INTEGRATED COMMUNICATIONS MANAGEMENT OF BROADBAND NETWORKS

Crete University Press, Heraklio, Greece
ISBN 960 524 006 8

Edited by David Griffin

First published in 1996 by
CRETE UNIVERSITY PRESS
Foundation for Research and Technology
P.O. Box 1527, Heraklio, Crete, Greece 711 10
Tel: +30 81 394235, Fax: +30 81 394236
email: pek@iesl.forth.gr

The ICM consortium consists of the following companies:

Alcatel ISR, France
Alpha SAI, Greece
Ascom Monetel, France
Ascom Tech, Switzerland
Centro de Estudos de Telecommunicações, Portugal
Cray Communications Ltd., United Kingdom (Prime contractor)
Danish Electronics, Light & Acoustics, Denmark
De Nouvelles Architectures pour les Communications, France
Foundation for Research and Technology - Hellas, Institute of Computer Science, Greece
GN Nettest AS, Denmark
National Technical University of Athens, Greece
Nokia Corporation, Finland
Queen Mary and Westfield College, United Kingdom
Unipro Ltd., United Kingdom
University College London, United Kingdom
University of Durham, United Kingdom
VTT - Technical Research Centre of Finland

Chapter 12

Experimental results

Editors: Dionysis Reisis, Richard Lewis
Authors: Panos Georgatsos, Peter Baxedale, Dionysis Reisis

This chapter presents the results of testing and evaluating the ICM TMN functional implementation, as the concluding phase of the project. This phase is important for two reasons: it validates the functionality of the developed system, and it shows the improvement in network performance resulting from operation of the TMN.

This chapter is organised as follows: Section 12.1 provides the rationale behind the tests and Section 12.2 introduces the evaluation procedures. Section 12.3 summarises the testing performed on real networks. Section 12.4 analyses in detail the tests and evaluation of the VPCM system applied on the ATM simulator. Section 12.5 draws conclusions from the experimental work performed.

Owing to space constraints only a subset of the tests performed is described here. Further details can be found in ICM Deliverable 21 [12.3].

12.1 Rationale

ICM's research activities included development of an experimental TMN Testbed, and concluded with evaluation of the TMN functional implementation. This phase is important in many ways. The evaluation results are useful to the people involved in the project: researchers can quantify the completion of their work; they can also identify the elements of the system which can be improved to lead to more successful realisation of their design. People external to the project will be able to understand the scope of the utilisation of the system; the proof of the proper system behaviour will encourage future use of the Testbed beyond the lifetime of the ICM project; measurements of the elements' performance can lead to the development of improved algorithms.

Experience from testing and redesigning can be of great importance to industry and to research institutes, in the future. An overall evaluation of the Testbed will show whether the architecture and the features of the developed system constitute the basis of future industrial products. Evaluation of a system consists mainly of showing that the system performance has been validated. Validation includes: design, execution and analysis of experiments in order to show that the product has been developed according to its specifications.

The evaluation process must be efficient. In this context, the term efficiency has a particular meaning. It requires the number of tests to be minimised, while retaining confidence that the functionality and the performance of the system can be accurately assessed. This process is quite complex, possibly justifying a project in itself, and requires an overall testing strategy to be followed. This must be thoroughly designed to lead to successful validation of the implementation.

Validation of the ICM Testbed functional implementation is essentially a system testing exercise. Traditionally, this task requires a significant proportion of the resources of a project. A methodical approach and detailed plan are therefore required. Testing is performed by an independent team which performs Black Box Testing, i.e. the test team is not required to know how the system works (the internal processes), the test team is only required to know what the system does (how it reacts to given stimuli). The test team is however, supported by system developers.

The test process is broken down into a number of test classes. A complete testing approach is defined, although it is not expected that all tasks are applicable to the Testbed. Indeed, they may not all be possible within the constraints of a research project. However, it is considered to be important to identify the tests and consciously decide their applicability, rather than risk the omission of key steps in the testing process. The test plan addresses the following subjects:

- Test Procedures (e.g. version control, fault reporting, conduct of tests),
- Test Classes (e.g. case study tests, functional tests, performance tests),
- Test Completion Criteria,
- Test Documentation,
- Test Results.

12.2 Test considerations

When all the Testbed components are operating together on a network, the conditions under which the tests are conducted can affect the results. The following paragraphs identify some of the conditions which may affect the test results.

12.2.1 Management applications

The Testbed implementation process can be assessed by the ease with which different applications can be introduced and tested on it. Clearly, as the effort in the project is finite, the number of applications produced will be, too. Thus each application produced must be tested as rigorously as possible.

From the Testbed viewpoint, a measure of its usefulness is the effort required to introduce new applications, so the manpower required should be carefully monitored. Measures should include not only the raw manpower required, but also an assessment of the experience of the people involved. Is it the first application, or the tenth, that they have attempted to implement and integrate?

12.2.2 Network traffic and load

For each application, there will clearly be a need to characterise its performance in a range of network and TMN configurations, under a range of different traffic mixes with different loads. Considering the network, if we have N traffic types from M source nodes, transmitting to M-1 destinations, we have an N*M*(M-1) dimensional problem space for a given topology, with each parameter varying from zero to some defined maximum and that is without considering the internal network states, such as VPC topologies and capacities.

Problems can occur due to the fact that realistic traffic cannot be directly scaled, but may be subject to quantification. For instance, a VCC carrying one video channel can not be scaled by 10% in a meaningful way, but 500 voice calls can.

One way to minimise the impact of traffic variations is to define typical type mixes, that is, the proportion of voice, video, etc. that are representative for a given topology, and then to scale all traffic by some factor. This corresponds to multiplying all the traffic matrix terms by a constant factor. This has the effect of removing N dimensions from the problem, while still giving a wide range coverage for the results. There are problems in terms of defining one scalable matrix, and there may be a case for more than one, but even if several are used, the problem space is still simplified.

12.2.3 Topologies

One could attempt to use a very large number (in the case of the simulator) of possible network topologies to test algorithms and management systems, but with limited resources, a small representative set is required.

Based on experience of currently installed network technologies, the following types of topology are included in the test scenarios:

1. A fully interconnected network. This is the type of network that is used between major national cities within a country, where it is assumed that traffic will be of a generally uniform nature. Much of the uniformity will be a result of the volume of traffic.
2. A fully interconnected main trunk with peripheral nodes. This type of network takes account of the minor nodes connected to a main distribution network, usually with multiple connection for resilience.

3. Sparse network with one or more main nodes. This type of network is typical of private networks where the major nodes are usually the companies' head offices or computer centres, depending on the traffic types in questions.
4. Interconnected island. This type of network is representative of the interconnections between countries or major centres of population separated by significant distances. Within a centre, normally well or fully interconnected, outbound traffic to a remote destination is funnelled onto a few lines.

12.2.4 TMN configuration

The TMN performance for a given network and loading condition should be assessed in several configurations. These should include the minimum practical number of processors, as well as a highly distributed one, if test resources allow. In considering TMN configurations, it should be borne in mind that resilience will be important in real networks, while processing power is rarely a major financial consideration today.

12.3 Experimental results

ICM focused on testing and evaluating its TMN functional implementation: initially the basic infrastructure functions of the TMN were tested on LAN FDDI network technology, followed by testing and evaluation of a more complex management system - Virtual Path Connection Management (VPCM). In parallel with the VPCM evaluation, ICM implemented and tested the TMN functionality on a Passive Optical Network (PON). Finally, the project concluded its work with the evaluation of the TMNs within an Intermediate Virtual Private Network scenario. In the paragraphs below we summarise the tests and evaluations of each of the above applications. This is followed by detailed descriptions of two of the major tests performed on the VPCM system - Load Balancing (LB) and Bandwidth Distribution (BD).

12.3.1 FDDI

The TMN platform and the fundamental management applications (e.g. monitoring) form the core of the management system and must be subject to tests before relying on them for complex management applications. During these tests the capabilities of the basic TMN modules were evaluated in terms of functionality, validity and performance. The main components were: the network including SNMP agents at each node; Q-adaptors; a Network Element Level monitoring OS per network node; a Network Level statistics OS; and a WS-OS.

The main steps followed for the functional testing were:

- test each component separately, to ensure it works properly and correctly (capability),
- test the integration of all components with the WS-OS(behaviour),
- test the reaction of the system to abnormal situations such as network unavailability (resolution tests).

For the first of the above steps the statistics generated by the TMN were validated against raw data collected directly from the network. The functional test showed that the TMN behaved as expected apart from some minor bugs met in abnormal situations. The TMN components exhibited the majority of the required features for these experiments, with the exception of some minor ones that would give more flexibility to the manager (presentation options, the ability to change the moving window size, event report post-processing). These requirements were fed back into the design an implementation areas of the project for enhancement in the later prototypes.

Regarding system performance, response times were measured for each TMN component under different loads. The main conclusion was that the measured response times were acceptable - comparable with a commercially available SNMP system. The management communications overhead was measured in terms of the bandwidth required by the TMN and it was demonstrated that the traffic generated by the TMN system does not impose a significant load in FDDI networks. Finally the performance tests showed that the monitoring functions of the TMN can follow the traffic characteristics only when properly configured. Parameters such as the EWMA and the sampling intervals significantly affect the quality of the monitored data.

The functional tests demonstrated that the network monitoring system functioned as intended. The performance tests measured how well the TMN achieves its objectives in terms of processing speed and communications resources used. Details are presented in ICM Deliverable 12 [12.1].

12.3.2 VPCM

The VPCM TMN system was tested in real and simulated environments. The system includes OSs which execute sophisticated algorithms, e.g. Load Balancing and Bandwidth Distribution. During operation, these OSs interact with the other OSs such as the Configuration Manager, Route Design and Performance Verification, in order to improve network performance.

The LB and BD experimental results and evaluation have been chosen to be presented in Section 12.4 in detail, as their tests implicitly exercise the other OSs.

12.3.3 PON

In parallel with the VPCM system testing process, ICM evaluated the TMN system applied to the COBRA Passive Optical Network, with the following results:

- The evaluation has shown that it is possible to reuse the ICM TMN components to develop a TMN system applied to the management of a network technology with different management requirements to those of ATM networks.
- Conformance to standards simplifies the task of managing heterogeneous networks, demonstrated by adoption of SNMP for management within the PON, with a Q-Adaptor providing the interface to the TMN.
- With respect to implementation and deployment, the tools developed within the project reduced the development time and maximised the reliability of the resulting system.

Finally, we can say that this phase was successfully developed and validated, and demonstrated the strengths of the ICM TMN implementation methods and tools, and how these technologies can be applied in the future.

12.3.4 I-VPN

The goal of the final experiment was to show that the TMN concepts, as developed by the project, are applicable to inter-domain management and interoperability. The experiments were conducted in a Pan-European environment consisting of multiple management domains of different network and service providers. This phase tested the iVPN implementation, using an extended platform which includes support for security (access control and authentication) and inter-domain management. The iVPN system was demonstrated in a real Pan-European environment between the ETB and RIA networks located in Switzerland and Portugal. In addition, it was shown that the ICM simulator provided support for more complex experiments involving CPNs as well as multiple interconnected networks.

12.4 VPCM experimental results

This section presents the functional validation tests results of the Load Balancing OS and the Bandwidth Distribution OS. The rationale for specifying the tests, the specific test description and the results of the tests are presented. The tests were designed to be carried out using the ICM ATM simulator.

12.4.1 Load Balancing application tests

The purpose of the Load Balancing (LB) application test is to test and verify the functionality (intelligence) of the LB OS in a network environment. The test is performed at a system level with the necessary components of the ICM testbed. The test is carried out in a variable traffic environment with the purpose to test and verify the functionality of the LB OS in all possible traffic conditions in which it might be activated.

It should be noted that the purpose of the test is not the detailed assessment of the performance of the LB OS, in terms of benefits and costs to the network operation, sensitivity, management overhead, scaling and processing time. This test aims at validating the functionality of the LB OS. Performance assessment is regarded as a huge task and it is not addressed by this work. However, the validation tests may provide indications on the performance of the LB OS.

12.4.1.1 Description of test cases

12.4.1.1.1 Analysis

The essence of the test, is to invoke the LB OS at instances where the network is at a given state and subsequently to observe the produced actions. These instances should cover all (most of) the possible cases in which LB could be activated. The definition of representative test cases that partition the network state space at LB invocation instances is the theme of this section. We reason as follows:

The functionality of LB depends on:

- the network CoSs and their bandwidth characteristics;
- the routing plan;
- the network load.

Therefore, the test cases should cover all possible values of the above parameters; an infinite number of combinations.

However, in order to be able to test the functionality of LB, for each test case we should be able to know the expected outcome i.e. the actions that LB would produce. The functionality of LB will be tested against the expected actions. For this reason, the test cases should not be chaotic, but controlled and comprehensive. The approach taken is the following:

We focus on a tagged CoS and a tagged source destination (s-d) pair. This is done in order to reduce unnecessary complexity. We further focus only on the LB actions arriving at the switches of the tagged SDClassRoute network, rather than observing the LB actions arriving at all switches in the network. For the tagged pair of the tagged CoS, we distinguish test cases, conditioning on the previous parameters as follows:

Test cases regarding network CoSs

We condition on the peak bandwidth of the CoSs. Hence, we partition the network CoSs into three categories according to whether their peak bandwidth is substantially larger, almost the same or substantially smaller than the peak bandwidth of our tagged CoS.

Note that the LB functionality does not depend on whether the network CoSs are CBR or VBR. Therefore, we assume the simplest type, the CBR type. Moreover note that the LB functionality does not depend on the actual number, given that is greater than one, of the CoSs in each category; the results would be similar, if we replaced the number of CoSs, say nc, of a particular category of peak bandwidth pb, with a new category with a single CoS, with peak bandwidth the product of nc and pb. Therefore, the following cases can be considered with respect to the network CoSs:

cases	smaller category CoS	same category CoS	larger category CoS
C000	NO	NO	NO
C100	YES	NO	NO
C010	NO	YES	NO
C001	NO	NO	YES
C110	YES	YES	NO
C101	YES	NO	YES
C011	NO	YES	YES
C111	YES	YES	YES

Table 12.1 Test cases for the network CoSs

Test cases regarding routing plan

The routing plan denotes the set of routes for each (s-d) pair and each network CoS. The functionality of LB depends on the following aspects of the routing plan:

- the topology of the routes,
- the multiplexing of the CoSs on the routes.

For the tagged (s-d) pair and network CoS, we consider a given set of routes provided that the following topology constraints are met:

- there are physically different routes as well as routes with common parts,
- there are routes with varying number of hops (VPCs up to the destination node),
- there is alternatibility not only at the access node but also at the transient nodes at every hop level.

It is believed that with the above constraints the SDClassRoute networks are representative of the route networks that will exist in the future networks. Given the SDClass-Route network of the tagged pair and CoS, and a given set of CoSs, corresponding to the test cases of the previous table (Table 12.1), further cases can be therefore considered, conditioning on the following parameters:

- For the cases with one extra CoS:
  - route alternatibility per CoS (1,>1).
- For the cases with two extra CoSs:
  - route alternatibility per CoS (1, >1),
  - route sharing (no, partial, total).
- For the case with three extra CoSs:
  - route alternatibility per CoS (1,>1),
  - route sharing per two CoSs (no, partial, total),
  - route sharing per three CoSs (no, partial, total).

