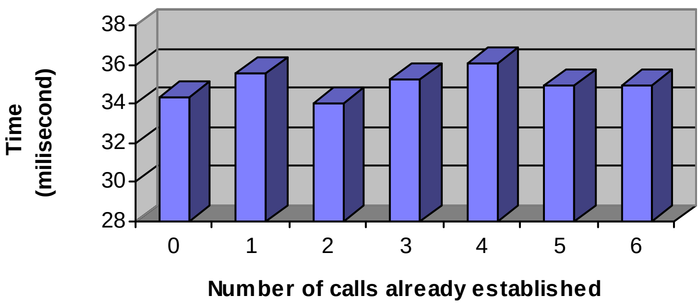

Figure 6 Average time spent with PNNI signalling when establishing a new CBR call.

Like for an UBR call, the figure shows that the time spent with PNNI signalling is always the same and is independent of the number of connections already established. The only restriction verified is about the CAC algorithm responsible for the total number of simultaneous calls over the network. In this way, whenever a user requires a new call, the network must verify if it have enough resources for it, i.e., verify the QoS requirements.

3.3.2 Call release

PNNI signalling associated to call release includes two messages: Release and Release Complete. The experiments under this section evaluated how much time PNNI needed to exchange these messages, whenever an end station pretends to terminate a call. To do this, UBR and CBR variable data rate calls were established over the scenarios presented in Table 3 and subsequently released. Figures 7 and 8 present the values obtained, for UBR and CBR calls, respectively.

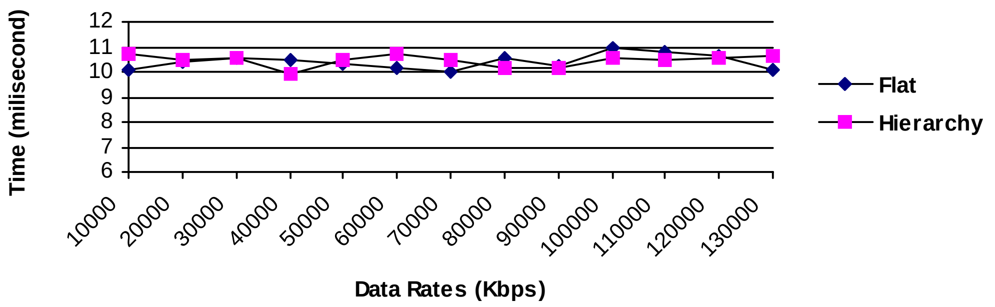

Figure 7 Average time spent with PNNI signalling when releasing an UBR call.

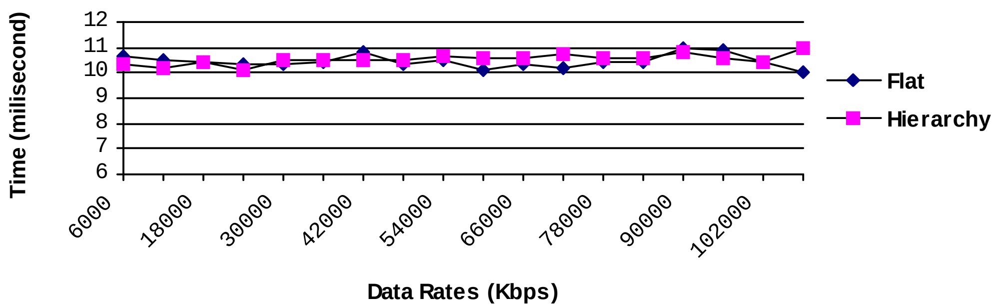

Figure 8 Average time spent with PNNI signalling when releasing a CBR call.

The constancy of values, approximately 10,5 ms, presented above suggests that releasing a call is independent of the type of traffic and the routing structure. These facts was expected, because when releasing a call, each node only needs to forward messages and free its allocated resources, without any additional processing, like the verification of QoS parameters. This behaviour is coherent with signalling outband, where the exchange of signalling data is transparent to user data and is not affected by network problems.

4. Conclusion

If PNNI protocol wants to be efficient whenever a call is established it must guarantee a group of premises to satisfy user requirements and manage resources correctly. Its routing component permits that all nodes in the network possess the same routing information, while its signalling component guarantee a fair usage of resources by the end users.

About PNNI routing, one of the observations was that the time spent in network convergence is minimally affected by the number of links and nodes in the Peer Group, but on the other hand routing domain logical configurations may be a critical factor in convergence time. If distribution of nodes per various Peers Groups does not increase network stability time, the

PGL election process and subsequently the hierarchy creation may increase the convergence time around 50%. Another evidence showed up was routing structure behaviour face to a failure. When a fault occurs it is not transparent to the other domain elements, because all of them must know what happen to readjust their own databases. This mechanism evolves the exchange of additional routing information.

About signalling PNNI it was noted that call setup information is exchanged over an appropriate signalling channel, where the connection parameters are evaluated. Particularly, the experiments showed that call setup is dependent of the type of user traffic, while call release is independent. Both are independent of the data rate requested and the subjacent routing structure. The behaviours presented corroborate the idea that the signalling channel only has peaks of utilisation, having bandwidth utilisation approximately zero after a request completion.

All the experiments were conduced with the objective of defining a baseline of PNNI values, to use during RiaXo experimentation. These tests also served to define a proposal to implement over RiaXo framework.

The next step is to apply all the knowledge gathered in the testing phase of RiaXo project in a way to obtain reliable results from experimentation.

Acknowledgements

This work was supported in part by the Portuguese Ministry of Science and Technology PRAXIS XII Research Program, under contract PRAXIS/P/EEI/10168/1998 (Project QoS II - Quality of Service in Computer Communication Systems).

This was also supported by the Fundação para a Computação Científica Nacional, the Portugal Telecom and the RiaXo project.

References

[1] The ATM Forum Technical Committee, "*Private Network–Network Interface Specification Version 1.0 (PNNI 1.0)*". af–pnni–0055.000, March 1996.

[2] Portugal Telecom, Fundação para a Computação Científica Nacional, "*Descrição Técnica do Projecto Riaxo*". Junho 1999.

[3] Fore Systems, "*ForeRunner™ Adapters Installation Manual, Software version 5.3.x*". Fore Systems Inc., September 1998.