In particular, from the many cases that can be found, the cases shown in Table 12.2 to Table 12.4 may be considered.

cases	alternatibility
TP:1.1	1
TP:1.2	2

Table 12.2 Test cases for the routing plan: One extra CoS case

cases	alternatibility		route sharing
TP:2.1.1.N	1	1	NO
TP:2.1.1.T	1	1	TOTAL
TP:2.2.2.P	2	2	PARTIAL

Table 12.3 Test cases for the routing plan: Two extra CoS case

cases	alternatibility			route sharing per 2	route sharing per 3
TP:1.1.1.N.N	1	1	1	PARTIAL	PARTIAL
TP:1.1.1.T.T	1	1	1	TOTAL	TOTAL
TP:2.1.2.P.N	2	1	2	PARTIAL	PARTIAL

Table 12.4 Test cases for the routing plan: Three extra CoS case

Test cases regarding network load

The network load is considered only at the VPCs that exist in the SDClassRoute network of the tagged pair and CoS. Further test cases are considered by conditioning the network load on the following two parameters:

- symmetry (symmetric, asymmetric),
- volume (low (L), medium (M), high (H)).

The asymmetric load cases may be partitioned conditioning on the hop levels where the asymmetry occurs. Hop level 1 denotes the VPCs that start from the access node. Hop level I=2.. denotes the VPCs that can be reached from hop level I-1. Note that a VPC may belong to consecutive hop levels. Asymmetry in hop level I means that the hop level I VPCs are asymmetrically loaded. Asymmetry may occur at any level. From the many combinations of symmetric and asymmetric load, the following may be selected:

cases	asymmetry			
Symmetric				
LD:S;l	NO			
LD:S;m	NO			
LD:S;h	NO			
Asymmetric	level 1	level 2	level 3	level 4
LD:A;1	YES	NO	NO	NO
LD:A.2	NO	YES	NO	NO
LD:A.3	NO	NO	YES	NO
LD:A.1.2	YES	YES	NO	NO
LD:A.1.3	YES	NO	YES	NO
LDA.1.2.3	YES	YES	YES	NO
LD:A.2.3	NO	YES	YES	NO
LD:A.3.4	NO	NO	YES	YES
LD:A.4	NO	NO	NO	YES

Table 12.5 Test cases for network load

12.4.1.1.2 Synthesis

So far the rationale for defining the test cases for validating the functionality of the LB OS have been described. These test cases will be running in the ICM testbed using the ATM network simulator. It should be noted that the use of the simulator makes feasible and cost effective the application of the defined test cases.

For each of the cases C:xxxx (see Table 12.1) and each of the cases TP:xxxxx appropriate networks are defined. The networks differ in the number of CoSs and in the set of routes for each CoS. The test cases LD:xxxx are generated by creating appropriate traffic profiles and user groups using these profiles. With these profiles, the load deviations corresponding to the defined LD:xxxx cases will be created.

For each experiment, the invocation times of the LB OS are specified together with the expected actions at each invocation instant. By comparing the actions taken with those expected at each invocation instant, the functionality of the LB OS is verified.

The configuration of the ICM testbed for the LB application test is as follows:

- SimMMI,
- Simulation engine,
- Simulator's QA,
- NL Configuration Manager OS,
- Current Load Model OS,
- Route Design OS,
- Load Balancing OS,
- TMN WS-OS.

12.4.1.2 Description of network experiments

12.4.1.2.1 Network experiment A

Physical topology

The physical topology is shown in Figure 12.1.

Connection and service classes

connection types	type	mean bandwidth (Mb/s)	peak bandwidth (Mb/s)
c1	CBR	5	5
c2	CBR	10	10
c3	CBR	2	2

Table 12.6 Supported connection types for experiment A

service classes	no of connection types	connection type
srv1	1 (forward)	c1
srv2	1 (forward)	c2
srv3	1 (forward)	c3

Table 12.7 Offered services for experiment A

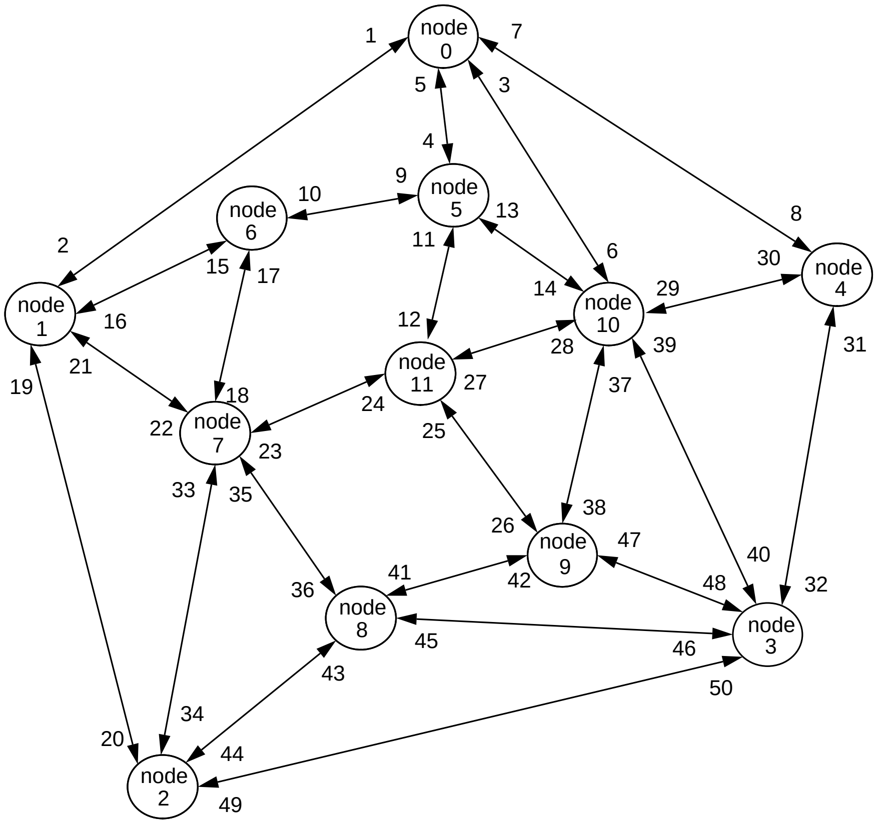

Figure 12.1 Reference network physical topology

Logical topology

VPC id	start node	destination node	links	capacity (Mb/s)
100	1	7	22	50
101	7	8	36	50
102	8	3	46	50
103	8	3	42,48	50
104	1	11	15,9,12	50
105	11	3	28,40	50
106	11	3	26,48	50
107	7	11	24	50
200	7	1	21	50
201	8	7	35	50
202	3	8	45	50

Table 12.8 VPCs for experiment A

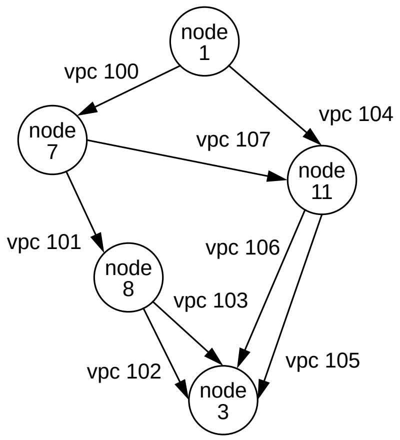

Figure 12.2 SDClassRoute network for (1-3) (s-d) pair and CoSs c1, c2, c3

VPC id	start node	destination node	links	capacity (Mb/s)
203	3	8	47, 41	50
204	11	1	11, 10, 16	50
205	3	11	39, 27	50
206	3	11	47, 25	50
207	11	7	23	50

Table 12.8 VPCs for experiment A

route	destination node	connection types	VPCs
R1	3	c1, c3	100, 101, 102
R2	3	c1	100, 101, 103
R3	3	c1	104, 106
R4	3	c1, c2	104, 105
R5	3	c1	100, 107, 105
R6	3	c1	100, 107, 106

Table 12.9 Routes for experiment A

	destination node	connection type	VPC	priority
node 1	3	c1	100	2
	3	c1	104	1
node 7	3	c1	101	2
	3	c1	107	1
node 8	3	c1	102	2

Table 12.10 Initial route selection priorities for experiment A

	destination node	connection type	VPC	priority
	3	c1	103	1
node 11	3	c1	105	2
	3	c1	106	1

Table 12.10 Initial route selection priorities for experiment A

Traffic definition

service profiles	service	start time	end time	call time	silence time
srv1-p1	srv1	900	11000	100	1000
srv3-p1	srv3	1000	11000	8000	1000
srv3-p2	srv3	1000	10000	6000	2000
srv3-p3	srv3	1000	9000	4000	3000
srv2-p1	srv2	4000	11000	5000	1000
srv2-p2	srv2	4000	11000	3000	2000
srv2-p3	srv2	4000	9000	1000	3000

Table 12.11 Service profiles for experiment A

start node	destination node	SRV	number of users
1	3	srv1-p1	1
1	3	srv2-p1	2
1	3	srv2-p2	1
1	3	srv2-p3	1
1	3	srv3-p1	10
1	3	srv3-p2	5
1	3	srv3-p3	5

Table 12.12 User groups for experiment A

Expected results

The LB OS will be invoked at the following times: 2100, 3100, 4100, 5100, 6100, 7100, 8100, 9100, 10100. The connection type c1 is our tagged network CoS and the (1-3) pair is our tagged (s-d) pair.

Figure 12.3 and Figure 12.4 show the number of active service calls over time for the (s-d) pair between which traffic has been defined.

The expected load is shown in Table 12.13.

invocation instant	Load on VPCs along Route R1	Load on VPCs along Route R4
1	L	no load
2	M	no load
3	H	no load

Table 12.13 Expected load per invocation instant for experiment A

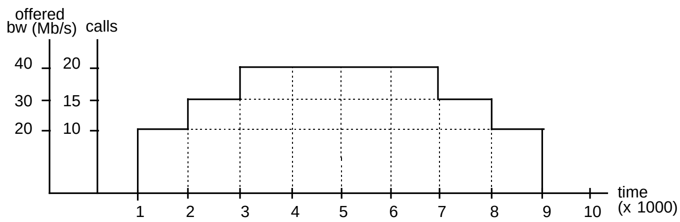

Figure 12.3 Active calls of c3 for the (1-3) pair

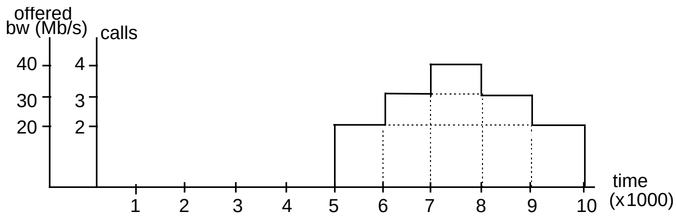

Figure 12.4 Active calls of c2 for the (1-3) pair

invocation instant	**Load on VPCs along Route R1**	**Load on VPCs along Route R4**
4	H	L
5	H	M
6	H	H
7	M	M
8	L	L
9	no load	no load

Table 12.13 Expected load per invocation instant for experiment A

L: low load (utilisation 20/50)   NB: VPCs vpc103 and vpc106 will have no load
M: medium load (utilisation 30/50)
H: high load (utilisation 40/50)

The experiments are concerned with traffic of different CoSs between a single (s-d) pair. The SDClassRoute networks have alternatibility 1 (for the other CoSs) or 2 (for the tagged CoS) in hop levels 1,2 and 3. With these test suites the functionality of LB, particularly regarding to the total route sharing and symmetric or asymmetric load (in hop levels 1, 2 and 3) test aspects, will be verified. The expected results regarding route selection priority changes for the routes of the tagged CoS, c1, per invocation instant are described next.

At invocation instant 1, the VPCs along route R1 are lightly loaded, whereas the other VPCs carry no load. Since the initial route selection priorities for CoS c1, in nodes 1, 7 and 8 suggest selection of the loaded VPCs, it is expected that LB will change the route selection priorities to suggest the selection of the non-loaded VPCs. Specifically, in node 1, vpc104 will be given a higher priority than vpc100 which is loaded, hence, a change priority action is expected. In node 7, the choice is between vpc107 (non-loaded) and vpc101 (loaded). Obviously, a priority change is expected in this node in favour of vpc107. Similarly, in node 8, LB is expected to change priorities so that to suggest VPC vpc103 (non-loaded), rather than VPC vpc102 (loaded). In node 11, both vpc106 and vpc105 are non-loaded. VPC105 has been initially set with a higher priority. However, vpc105 is to be shared by CoSs c1 and c3, whereas VPC vpc106 is to exclusively accommodate traffic from CoS c1. Therefore, it is expected LB to change priorities in favour of vpc106, as the LB functionality has been designed to take into account the alternatibility of the different CoSs, protecting CoSs with less alternatibility over CoSs with higher alternatibility.

The load at invocation instances 2 and 3 is similar with the load in the previous invocation instances but the load on the VPCs along route R1 has increased to medium (instant 2) and to high (instant 3). Obviously, the non-loaded VPCs should be suggested for routing. Since this selection preference is already reflected with the current route selection priorities, it is expected LB not to determine any action.

At invocation instances 4 and 5, the load on the VPCs along route R1 is the same as in the previous instant and the VPCs along route R4 have light (instant 4) or medium (instant 5) load. At node 1, both vpc100 and vpc104 are loaded but vpc100 has a greater load than vpc104, hence vpc104 should be preferred. Vpc104 already has higher priority than vpc100 (from the first invocation), therefore no action is expected. In the other nodes, there is a choice between loaded and non-loaded VPCs. Since the non-loaded VPCs have been suggested for routing from the first invocation instant, it is expected LB will not take any action.

At invocation instant 6, the VPCs along routes R1 and R4 are all highly loaded. At node 1, both vpc100 and vpc104 have the same load; vpc100 is shared between CoSs c1 and c2 and vpc104 is shared by CoSs c1 and c3. However, the bandwidth requirements of CoS c2 are higher than the bandwidth requirements of CoS c3. The design of the LB algorithm, in addition to the load and alternatibility criteria, favours routes that are shared by lower bandwidth CoSs than routes shared by higher bandwidth CoSs. Therefore, since vpc104 had higher priority than vpc100, it is expected LB to change priorities so that to favour vpc100. In the other nodes, since there is a choice between a non-loaded VPC and a loaded VPC, the non-loaded VPC should be preferred. Given that the priorities as set by the first invocation instant prioritise these non-loaded VPCs, no LB action is expected.

Invocation instances 7, 8 and 9 are similar to the previous invocation instant but the load on the routes has been decreased to medium (instant 7), light (instant 8) and no load at all (instant 9) respectively. Following the arguments presented for the previous instant, the priorities should remain unchanged. Therefore, no LB action is expected.

The expected LB actions are summarised in Table 12.14.

invocation instances	node 1 vpc priority	node 7 vpc priority	node 8 vpc priority	node 11 vpc priority
1	104 2* 100 1*	107 2* 101 1*	103 2* 102 1*	106 2* 105 1*
2	as previous	as previous	as previous	as previous
3	as previous	as previous	as previous	as previous
4	as previous	as previous	as previous	as previous
5	as previous	as previous	as previous	as previous
6	100 2* 104 1*	as previous	as previous	as previous
7	as previous	as previous	as previous	as previous
8	as previous	as previous	as previous	as previous
9	as previous	as previous	as previous	as previous

Table 12.14 Expected actions for experiment A

Results

The results per invocation instant as produced by the actual invocation of the LB OS are presented below:

```
#file: lb_log.dat
#Experiment: lbapltstexp1a-
###invocation Instant: 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_T-EU7, CoSId 1, Dest EU_1_A-
EU3, VPC vpc101, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_T-EU7, CoSId 1, Dest EU_1_A-
EU3, VPC vpc107, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_T-EU8, CoSId 1, Dest EU_1_A-
EU3, VPC vpc102, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_T-EU8, CoSId 1, Dest EU_1_A-
EU3, VPC vpc103, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc106, Prio 3 OldPrio 1
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc105, Prio 1 OldPrio 3
###invocation Instant: 2
###invocation Instant: 3
###invocation Instant: 4
###invocation Instant: 5
###invocation Instant: 6
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 1 OldPrio 2
###invocation Instant: 7
###invocation Instant: 8
###invocation Instant: 9
```

As it can be seen (see Table 12.14), LB functions as expected.

12.4.1.2.2 Network experiment B

Physical topology

As in experiment A (see Figure 12.1).

Connection and service classes

As in experiment A.

Logical topology

VPC id	start node	destination node	links	capacity (Mb/s)
100	1	7	22	50
101	7	8	36	50
102	8	3	46	50
103	8	3	42, 48	50
104	1	11	15, 9, 12	50
105	11	3	28, 40	50
106	11	3	26, 48	50
107	7	11	24	50
110	2	8	43	50
111	2	7	33	50
113	8	4	42, 37, 30	50
200	7	1	21	50
201	8	7	35	50
202	3	8	45	50
203	3	8	47,41	50
204	11	1	11, 10, 16	50
205	3	11	39, 27	50
206	3	11	47, 25	50
207	11	7	23	50
210	8	2	44	50
211	7	2	34	50
213	4	8	29, 38, 41	50

Table 12.15 VPCs for experiment B

route	destination node	CoS	VPCs
R1	3	c1	100, 101, 102
R2	3	c1, c2	104, 106
R3	3	c1	104, 207, 101, 102
R4	3	c3	110, 102
R5	4	c3	111, 101, 113

Table 12.16 Routes for experiment B

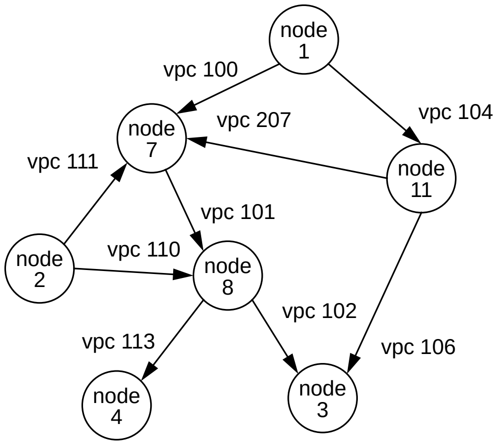

Figure 12.5 S-DClassRoute networks for (1-3), (2-3), (2-4) (s-d) pairs and CoSs c1, c2, c3

	destination node	connection type	VPC	priority
node 1	3	c1	100	2
	3	c1	104	1
node 11	3	c1	207	2
	3	c1	106	1

Table 12.17 Initial route selection priorities for experiment B

Traffic definition

service profiles	service	start time	end time	call time	silence time
srv1-p1	srv1	900	11000	100	1000
srv3-p1	srv3	1000	3500	1000	1000
srv3-p2	srv3	1000	7000	3000	2000
srv2-p1	srv2	1000	7000	1000	3000
srv2-p2	srv2	3000	7000	2000	1000

Table 12.18 Service profiles for experiment B

start node	destination node	SRV	number of users
1	3	srv2-p1	1
1	3	srv2-p2	3
2	3	srv3-p1	15
2	4	srv3-p2	15

Table 12.19 User groups for experiment B

Expected results

The LB OS will be invoked at the following times: 2100, 3100, 4100, 5100, 6100. The connection type c1 is our tagged network CoS and the (1-3) pair is our tagged (s-d) pair.

The following figures show the number of active service calls in time for the (s-d) pairs between which traffic has been defined.

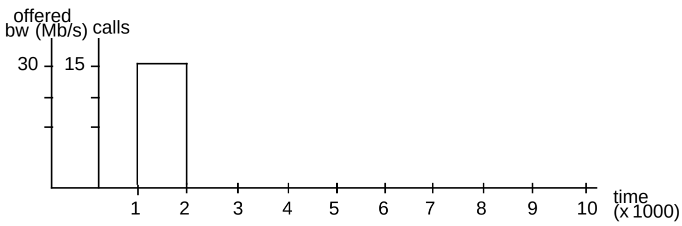

Figure 12.6 Active calls of c3 for the (2-3) pair

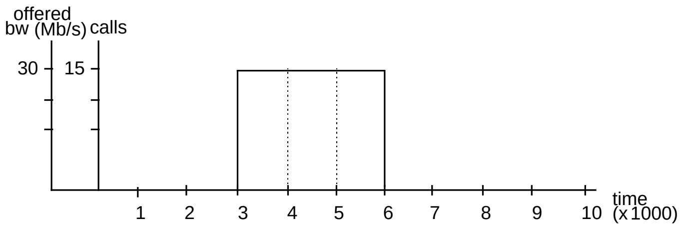

Figure 12.7 Active calls of c3 for the (2-4) pair

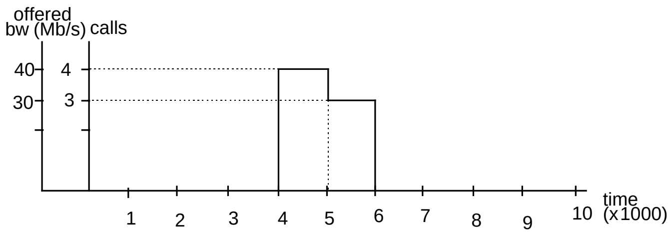

Figure 12.8 Active calls of c2 for the (1-3) pair

The expected load on network resources is shown Table 12.20.

invocation instances	load on VPCs along route R2	load on VPCs along route R4	load on VPCs along route R5
1	no load	M	no load
2	no load	no load	M
3	H	no load	M
4	M	no load	M
5	no load	no load	no load

Table 12.20 Expected load per invocation instant for experiment B

L: low load (utilisation 20/50)
M: medium load (utilisation 30/50)
H: high load (utilisation 40/50)

The experiments are concerned with traffic of different CoSs between multiple source destination pairs. The SDClassRoute networks have alternatibility 1 (for the other CoSs) or 2 (for the tagged CoS) in hop levels 1,2. With these test suites the functionality of LB, particularly regarding to the total and partial route sharing and asymmetric load (between hop levels 1, 2 and 3) test aspects, will be verified. The expected results regarding route selection priority changes for the routes of the tagged CoS, c1, per invocation instant are described next.

At invocation instant 1, vpc102 which belongs to routes R1 and R3 of the tagged CoS is medium loaded. The priorities in the nodes 1 and 11 have been initially set to suggest routing over these routes; therefore, LB when invoked should take actions so that to suggest routing over route R2. In particular the following LB actions are expected: in node 1 vpc104 should be selected against vpc100, since congestion exists in a remote area of the routes starting with vpc100. For the same reason, in node 11 vpc106 should be selected against vpc 207.

Invocation instant 2 is similar to the previous invocation instant, except that vpc101 and not vpc102 is medium loaded. VPC101 also belongs to routes R1 and R3 of the tagged CoS. Since in the previous invocation instant, LB changed the priorities so that to suggest routing over the VPCs of route R2, no LB action is expected at this instant.

In invocation instant 3, vpc101 (belonging to the routes R1 and R3) is still remains medium loaded but vpc104 and vpc106 (which both belong to route R2) are now highly loaded. The priorities from the previous instant have been set to suggest routing along route R2. Since LB has been designed so that to suggest routing over less loaded routes as well as over routes which are shared with lower bandwidth CoSs, it is expected LB to changes the priorities so that to avoid route R2. In particular, the following actions are expected: in node 1 vpc100 should be selected against vpc104 since vpc104 leads to significantly more congested network areas. Moreover, vpc104 is shared by another CoS, whereas vpc100 is dedicated to routing the tagged CoS traffic. For the same reason, in node 11 vpc106 should be selected against vpc207.

Invocation instant 4 is similar to the previous one, but now the VPCs have the same volume of traffic. Since LB has been designed to favour routes that are shared by lower bandwidth CoSs, routing over route R2 should be avoided -route R2 is completely shared by CoS c2 requiring 10 Mb/s, whereas routes R1 and R3 are partially shared by CoS c3 requiring 2 Mb/s. Therefore, no LB action is expected since routing over routes R1 and R2 has been suggested from the previous invocation instant.

In invocation instant 5, all VPCs are not loaded. Following similar arguments with the previous instant, no LB action is expected.

The expected LB actions are summarised in Table 12.21.

invocation instances	node 1 vpc priority	node 11 vpc priority
1	104 2* 100 1*	106 2* 207 1*
2	as previous	as previous
3	100 2* 104 1*	106 1* 207 2*
4	as previous	as previous
5	as previous	as previous

Table 12.21 Expected actions for experiment B

Results

The results per invocation instant as produced by the actual invocation of the LB OS are presented below:

```
#file: lb_log.dat
#Experiment: lbapltstexp3-
###invocation Instant: 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc106, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc207, Prio 1 OldPrio 2
###invocation Instant: 2
###invocation Instant: 3
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc106, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_T-EU11, CoSId 1, Dest EU_1_A-
EU3, VPC vpc207, Prio 2 OldPrio 1
###invocation Instant: 4
###invocation Instant: 5
```

As it can be seen (see Table 12.21), LB functions as expected.

12.4.1.2.3 Network experiment C

Physical topology

The same as in experiment A (see Figure 12.1).

Connection and service classes

As in experiment A.

Logical topology

VPC id	start node	destination node	links	capacity (Mb/s)
100	1	7	22	50
101	7	8	36	50
102	8	3	46	50
103	8	3	42, 48	50
104	1	11	15, 9, 12	50
105	11	3	28, 40	50
106	11	3	26, 48	50
110	2	8	43	50
200	7	1	21	50
201	8	7	35	50
202	3	8	45	50
203	3	8	47, 41	50
204	11	1	11, 10, 16	50
205	3	11	39, 27	50
206	3	11	47, 25	50
210	8	2	44	50

Table 12.22 VPCs for experiment C

route	destination node	CoS	VPCs
R1	3	c1	100, 101, 102
R2	3	c1	104, 106
R3	3	c3	110, 102

Table 12.23 Routes for experiment C

	destination node	connection type	VPC	priority
node 1	3	c1	100	2
	3	c1	104	1

Table 12.24 Initial route selection priorities for experiment C

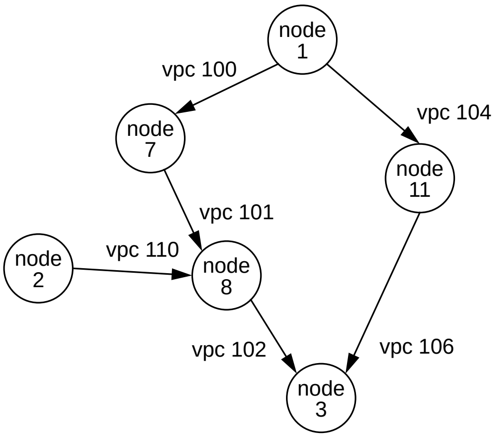

Figure 12.9 S-DClassRoute networks for (1-3), (2-3) (s-d) pairs and CoSs c1, c3

Traffic definition

service profiles	service	start time	end time	call time	silence time
srv1-p1	srv1	600	2600	800	800
srv1-p2	srv1	800	2600	400	1000
srv3-p1	srv3	600	2600	1200	400

Table 12.25 Service profiles for experiment C

start node	destination node	service	number of users
1	3	srv1-p1	4
1	3	srv1-p2	4
2	3	srv3-p1	10

Table 12.26 User groups for experiment C

Expected results

The LB OSF will be invoked at the following times: 1050, 1450, 1850, 2250. The connection type c1 is our tagged network CoS and the (1-3) pair is our tagged (s-d) pair.

Figures 12.10 and 12.11 show the number of active service calls in time for the (s-d) pairs between which traffic has been defined.

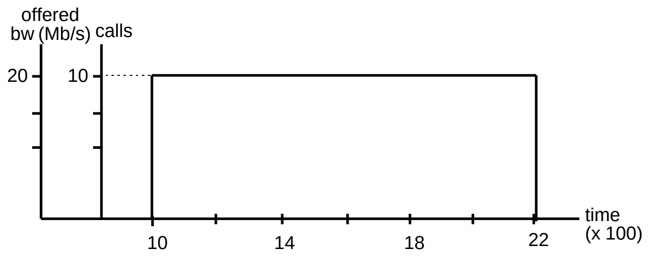

Figure 12.10 Active calls of c3 for the (2-3) pair

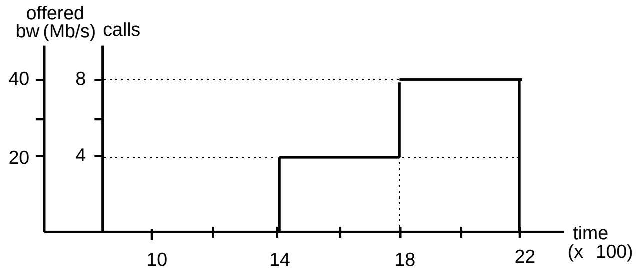

Figure 12.11 Active calls of c1 for the (1-3) pair

The expected load on network resources is shown in the table below.

invocation instances	load on VPCs along route R3	offered load of tagged CoS generated by S-D 1-3
1	L	no load
2	L	L
3	L	H
4	no load	no load

Table 12.27 Expected load per invocation instant for experiment C

L: low load (utilisation 20/50)
M: medium load (utilisation 30/50)
H: high load (utilisation 40/50)

The experiments are concerned with traffic of different CoSs between multiple source destination pairs. The SDClassRoute networks have alternatibility 1 (for the other CoSs) or 2 (for the tagged CoS) in hop level 1. With these test suites the functionality of LB, particularly regarding to the partial route sharing and asymmetric load (between hop levels 1, 2 and 3) test aspects, will be verified. The expected results regarding route selection priority changes for the routes of the tagged CoS, c1, per invocation instant are described next.

At invocation instant 1, only vpc102, which belongs to route R1 of the tagged CoS, is lightly loaded. Initially the priorities have been set to suggest routing over route R1.

Since LB has been designed to favour the least congested routes, it is expected to take actions so that to avoid route R1 which uses vpc102. In particular, it is expected to change priorities in node 1 so that to favour vpc104 rather than vpc100 which leads to more congested network areas.

At invocation instant 2, vpc102 is still lightly loaded. A part of route R1 is also loaded with the same amount of load; due to load of the tagged CoS, offered by the (1-3) pair, which followed route R2, recommended from the previous invocation instant. However, route R1 is partially used by another CoS, c3, of another (s-d) pair, which uses that route with alternatibility 1. Since LB favours routes which are exclusively dedicated to a particular CoS at the expense of routes that are shared by other CoSs, it is expected to favour routing over route R2. Since routing over route R2 has been recommended from the previous invocation instant, no LB action is expected to occur.

At invocation instant 3, the load in route R2 has increased to high. Although route R2 is not shared by another CoS whereas route R2 is, the load difference between these two routes is expected to prevail, and therefore LB should change the priorities so that to suggest routing over route R1 (the least congested ones). Specifically, in node 1, the route selection priorities are expected to change so that to favour vpc100 (leading to the least congested network area, even if is to create congestion to the route of another CoS of another (s-d) pair), than vpc104 (which is already overloaded).

At invocation instant 4, the load on route R2 is still high. In addition, the load on vpc102, part of route R1, is also high. From the previous invocation instant route R1 has been set higher priorities that route R1. But because route R1 is to be used by another CoS of another (s-d) pair, whereas route R2 is not, according to the LB design, it is expected the priorities to change so that to suggest routing over route R2. Specifically, in node 1, the priorities will change so that to favour vpc104 (leading to the dedicated route R2), rather than vpc100 (leading to a route that affects routes of other (s-d) pair CoSs).

The expected LB actions are summarised in Table 12.28.

invocation instances	node 1 vpc priority
1	104 2* 100 1*
2	as previous
3	100 2* 104 1*
4	104 2* 100 1*

Table 12.28 Expected actions for experiment C

Results

The results per invocation instant as produced by the actual invocation of the LB OS are presented below:

```
#file: lb_log.dat
#Experiment: lbapltstexp3-
###invocation Instant: 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 2 OldPrio 1
###invocation Instant: 2
###invocation Instant: 3
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 2 OldPrio 1
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 1 OldPrio 2
###invocation Instant: 4
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc100, Prio 1 OldPrio 2
Sending Update NRT start node EU_1_A-EU1, CoSId 1, Dest EU_1_A-
EU3, VPC vpc104, Prio 2 OldPrio 1
```

As it can be seen (see Table 12.28), LB functions as expected.

12.4.1.3 Conclusions on the Load Balancing capability tests

The functionality of LB has been validated in a number of test suites defined in the course of different experiments. The test suites considered, cover a significant number of representative test cases. The rationale behind the definition of representative test cases partitioning the space of network states in which the LB OS might be invoked, has also been presented. Specifically, the test suites considered covered different cases regarding route alternatibility in different hop levels, different cases regarding CoSs that may use the routes and different cases regarding the supported and different cases regarding route sharing (total, partial).

The functionality of LB has been validated by checking its outcome against the expected outcome in each test suite.

The results indicate that the LB OS indeed functions correctly according to the design specifications of its route selection priority management functionality. In particular, it was shown that for a particular network CoS, LB takes actions so that routes combining the following merits are favoured:

- less load,
- higher alternatibility factor regarding the occasions under which the different CoSs may be using them,
- sharing by CoSs with lower total bandwidth requirements.

which indeed constitute the main criteria for prioritising selection of alternative routes.

Conclusively, it is believed that LB OS indeed according to its design specifications.

12.4.2 LB behavioural tests

12.4.2.1 Overview

The purpose of the LB behavioural tests is to verify the performance of the LB OS in the network by assessing the effect of its actions on network operation.

The effect of LB on network operation is captured in terms of:

- connection rejection ratio, and,
- deviation of links' utilisation around mean link utilisation.

For a given network configuration, the performance of LB is assessed, by relatively comparing network-wide connection rejection ratio, and/or link utilisation deviation as measured in the cases where the network employs a TMN system with the LB OS and where the network does not employ a TMN system.

Overall assessment of the performance of LB can then be achieved by consolidating the results of its performance per network case, over a number of network configurations.

Following the test case decomposition principles presented for the LB functional tests specification, the different network configurations may be partitioned in terms of the following parameters:

- logical topology configuration; and specifically in terms of:
  - route alternatibility (1, 2,..),
  - route sharing among different CoSs (complete, partial, none),
- offered network load (uniform among all source-destination pairs, asymmetric).

The performance of the LB OS is assessed in both the static and dynamic aspects of its functionality. The static aspect of its functionality is related to the initial assignment of the route selection priorities, for the routes that have been defined by the Route Design OS. The static part of LB functionality may be performed at predicted usage prediction changes epochs. Initial route selection priority assignment does not take into account the dynamics of actual network resource usage, since the newly derived routes may be defined in terms of newly created VPCs on which no usage record exists as yet. The dynamic aspect of its functionality is related to the dynamic assignment of the route selection priorities, according to the actual network load. This part of LB functionality is performed within the time-frame of network usage predictions, taking into account the dynamics of network resource usage. In this case, the performance of LB is assessed as a function of its main operational design variable; the activation interval. It is reminded that the current design of the LB OS assumes that the LB OS will be invoked by the TMN operator, either periodically, every specified time-intervals, or at exception.

It is worth noting that the LB behavioural tests show how the developed ICM TMN testbed can be used to assess the performance of various management functions. Obviously, a prerequisite for such tests is the management functionality under test to comply with the Q3 interfaces between the components of the VPCM management system. For instance, the behavioural tests used for the assessing the behaviour of the developed LB OS could well be used for assessing the performance of any other OS complying with the functional and architectural specifications of the LB OSF.

The following section presents the network configurations used and specific experiments carried out, for assessing the performance of the developed LB OS.

12.4.2.2 Network experiments description

12.4.2.2.1 Static case: initial assignment of route selection priorities

The network configurations in which the static aspects of the LB functionality will be tested are described in the following. All test networks are based on the same VPC network and support the same CoSs. They differ in terms of routing plan (admissible routes per each CoS) and in terms of traffic generation patterns. As far as the routing plan is concerned, the following cases have been considered: complete and partial route sharing amongst CoSs of the same source-destination pairs. In either case the alternatibility in the SDClassRoute networks is of the order of 1, 2 and 3. As far as the traffic generation patterns are concerned, three cases have been considered: asymmetric traffic generation amongst the traffic sources, in terms of time-epochs where the load is offered to the network, symmetric traffic generation and bursty traffic generation patterns. The test networks run for various combinations of route selection priority settings including the settings resulted from the LB OS and by relatively comparing the connection rejection ratios per CoS as produced per each priority setting case, the effect of LB is assessed.

The logical topology (VPC network) of the test networks used is shown in Figure 12.12. All VPCs have capacity of 100 Mb/s. Table 12.29 shows the supported CoSs.

CoS	Type	Mean, Peak bandwidth (Mb/s)
c1	CBR	5
c2	CBR	10
c3	CBR	2

Table 12.29 Supported CoSs

Based on the logical topology shown in Figure 12.12, two network cases are defined, each having different sets of routes for the supported CoSs. In the first case, network case N1, all CoSs share the same routes; complete route sharing. The defined routes are shown in Table 12.30; for a given source-destination pair the defined routes are all possible routes that can be derived from the VPC network (Figure 12.12). The alternatibility in all SDClassRoute networks is of the order of 2 and 3. In the second case, network case N2, there is partial route sharing amongst the CoSs. CoS c1 is admitted to routes comprising the VPCs connecting the access nodes to the transit nodes and the VPCs along the circuit connecting the transit nodes but which do not lie in the circuit connecting the transit nodes; CoS c2 is admitted to routes comprising the VPCs connecting the access and transit nodes and the VPCs connecting the transit nodes; CoS c3 is sharing all routes of CoSs c1 and c2. The defined routes are shown in Table 12.31. The

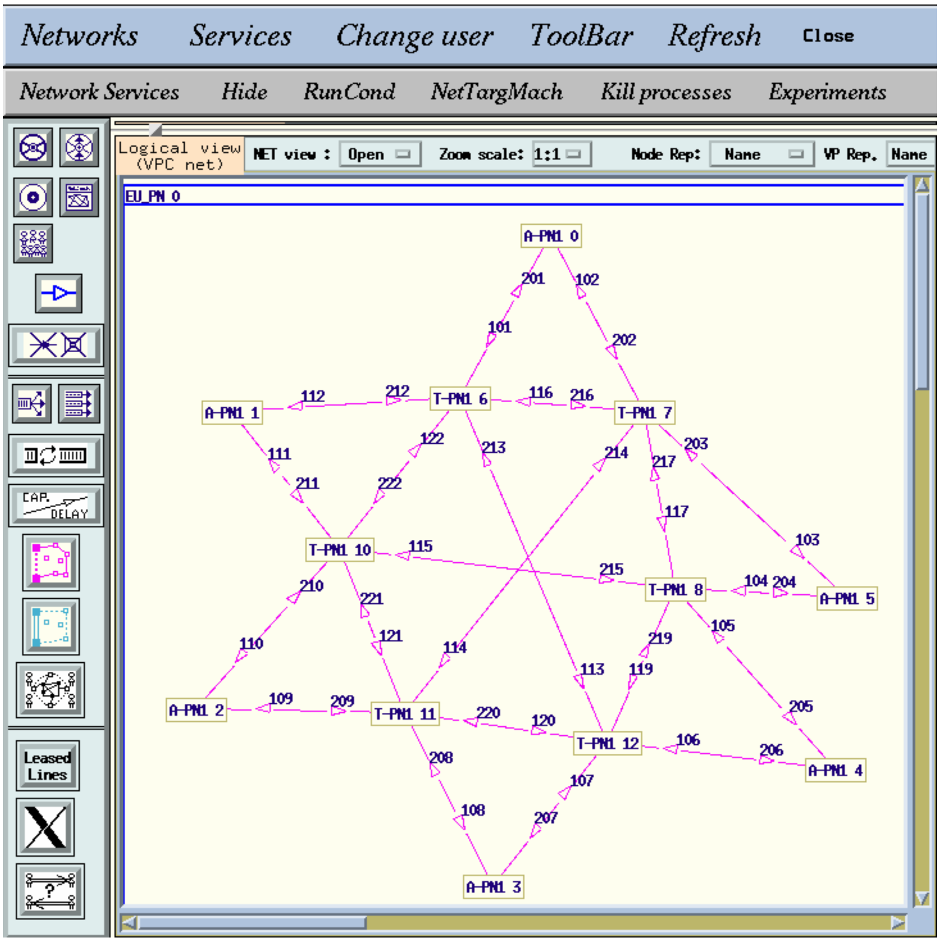

Figure 12.12 VPC network for static case

alternatibility in the SDClassRoute networks is in the order of 1 and 2 for CoSs c1 and c2 and in the order of 1, 2 and 3 for CoS c3.

Routes	destination node	CoS	VPCs
R1	3	c1, c2, c3	101, 222, 121, 108
R2	3	c1, c2, c3	101, 113, 207
R3	3	c1, c2, c3	202, 117, 119, 207
R4	3	c1, c2, c3	101, 216, 117, 119, 207
R5	4	c1, c2, c3	211, 121, 120, 206
R6	4	c1, c2, c3	211, 122, 216, 117, 205
R7	4	c1, c2, c3	212, 216, 117, 205
R8	4	c1, c2, c3	211, 122, 113, 206

Table 12.30 Route definitions for network case N1

Routes	destination node	CoS	VPCs
R9	5	c1, c2, c3	210, 122, 216, 103
R10	5	c1, c2, c3	209, 120, 219, 204
R11	5	c1, c2, c3	210, 122, 113, 219, 204
R12	5	c1, c2, c3	209, 221, 122, 113, 219, 204
R13	0	c1, c2, c3	107, 219, 217, 102
R14	0	c1, c2, c3	107, 213, 201
R15	0	c1, c2, c3	208, 120, 213, 201
R16	0	c1, c2, c3	208, 221, 122, 201
R17	1	c1, c2, c3	106, 220, 221, 111
R18	1	c1, c2, c3	106, 213, 112
R19	1	c1, c2, c3	105, 119, 213, 112
R20	1	c1, c2, c3	105, 217, 116, 112
R21	2	c1, c2, c3	104, 119, 213, 222, 110
R22	2	c1, c2, c3	104, 119, 220, 109
R23	2	c1, c2, c3	203, 117, 119, 213, 222, 110
R24	2	c1, c2, c3	203, 116, 222, 110
R25	3	c1, c2, c3	202, 114, 108
R26	3	c1, c2, c3	202, 117, 115, 121, 108
R27	0	c1, c2, c3	208, 214, 102
R28	0	c1, c2, c3	208, 221, 215, 217, 102
R29	4	c1, c2, c3	211, 215, 205
R30	4	c1, c2, c3	212, 216, 114, 120, 206
R31	1	c1, c2, c3	106, 220, 214, 116, 112
R32	1	c1, c2, c3	105, 115, 111
R33	5	c1, c2, c3	210, 215, 204
R34	5	c1, c2, c3	209, 214, 103
R35	2	c1, c2, c3	203, 114, 109
R36	2	c1, c2, c3	104, 115, 110

Table 12.30 Route definitions for network case N1

Routes	destination node	CoS	VPCs
R1	3	c1, c3	101, 222, 121, 108
R2	3	c2, c3	101, 113, 207
R3	3	c1, c3	202, 117, 119, 207
R4	3	c1,c3	101, 216, 117, 119, 207
R5	4	c1, c3	211, 121, 120, 206
R6	4	c1, c3	211, 122, 216, 117, 205
R7	4	c1, c3	212, 216, 117, 205
R8	4	c2, c3	211, 122, 113, 206
R9	5	c1, c3	210, 122, 216, 103
R10	5	c1, c3	209, 120, 219, 204
R11	5	c2, c3	210, 122, 113, 219, 204

Table 12.31 Route definition for network case N2

Routes	destination node	CoS	VPCs
R12	5	c2, c3	209, 221, 122, 113, 219, 204
R13	0	c1, c3	107, 219, 217, 102
R14	0	c2, c3	107, 213, 201
R15	0	c2, c3	208, 120, 213, 201
R16	0	c1, c3	208, 221, 122, 201
R17	1	c1, c3	106, 220, 221, 111
R18	1	c2, c3	106, 213, 112
R19	1	c2, c3	105, 119, 213, 112
R20	1	c1, c3	105, 217, 116, 112
R21	2	c2, c3	104, 119, 213, 222, 110
R22	2	c1, c3	104, 119, 220, 109
R23	2	c2, c3	203, 117, 119, 213, 222, 110
R24	2	c1, c3	203, 116, 222, 110
R25	3	c2, c3	202, 114, 108
R26	3	c2, c3	202, 117, 115, 121, 108
R27	0	c2, c3	208, 214, 102
R28	0	c2, c3	208, 221, 215, 217, 102
R29	4	c2, c3	211, 215, 205
R30	4	c2, c3	212, 216, 114, 120, 206
R31	1	c2, c3	106, 220, 214, 116, 112
R32	1	c2, c3	105, 115, 111
R33	5	c2, c3	210, 215, 204
R34	5	c2, c3	209, 214, 103
R35	2	c2, c3	203, 114, 109
R36	2	c2, c3	104, 115, 110

Table 12.31 Route definition for network case N2

For each of the above network cases, traffic is generated between the following source-destination (s-d) node pairs: node0-node3, node3-node0, node1-node4, node4-node1, node2-node5 and node5-node2. Specifically, for a given (s-d) pair, traffic is generated according to the following pattern:

$$T_{(s\text{-}d),\,c}(k) = A + 200 * U_{(s\text{-}d)} * CR_c * k,\ c=c1,c2,c3,\ k=1,2,3. \qquad (1)$$

where:

A is an initial time period required for the TMN system to initialise the network simulator in terms of VPCs and routes.

$U_{(s\text{-}d)}$ is a uniformly distributed random variable in a specified interval which is used to differentiate traffic generation between different (s-d) pairs.Three cases are considered:

- $U_{(s\text{-}d)}$ uniformly distributed in (0,2),
- $U_{(s\text{-}d)}$ uniformly distributed in (1,2),
- $U_{(s\text{-}d)}$ uniformly distributed in (0.9,1.1),

with each case corresponding to a different network case.

CR_c is a variable differentiating traffic generation between different CoSs for a given (s-d) pair. It takes the following values: $CR_{c1} = 1$; $CR_{c2} = 0.9$; $CR_{c3} = 1.1$.

k index corresponding to the traffic generation instances per (s-d).

All (s-d) pairs generate the same number of connection requests at each generation instance; the time of generation differs amongst the (s-d)s and the CoSs (according to formula (1) above). Table 12.32 shows the number of connection requests per CoS and per generation instance. It is assumed that each successful connection call lasts for the whole duration of the simulation run.

Generation instance	# c1 (CBR 5Mb/s) reqs	# c2 (CBR 10 Mb/s) reqs	#c3 (CBR 2 Mb/s) reqs
k = 1	6	4	15
k = 2	5	3	12
k = 3	8	4	20

Table 12.32 Connection requests per generation instance and per CoS

Table 12.33 summarises the different networks defined for the LB static tests which differ in terms of the sets of routes per CoS and the traffic generation pattern.

Networks	Logical topology	Traffic generation pattern
N1-A	VPCs: as in Figure 12.12 Routes: as in Table 12.30	as in formula (1) with Us-d~(0,2) and in Table 12.32
N1-B	VPCs: as in Figure 12.12 Routes: as in Table 12.30	as in formula (1) with Us-d~(0,2) and in Table 12.32
N1-C	VPCs: as in Figure 12.12 Routes: as in Table 12.30	as in formula (1) with Us-d~(0,2) and in Table 12.32
N2-A	VPCs: as in Figure 12.12 Routes: as in Table 12.31	as in formula (1) with Us-d~(0,2) and in Table 12.32
N2-B	VPCs: as in Figure 12.12 Routes: as in Table 12.31	as in formula (1) with Us-d~(0,2) and in Table 12.32
N2-C	VPCs: as in Figure 12.12 Routes: as in Table 12.31	as in formula (1) with Us-d~(0,2) and in Table 12.32

Table 12.33 Different network cases for the LB static tests

Each of the above networks will run under various route selection priority settings. Specifically, apart from the route selection priority settings resulted from the LB OS, the following four other settings were considered:

1st: all nodes favour the interior VPCs (the VPCs that connect the transit nodes but which do not lie in the circuit connecting the transit nodes) and then the VPC of the shortest.

2nd: all nodes prioritise VPC with the exact reverse priorities of the previous case.

3rd: half the transit nodes taken per 3, prioritise VPCs according to the 1st priority setting and half according to the 2nd priority setting.

4th: half the transit nodes taken one-by-one prioritise VPCs according to the 1st priority setting and half according to the 2nd priority setting.

For each of the above networks the results of the connection rejection ratios per CoS for each route selection priority setting can be found in Section 12.4.2.3.1.

12.4.2.2.2 Dynamic case: dynamic assignment of route selection priorities

The network configurations in which the dynamic aspects of the LB functionality will be tested are described in the following. Two different network logical topologies are considered, covering different cases regarding route alternatibility and partial route sharing amongst different source-destination pairs. Furthermore, different traffic generation (symmetric or asymmetric) patterns amongst the traffic sources are considered. Each of the test networks run with or without the LB management system and furthermore in the former case under various LB invocation periods. By relatively comparing the produced results for the connection rejection ratios per network case, the LB performance was verified and assessed.

Figure 12.13 and Figure 12.14 show the two different VPC topologies used in the test network cases. Tables 12.34 and 12.35 show the bandwidth of the VPCs and tables 12.36 and 12.37 show the defined routes for each network case. The networks support one CoS which is CBR with required bandwidth of 1Mb/s.

VPCs	bandwidth (Mb/s)
100	60
200	60
101	100
201	100
103	100
203	100
104	60
204	60
105	120
205	120
106	120
206	120

Table 12.34 VPC bandwidth of VPC network case A

VPCs	bandwidth (Mb/s)
100	30
200	30
101	80
201	80
103	80
203	80
104	60
204	60
105	80

Table 12.35 VPC bandwidth of VPC network case B

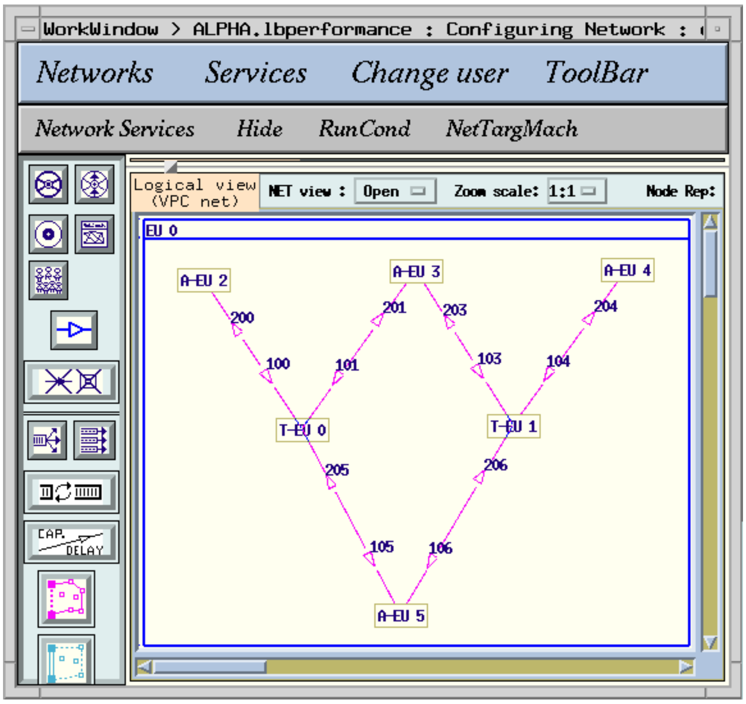

Figure 12.13 VPC network for network A

VPCs	bandwidth (Mb/s)
205	80
106	80
206	80
107	30
207	30
108	60
208	60
109	80
209	80

Table 12.35 VPC bandwidth of VPC network case B

Routes	destination node	CoS	VPCs	initial route selection priorities
R1	5	c1	100, 105	1
R2	5	c1	101, 105	2
R3	5	c1	103, 106	1
R4	5	c1	104, 106	1

Table 12.36 Route definitions for network case A

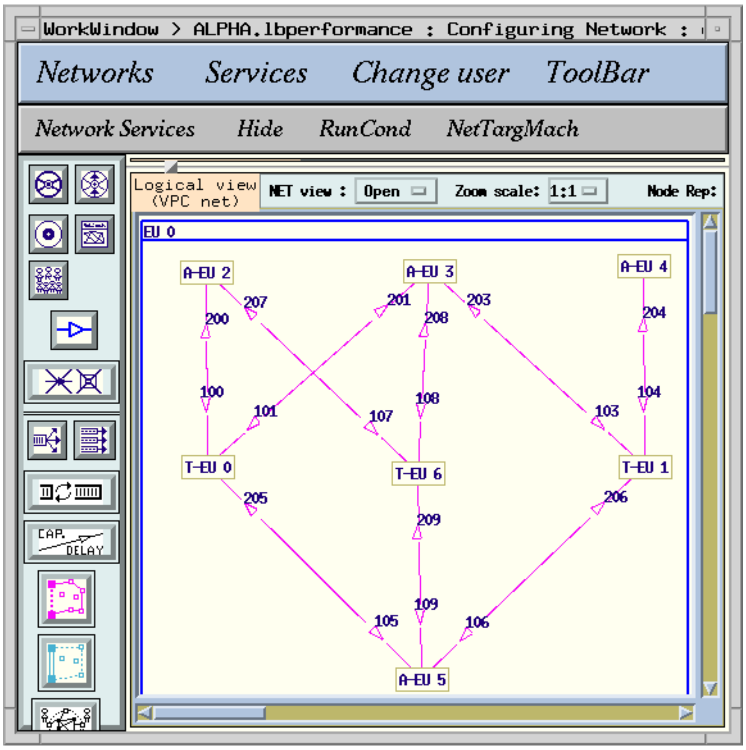

Figure 12.14 VPC network for network B

Routes	destination node	CoS	VPCs	initial route selection priorities
R1	5	c1	100, 105	2
R2	5	c1	107, 109	1
R3	5	c1	101, 105	1
R4	5	c1	108, 109	3
R5	5	c1	103, 106	1
R6	5	c1	104, 106	1

Table 12.37 Route definitions for network case B

For each of the above network cases traffic is generated from the access nodes 2, 3 and 4 to node 5. For a given (s-d) pair, traffic is generated according to the following pattern:

$$T_{(s\text{-}d),\,c}(k) = A + GP_{(s\text{-}d)} * k,\ c=c1,\ k=1,2,3,... \qquad (2)$$

where:

A is an initial time period required for the TMN system to initialise the network simulator in terms of VPCs and routes.

$GP_{(s\text{-}d)}$ is the connection generation period, used to differentiate traffic generation between different (s-d) pairs. Three possible values are considered: 300 (denoted by 3), 600 (denoted by 2), 900 (denoted by 1) time units. Different combinations of $GP_{(s\text{-}d)}$s per the three (s-d) pairs result in different network cases. For instance, the combination GP(2-5)=2, GP(3-5)=3, GP(4-5)=1 means that source node 2 generates traffic every 600 time units, source node 3 every 300 time units and source node 4 every 900 time units.

k index corresponding to a traffic generation instance.

All (s-d) pairs generate connection requests in bulks of 10 connections at each traffic generation epoch. Nodes 2 and 4 generate a total of 50 connections whereas node 3 generates a total of 100 connections. The invocation instances per (s-d) pair are as many it takes for the total number of connection requests to be made taking into account that at a given instant 10 connection requests are generated simultaneously. It is assumed that each successful connection call lasts for the whole duration of the simulation run.

Table 12.38 summarises the different networks defined for the LB dynamic tests which differ in terms of logical topology and of traffic generation patterns.

Networks	**Logical topology**	**Traffic generation pattern**
A-111	VPCs: as in Figure 12.13, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=1, GP(3-5)=1, GP(4-5)=1
A-121	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=1, GP(3-5)=2, GP(4-5)=1
A-131	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=1, GP(3-5)=3, GP(4-5)=1
A-311	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=3, GP(3-5)=1, GP(4-5)=1
A-221	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=2, GP(3-5)=2, GP(4-5)=1
A-311	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=3, GP(3-5)=1, GP(4-5)=1
A-321	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=3, GP(3-5)=2, GP(4-5)=1
A-331	VPCs: as in Figure 12.12, Table 12.34 Routes: as in Table 12.36	as in formula (2) with GP(2-5)=3, GP(3-5)=3, GP(4-5)=1
B-111	VPCs: as in Figure 12.12, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=1, GP(3-5)=1, GP(4-5)=1

Table 12.38 Different network cases for the LB dynamic tests

Networks	Logical topology	Traffic generation pattern
B-121	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=1, GP(3-5)=2, GP(4-5)=1
B-131	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=1, GP(3-5)=3, GP(4-5)=1
B-311	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=3, GP(3-5)=1, GP(4-5)=1
B-221	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=2, GP(3-5)=2, GP(4-5)=1
B-311	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=3, GP(3-5)=1, GP(4-5)=1
B-321	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=3, GP(3-5)=2, GP(4-5)=1
B-331	VPCs: as in Figure 12.14, Table 12.35 Routes: as in Table 12.37	as in formula (2) with GP(2-5)=3, GP(3-5)=3, GP(4-5)=1

Table 12.38 Different network cases for the LB dynamic tests

Each of the above networks will run with or without a TMN system incorporating the LB OS and further for various LB invocation periods. Specifically, the following LB invocation periods are considered:

LB(1): The LB invocation period is smaller than the smallest traffic generation period (specifically every 250 time units)

LB(2): The LB invocation period is in between the smallest and largest traffic generation periods (every 500 time units)

LB(3): The LB invocation period is greater than the largest traffic generation period (every 1000 time units)

LB(4): The LB invocation period is far greater than the largest traffic generation period (every 2000 time units)

The results can be found in Section 12.4.2.3.2.

12.4.2.3 Results

12.4.2.3.1 Static case: initial assignment of route selection priority

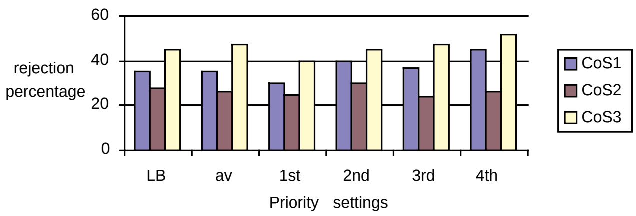

Figure 12.15 Connection rejection ratios under various route selection priority settings for network N1-A

As it can be seen from Figures 12.15 to 12.20 network performance under LB route selection priority settings is as good as in any of the other settings per each CoS. On average (compare with column labelled by ‘av’), the LB priority settings outperform.

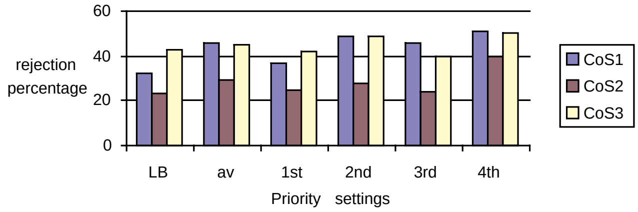

Figure 12.16 Connection rejection ratios under various route selection priority settings for network N1-B

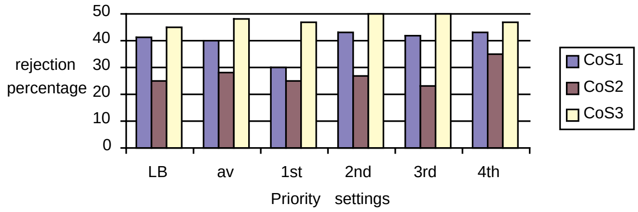

Figure 12.17 Connection rejection ratios under various route selection priority settings for network N1-C

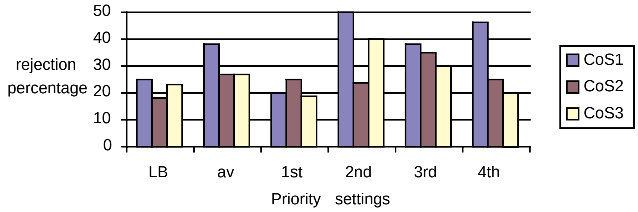

Figure 12.18 Connection rejection ratios under various route selection priority settings for network N2-A

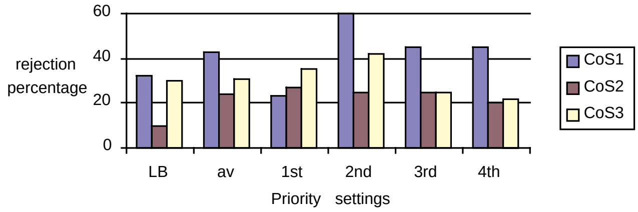

Figure 12.19 Connection rejection ratios under various route selection priority settings for network N2-B

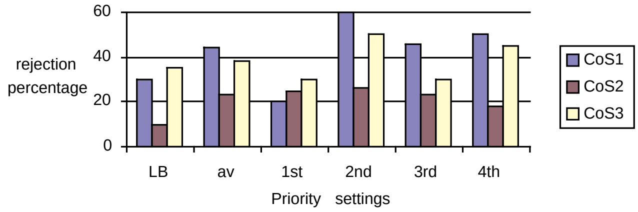

Figure 12.20 Connection rejection ratios under various route selection priority settings for network N2-C

The gain in the reduction of blocking probability is greater in the case of symmetric (uniform) traffic generation patterns (see Figures 12.16 and 12.19) and it is becoming smaller as we are moving to more asymmetric traffic generation cases.

Furthermore, it can be seen that network performance is indeed sensitive to the particular arrangement of route selection priorities per CoS. This observation justifies and further advocates the existence of a management functionality as the LB OSF.

It is worth noting that the marginal blocking probability of CoS c2, which has the largest bandwidth requirements from the other two CoS, is improved best under the LB route selection priority settings than under any other setting. This is attributed to the fact that the LB algorithm takes into account the bandwidth requirements of the CoSs that have common routes in the process of prioritising the available routes per CoS. In this sense, LB protects the higher bandwidth CoSs at the expense of the lower bandwidth CoSs when sharing common (parts of) routes. The other route selection priority settings treat all CoSs the same. This feature is useful in a multi-class network environment especially when different billing policies apply to different CoSs.

To better highlight this point, we aggregate the marginal connection rejection ratios of each CoS using weights proportional to the bandwidth requirements of each CoS, namely we consider the following aggregate of the marginal connection rejection ratios:

$$\frac{\sum_{c} B_c R_c}{\sum_{c} B_c}$$

where B_c, R_c are the mean (or peak or effective) bandwidth requirements, connection rejection ratio, respectively of CoS c. This aggregate rejection ratio reflects the network loss in terms of offered bandwidth; hence, we call it offered bandwidth loss percentage. Considering that whenever a connection is rejected, the network loses an amount of revenue proportional to the bandwidth requirements of the connection, the previously introduced measure is indicative of the losses in network revenue, considering that each CoS is charged only on the basis of its bandwidth requirements (not on the basis of duration, quality etc.).

Figures 12.21 to 12.26 show the offered bandwidth loss percentage under the various route selection priority settings, for each of the test networks considered. As it can be seen LB outperforms the other route selection priority methods. It should be noted that although there might be a particular route selection priority method performing better than the LB method for a specific network (e.g. see Figure 12.21, LB vs ‘1st’), throughout all networks this particular method is not constantly better than the LB method. Figure 12.27 summarises the results shown in Figures 12.21 to 12.26, by providing a comparison between the performance, in terms of offered bandwidth loss percentage, of the LB route selection priority method and the average performance of the four other priority methods used. As it can be seen LB provides better performance throughout all networks. The greatest gains are attained in the case of partial route sharing (test network case N2, see also Figures 12.24 to 12.26).

The above results show that the LB functionality indeed contributes positively to network operation especially considering the multi-service nature of the network environment. Most importantly, through its functionality LB may ‘protect’ certain CoSs by appropriately influencing the routing decisions, which may lead to desired levels of overall network optimality. The current implementation of LB takes into account the bandwidth requirements of the different CoSs that might share common parts of routes. The results showed network improvement in terms of lost offered bandwidth. Other criteria differentiating access of different CoSs to common resources (routes) might be incorporated to its route selection functionality. Therefore, LB offers the network the opportunity to optimise its overall performance according to its business objectives, by adjusting routing between the different CoSs. This is yet another evident proving the need for a management component with the specifications of the LB OSF.

Another point to observe is that overall network performance improves under partial route sharing as opposed to complete route sharing, under any route selection priority setting. The best network performance is achieved under partial route sharing and LB (see Figure 12.27). Moreover the benefits from LB are greater in the case of partial route sharing (see Figures 12.18 to 12.20 and 12.24 to 12.27) than in the complete route sharing case (see Figures 12.15 to 12.17 and 12.21 to 12.23). This is attributed to the fact that the LB functionality takes into account the occasions with which a particular VPC is selected for routing. Note that in the partial route sharing case all SDClass-Route networks have the same alternatibility.

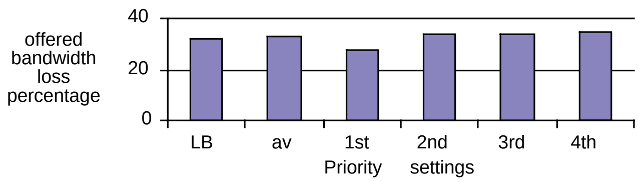

Figure 12.21 Offered bandwidth loss percentage under various route selection priority settings for network N1-A

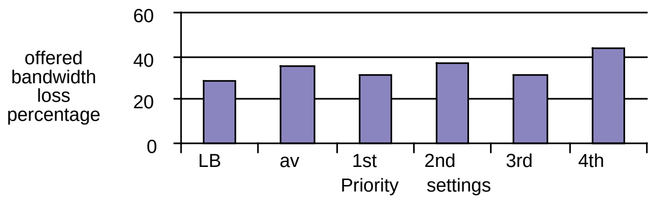

Figure 12.22 Offered bandwidth loss percentage under various route selection priority settings for network N1-B

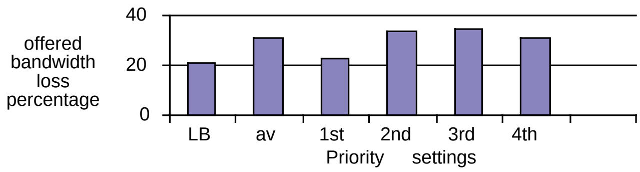

Figure 12.24 Offered bandwidth loss percentage under various route selection priority settings for network N2-A

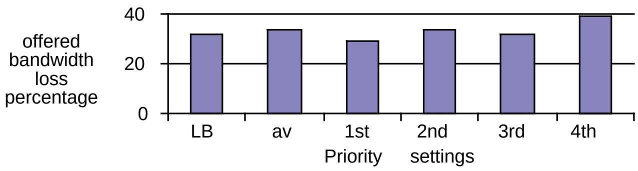

Figure 12.23 Offered bandwidth loss percentage under various route selection priority settings for network N1-C

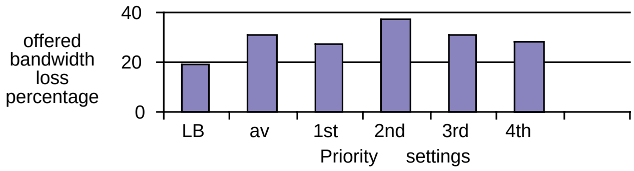

Figure 12.25 Offered bandwidth loss percentage under various route selection priority settings for network N2-B

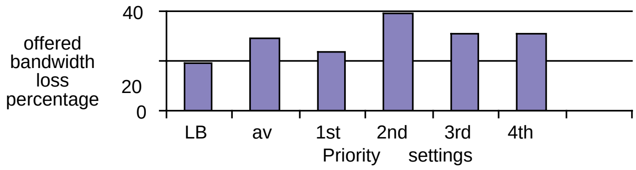

Figure 12.26 Offered bandwidth loss percentage under various route selection priority settings for network N2-C

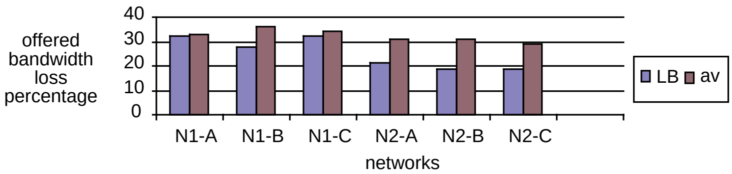

Figure 12.27 Performance of LB route selection priority settings (LB) against average performance of four other settings (av) over different network configurations

12.4.2.3.2 Dynamic case: dynamic assignment of route selection priorities

As it can be seen from the above figures, network performance in terms of connection call blocking probability, improves under LB. The results demonstrate that the LB management system can indeed improve network performance. The rate of improvement depends on the LB invocation period. Although the network environment of the tests carried out is simple, it might be argued that it is representative of the vast major-

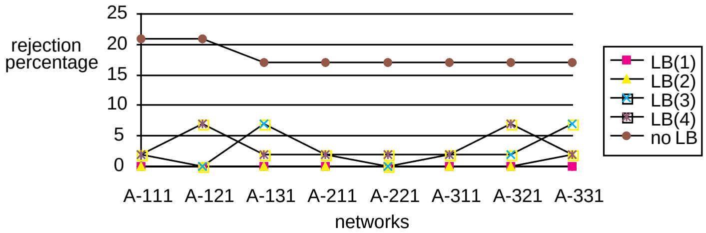

Figure 12.28 Connection rejection ratios, under various LB invocation frequencies, for networks A-111 to A-311

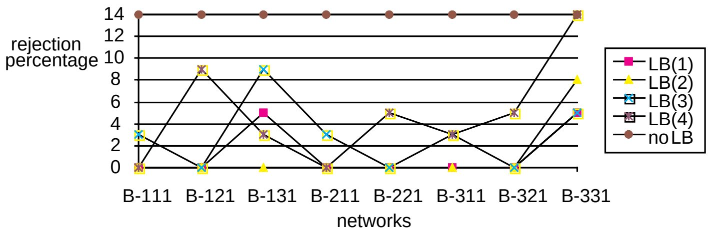

Figure 12.29 Connection rejection ratios, under various LB invocation frequencies, for networks B-111 to B-331

ity of network cases, as it assumes different alternatabilities and degrees of sharing common resources, between different SDClassRoute networks.

The benefits of LB are greater in the test network case B than in the network case B. This indicates that LB improves its performance in cases where all SDClassRoute networks have more or less the same alternatibility, with its performance being deteriorating as we are moving to cases with significant asymmetry in the alternatibility of the defined SDClassRoute networks, especially when combined with traffic generation asymmetry in favour of the (s-d) pairs with the highest alternatibility in their SDClassRoute networks. This behaviour is explained in the following paragraphs.

It is worth noting that there is not a particular activation period that offers constantly, throughout all networks, the best performance. This indicates that the LB should be activated asynchronously based on actual network conditions.

Although it is intuitively anticipated better performance in the cases where LB is activated more frequently as opposed to the cases where it is activated less frequently, the results showed that this is not the general case. This is especially evident in the network case B (see Figure 12.29), which has more asymmetry in the alternatibility between the SDClassRoute networks than network case A. This behaviour can be explained as follows. The frequent invocation of LB results in continuous balancing the load amongst the alternative routes of certain CoSs; if the load on these alternative routes becomes high enough due to traffic from certain source-destination pairs, this continuous refinement blocks the connections between other source-destination pairs that might use parts of these routes as unique alternatives. This can be clearly seen in the case of networks B-131 and B-331 in Figure 12.29, where the LB being activated faster than even the traffic is generated yields worst performance than the LB being activated far less frequent than the traffic is generated. The current implementation of

LB takes into account the asymmetry in alternatibility that might exists between the defined SDClassRoute networks; however, the load of the VPCs along the routes prevails.

The results indicate that LB should base its functionality on some sort of predictions regarding anticipating network usage. These predictions should be made with medium-term perspective and should refer within the time-frame of the predictions based on which the VPC and route networks were built (made by PUM and used by RD). Based on the insight gained from the results, it is believed that LB should utilise information as the following for making its medium-term predictions: actual VPC usage (from the network) for the purpose of deriving increasing or decreasing trends, maximum number of active connections per source-destination pair and per CoS that the network should provide so that to guarantee the acceptable blocking probability per CoS (from RD) and actual number of active connections per (s-d) pair and CoS (from the network). The latter two measures should allow LB to forecast the anticipating network usage from the various traffic sources and therefore make better judgements on the sharing of routes. These measures are not taken into account by the current specifications of the LB OSF and are suggested as updates to the VPCM architecture.

12.4.2.4 Conclusions on the Load Balancing performance tests

This section presented the results of the tests regarding the performance of the LB OS in the network level. The purpose of these tests was to establish an indication of the effect of the LB OS in the network, justifying therefore the need for a management component such the LB component. Moreover, the tests undertaken demonstrated the use of the developed ICM TMN testbed for testing the functionality and performance of various TMN systems/functions. Furthermore, these tests as well as the functional (capability) tests, presented in previous sections, served as a means for validating the ICM methodologies for specifying, designing and testing management functionality.

The performance of LB was assessed in a number of network configurations regarding logical topologies and network offered load patterns. LB performance was assessed in both of its functional modes; static and dynamic. The test cases considered are not exhaustive of the many different possible network configurations. However, the tests carried out, indicated network performance improvement for networks employing a management system including the LB OS, over networks which do not employ such management schemes.

Specifically, the following conclusions may be drawn regarding LB performance:

- There is indeed scope for route selection management through a management functional component as the LB OSF, especially considering the multi-service nature of the network environment. The tests showed network performance improvement under LB. Gains in network performance may be obtained in terms of connection rejection ratio and network load balancing at link level. In a multi-class environment, LB improves the marginal connection blocking probabilities per CoS, in the sense that significantly improves overall network performance in terms of the offered bandwidth loss probability. Marginal connection blocking probability is also improved for some CoSs, especially for those of higher bandwidth requirements. This behaviour is particularly useful for guaranteeing acceptable performance in multi-class networks and most importantly shows how routing management policies may be used to achieve the business objectives of the network operators.

- The gains in network performance depend on
  - Route alternatibility.

    It was shown that the higher the alternatibility in the routes of the supported network CoSs, the higher the performance gains. This is attributed to the fact that route alternatibility increases LB management flexibility, therefore the positive potential of the impact of the routing decisions as influenced by the LB OS. However, LB performance deteriorates in cases where there is asymmetry in the alternatibility of the defined SDClassRoute networks, especially combined with asymmetry in traffic generation particularly in favour of the (s-d) pairs with the higher alternatibility n their SDClassRoute networks.
  - Route sharing.

    It was shown that network performance under LB is more improved in cases where there is partial sharing of the defined routes amongst the supported network CoSs. This is attributed to the fact that the route selection management algorithm of the developed LB OS takes into account the routing occasions in which the different network CoSs may use the defined routes; prioritising routes in favour of CoSs using these routes with less alternatibility as opposed to CoSs using the same routes with higher alternatibility. Considering the multi-class environment of IBC networks, this feature is yet another justification of a management component like LB.
  - LB activation epochs.

    It was shown that LB performance depends primarily on the time epochs at which it will be activated. The intuitive argument that LB should be activated sufficiently prior to instances where offered network changes, was shown that it is not generally true. No fair statement can be made regarding the most appropriate LB activation mechanism. However, this does not make obsolete the existence of a management component like LB. Associated with this issue is the ‘usual’ problem of routing, namely it is difficult to assess the impact of routing decisions taken at one time, to the future i.e. how can it be guaranteed that routing decisions that seemed ‘good’ at one time will continue to be ‘good’ for the next period of time. There is a need for enhancing LB functionality to make some sort of predictions regarding network usage. Based on these predictions route prioritisation and determination of next activation epoch should be made. Note that these predictions should refer within the time-frame of the source-destination traffic predictions based on which the VPC and route networks were built. From the insight gained from the results there are two kinds of information that should be utilised in making such predictions: actual VPC usage and actual number of active connections per (s-d) pair and CoS.

Further tests are required in order to assess in full extent the effect of LB behaviour in network operation, gain more insight into its functional aspects and support with more confidence the above mentioned conclusions. The preliminary results presented in this section encourage the undertaking of such tasks. An important dimension of future testing is to test LB performance in a more complicated network management environment, involving interactions with other performance management components, like BD.

12.4.3 Bandwidth Distribution tests

This section presents tests for validating the functionality of the Bandwidth Distribution OSF at a system level. The tests have been designed to be carried out using the ATM simulator.

12.4.3.1 Introduction

The purpose of Bandwidth Distribution (BD) application test is to examine and verify the functionality of the BD OSF in a network environment. We create variable traffic with the purpose to examine and verify the functionality of BD OSF in all (most of) possible traffic conditions in which the BD OSF might be activated.

It should be noted that the purpose of the test is not the detailed assessment of the performance of the BD OSF, in terms of benefits and costs to the network operation, sensitivity, management overhead, scaling and processing time. This test aims at validating the functionality of the BD OSF. Performance assessment is regarded as a huge task and it is not addressed in detail by this work. However, the validation tests may provide indications on the performance of the BD OSF.

12.4.3.2 Description of test cases

The functionality of BD depends on:

- the load on the VPCs,
- the required bandwidth of a VPC,
- the allocated bandwidth of a VPC,
- the topology of the VPCs.

The functionality of BD does not depend on the particular network CoSs and their bandwidth characteristics. Rather it depends on their aggregated load on the VPCs and the variation in their arrival pattern to the VPCs.

The test cases should cover all possible values of the above parameters but to reduce unnecessary complexity, we focus on a limited set of VPCs (tagged VPCs), with different topologies but spanning over common links. Moreover, the allocated bandwidth and the required bandwidth of the VPCs is assumed to be constant. Creating possible combinations of load on the selected VPCs, the functionality of the BD OSF can be verified, by checking the actions taken against to what is expected. For a given load situation, the expected BD actions can be deduced following its functional principle of redistributing links' "spare" bandwidth to the more congested VPCs.

12.4.3.3 Tests and results

Test conditions

The physical topology of the test network used was as described in Section 12.4.1.2. Only two links were used for the tests: link 2 and link 20. These were both assigned a bandwidth of 155Mb/s.

The BD OS was run in simulated time. As well as the simulator, the following OSs were also running at the same time as BD: Simulator QAF (Simqaf), Configuration Manager (CM), Current Load Model (CLM), Load Balancing (LB), Route Design (RD). BD was invoked initially by RD and then periodically every 10 seconds of simulated time.

In addition to the above, a data collection manager was developed to acquire the results of the tests. This manager created the necessary objects in the Simqaf to collect call statistics and logged selected statistics (including VPC effective and allocated bandwidths) to a file. These results were then analysed and plotted using a spreadsheet to produce the graphs which follow.

The following test scenarios were used:

VPC Scenario 1:

VPC10 from node1 to node2, using link 20
VPC11 from node0 to node2, using link 2 and link 20

VPC Scenario 2:

As for scenario 1, but in addition:
VPC12 from node0 to node1 over link 2.

Traffic Scenario 1:

per user, 1Mb/s CBR, call time 5 seconds, silence time 5 seconds.
User groups arranged to generate the following:

- On VPC10:
  - time 100-200 seconds 50 users,
  - time 200-300 seconds 100 users,
  - time 300-400 seconds 150 users,
  - time 400-500 seconds 100 users,
  - time 500-600 seconds 50 users.
- On VPC11:
  - time 100-200 seconds 150 users,
  - time 200-300 seconds 100 users,
  - time 300-400 seconds 50 users,
  - time 400-500 seconds 100 users,
  - time 500-600 seconds 150 users.

Traffic Scenario 2:

As for Traffic Scenario 1 but in addition:

- on VPC12:
  - time 100-350 seconds 0 users,
  - time 350-600 seconds 100 users.

Traffic Scenario 3:

- On VPC10:
  time 100-350 seconds 50 users,
  time 350-600 seconds 150 users.
- On VPC11:
  time 100-600 seconds 75 users.

With the link capacities and traffic types used, 200 users corresponds to approximately 100% load on the physical links. The above scenarios thus represent steps in the range of 25% to 75% potential link utilisation on both VPCs. The traffic for Traffic Scenario 1 is experimentally measured and shown in Figure 12.30. These results were obtained by making link and VPC bandwidths so large that no calls were rejected. The graph thus shows the load offered to the network.

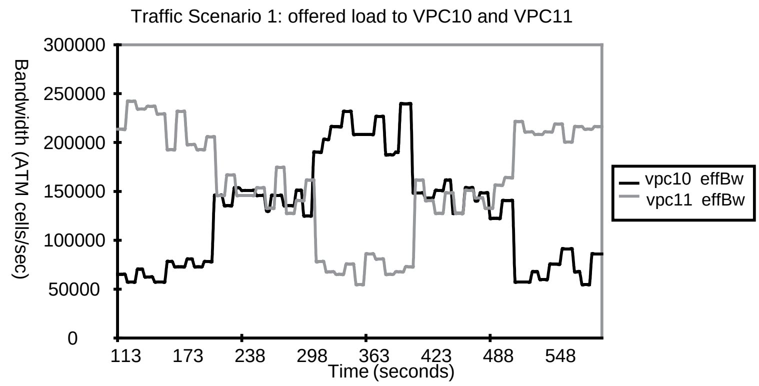

Figure 12.30

Although, strictly, traffic is generated with a source/destination node pair rather than for a particular VPC, the routing options available in the test scenario are so restricted that the selected traffic can only pass down the VPC for which the test is intended. At a particular node, there is only one route to a particular destination. This restriction was necessary in order to remove any effects of routing: these tests are purely designed for BD.

The following tests were carried out:

1. VPC Scenario 1, Traffic scenario 1, VPC10 and VPC11 both class 1 VPCs (leased lines) with allocated bandwidth for each set at half the total link capacity.
2. As for test 1 but with both VPCs class 3.
3. As for test 1 but with VPC10 class 3, VPC11 class 2.
4. VPC Scenario 2, Traffic Scenario 1, VPC10 class 2, VPC11 class 3 VPC, VPC12 class 1 with allocated bandwidth half the physical link capacity.
5. As test 4 but with all VPCs class 3 and Traffic Scenario 2.
6. VPC Scenario 1, Traffic Scenario 3, VPC10 class 3, VPC11 class 4.

Test 1 results

For this test, both VPCs are class 1, and according to the algorithm for BD, class 1 VPCs should not have their bandwidth changed. The allocated and effective bandwidths for VPC10 and VPC11 are shown in Figure 12.31 and Figure 12.32, from which it can be seen that the correct initial value of bandwidth is allocated to the VPCs (180,000 cells/sec) and that the allocated bandwidth has not changed during the test run for either VPC.

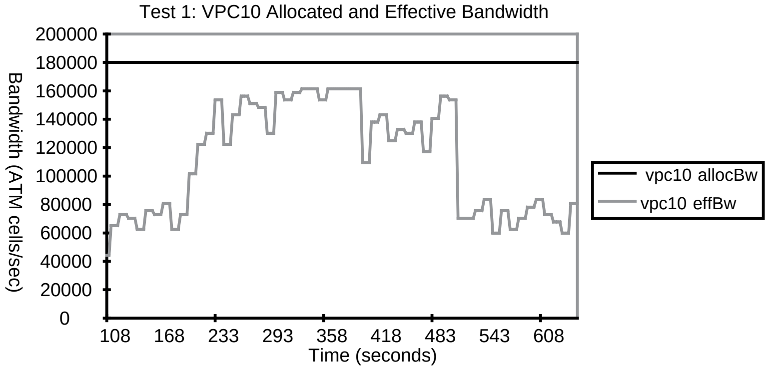

Figure 12.31

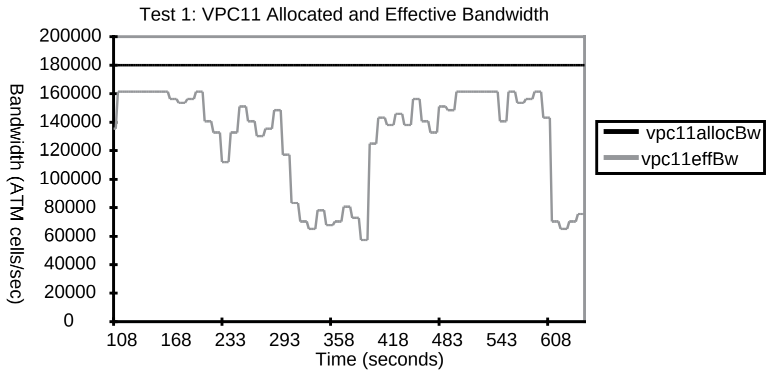

Figure 12.32

This test shows the correct behaviour of BD for class 1 VPCs.

Since the offered load for both VPCs is significantly above the available bandwidth on the VPCs for some parts of the test, it can be expected that calls are rejected during these times. This result is shown in Figure 12.33.

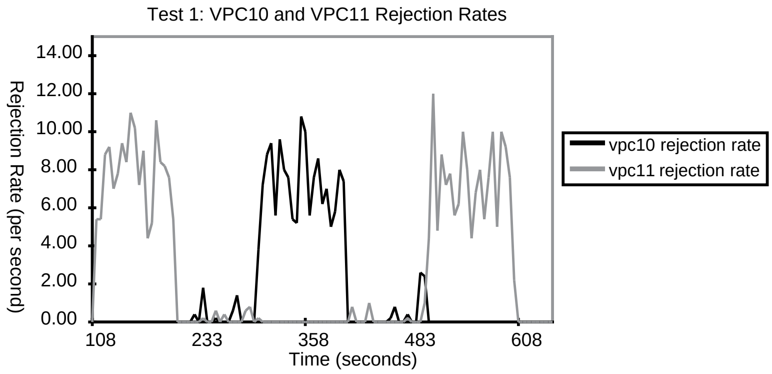

Figure 12.33

Test 2 results

For this test, both VPCs are set to class 3. According to the BD algorithm, any spare bandwidth on links is distributed equally (as far as is possible) amongst the class 3 VPCs on the link. In this case, VPC11 is the only VPC on link 2, but shares link 20 with VPC10. We would expect BD to share the total available bandwidth on link 20 equally between VPC10 and VPC11. It is not possible to allow VPC11 to use all the spare bandwidth on link 2 because this would leave no spare bandwidth to share with VPC10 on link 20. BD is therefore unable, in this case, to distribute all the available bandwidth on link 2.

Figure 12.34 and Figure 12.35 show the result of the test. The allocated and effective bandwidths are plotted for each of the VPCs.

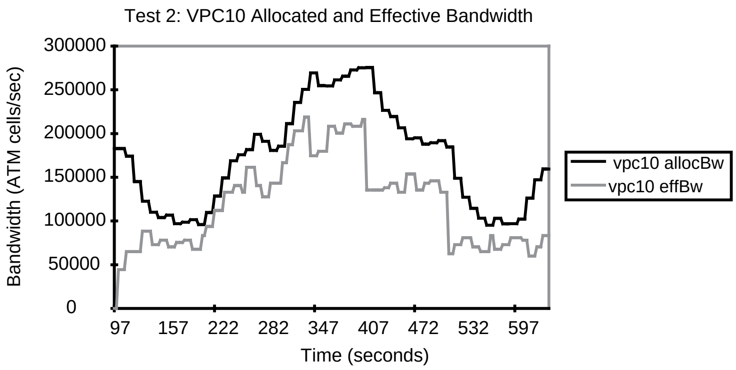

Figure 12.34

It can be seen that initially (time 97 seconds) the total bandwidth available on the link was shared equally between the two VPCs (at about 180,000 cells/sec). Later, BD has allocated more bandwidth to a VPC as its load increases, reducing the bandwidth

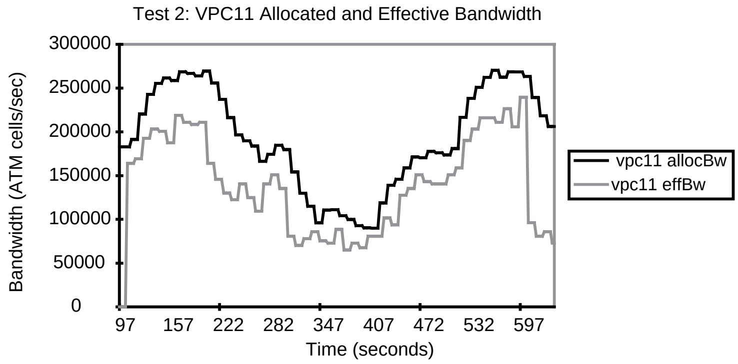

Figure 12.35

available to the other VPC. This is in accordance with the rule that when the effective bandwidth exceeds an upper threshold value (75% of allocated bandwidth), the allocated bandwidth is increased, if possible. This reduces the pool of spare bandwidth available for sharing amongst class 3 VPCs.

A further rule for BD is that the allocated bandwidth cannot be reduced below the initial value allocated to the VPC by Route Design. In this case, the initial allocated value was set at 90,000 cells/sec. It can be seen that this rule is not violated in the test.

This test shows the correct behaviour of BD for class 3 VPCs.

The call rejection rates for the two VPCs are plotted in Figure 12.36

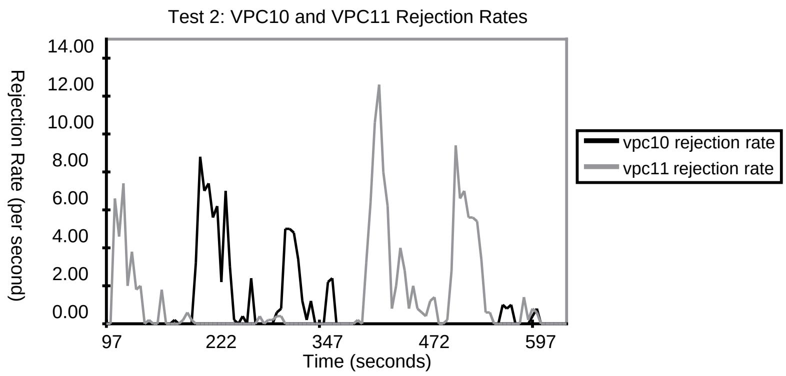

Figure 12.36

The call rejection rates peak when the traffic on a VPC increases faster than the BD algorithm can deal with. The imposed traffic pattern is in fact rather severe, representing near step functions in load changes.

A reasonable measure of the performance of BD in adapting to changing loads is to calculate the blocking probability seen with and without BD. This is calculated as:

blocking probability = calls rejected / calls accepted

The values over the whole period of the test (600 seconds) are calculated for both VPCs for this test and also for the previous test, where BD was effectively prevented from distributing bandwidth by virtue of the fact that both VPCs were leased lines (class 1 VPCs). These results are shown below.

VPC	test 1 blocking probability	test 2 blocking probability
10	0.222	0.109
11	0.363	0.156

Table 12.39

It can be seen that when BD distributes bandwidth according to the traffic, a considerably improved performance is seen.

Test 3 results

This test is a repeat of the above test, but with VPC10 as a class 2 VPC and VPC11 as class 3. The algorithm for BD states that spare bandwidth on a link is to be evenly distributed amongst class 3 VPCs on the link (but not to class 2 VPCs). Also, when the effective bandwidth on a class 2 VPC crosses the high (75%) threshold mark, more bandwidth is to be allocated to it, with any remaining then distributed to class 3 VPCs. Only when the low threshold (25%) is crossed will bandwidth be taken back from a class 2 VPC and become available for class 3 VPCs.

The results are plotted in Figure 12.37, Figure 12.38 and Figure 12.39.

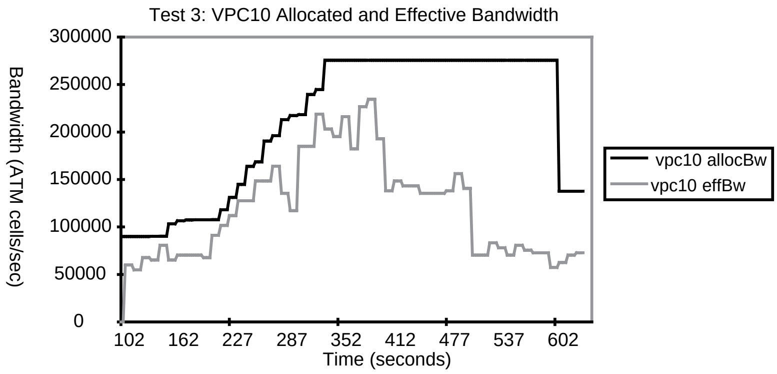

Figure 12.37

It can be seen that at the start of the test, BD has correctly allocated the required bandwidth (90,000 cells/sec) to the class 2 VPC (VPC10) and distributed all the remaining bandwidth to class 3 VPCs, in this case just VPC11. (By good fortune, this happens to match the traffic quite well at this point in the test.) As the traffic on VPC10 increases and crosses the upper threshold (75% of the allocated bandwidth) more bandwidth is allocated to this VPC, leaving less to be distributed to VPC11. The bandwidth allocated to VPC10 is not then reduced until the effective bandwidth drops below the

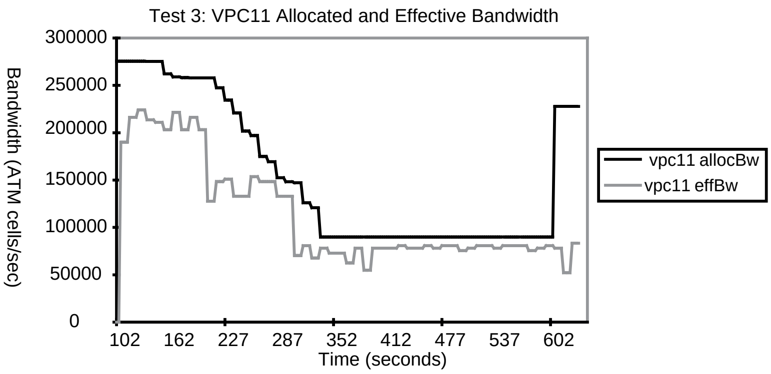

Figure 12.38

low threshold (25% of allocated bandwidth). This occurs at about time 602 seconds (Figure 12.37). At this time the freed bandwidth on the link is reallocated to VPC11.

This test shows that BD behaves correctly for class 2 VPCs, and that precedence is correctly given to class 2 VPCs over class 3 VPCs.

The VCC rejection rates are plotted in Figure 12.39. This shows a consistently low value for the class 2 VPC10 (which has been allocated most of the available bandwidth). A very high figure is shown for the class 3 VPC11 in the second period of the test (400 seconds on), where the traffic demand on VPC11 increases substantially, but no bandwidth is available on the link until the bandwidth allocated to the class 2 VPC can be reduced.

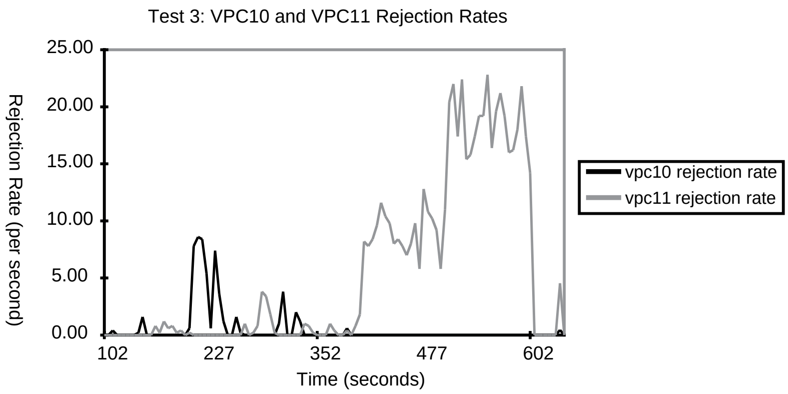

Figure 12.39

Test 4 results

For this test, VPC10 remains a class 2 VPC and VPC11 a class 3 VPC. However, in this case there is a leased line (VPC12) sharing link 2 with VPC12. The leased line has been allocated half the bandwidth of the link. BD is therefore restricted in how much bandwidth it can allocate to VPC12 - it is limited by what is available to VPC12 on link 2.

The results of this test are plotted in Figures 12.40, 12.41 and 12.42.

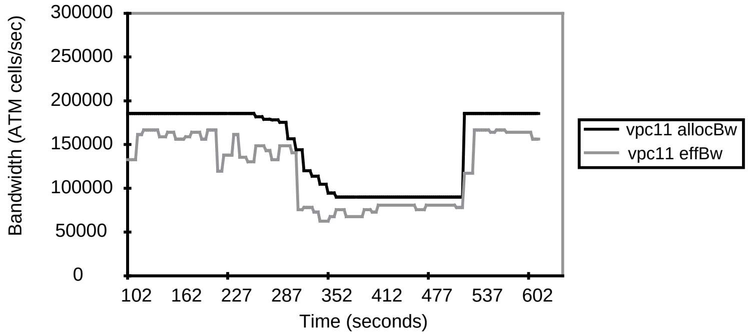

Figure 12.40

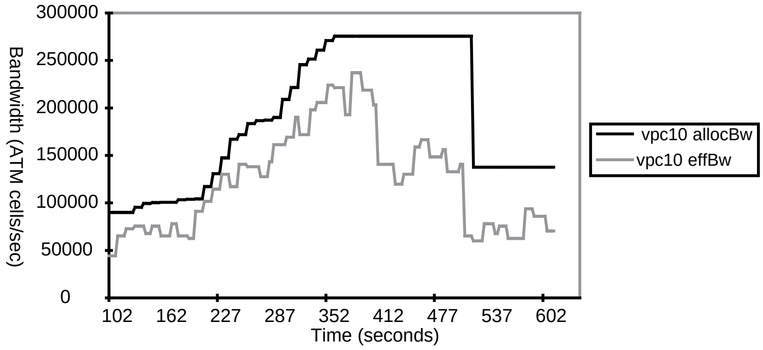

Figure 12.41

Comparing Figures 12.41 and 12.37 it can be seen that the behaviour of BD for the class 2 VPC (VPC10) is unaffected. (Any differences seen are due to the random nature of the traffic.) Comparing Figures 12.40 and 12.38, it can be seen that the maximum bandwidth allocated to VPC11 is now reduced, due to the restriction on what is available on link 2.

This demonstrates that BD correctly takes into account available bandwidth on all the links which a VPC traverses when distributing spare bandwidth to a class 3 VPC.

The result of this can be seen in much increased reject rates on VPC11 at the start of the test (when the traffic demands are high but the bandwidth is restricted).

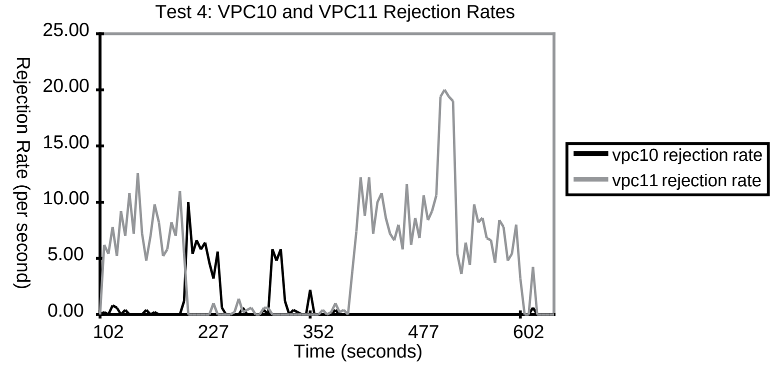

Figure 12.42

Test 5 results

For this test, both VPCs are class 3, allowing any spare bandwidth to be distributed evenly between them. There is also now another VPC, also of class 3, (VPC12) sharing link 2 with VPC11. VPC12 has traffic which is zero for the first half of the test, and rises to about 50% of the link capacity for the second half of the test. BD must therefore share bandwidth amongst the class 3 VPCs on both link2 and link 20, with one of the VPCs being common to both links.

The results of this test are plotted in Figure 12.43, Figure 12.44, Figure 12.45.

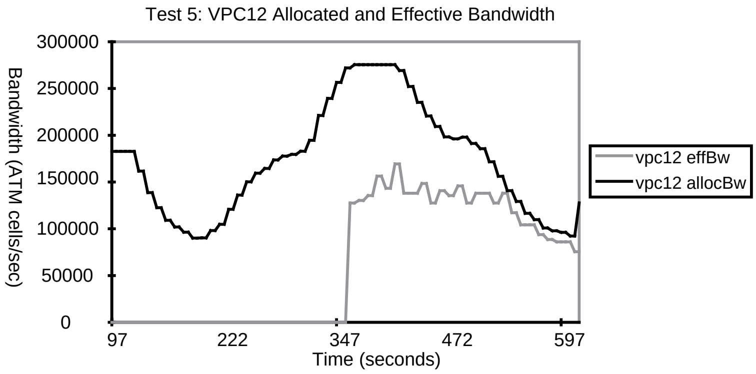

Figure 12.43

It can be seen that BD initially shares the link bandwidth equally between VPC10 and VPC11 on link 20 and is also able to share the bandwidth equally between VPC11 and VPC12 on link 2. As the effective bandwidth on VPC11 crosses the high threshold, more bandwidth is allocated to it, reducing the spare pool available for both VPC10

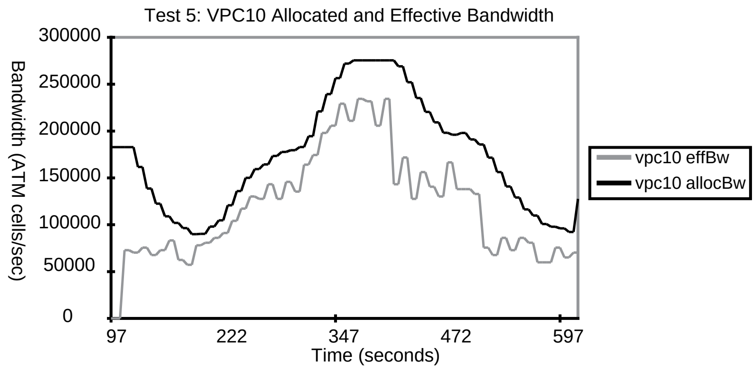

Figure 12.44

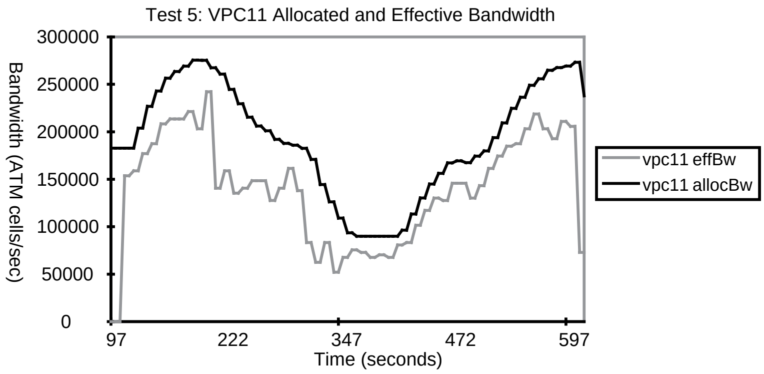

Figure 12.45

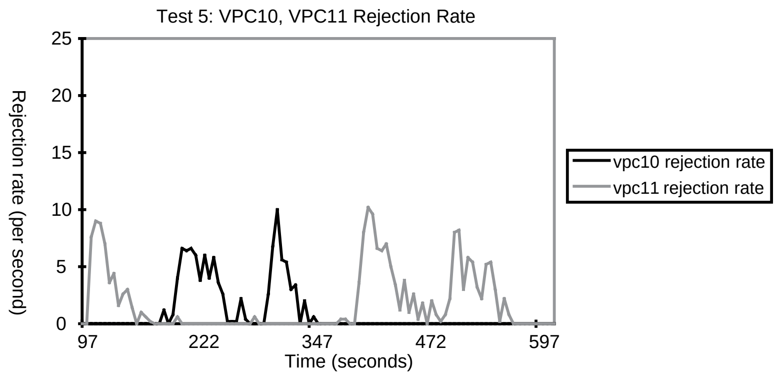

Figure 12.46

and VPC12, so the allocated bandwidth for these two drops. Later, as the effective bandwidth on VPC10 rises, the spare bandwidth available to VPC11 is reduced, and its allocated bandwidth drops. This increases the spare bandwidth available to VPC12 on the other link, so its allocated bandwidth increases.

Later, as the traffic on VPC12 starts, there is initially plenty of spare bandwidth on this VPC. However, as the load on VPC11 also increases, the spare bandwidth on link 2 reduces to zero. It can be seen that towards the end of the test (Figure 12.43), bandwidth continues to be taken away from VPC12 and given to VPC11.

Test 6 results

This test is designed to show the effect of mixing a class 4 and a class 3 VPC. BD should allocate spare bandwidth on a link to class 3 VPCs. When more bandwidth is required on a VPC due to a high threshold being crossed, VPCs of class 3 should take precedence over those of class 4.

For this test, VPC scenario 1 is used with traffic scenario 3. VPC10 is set to class 3 with an initial bandwidth allocation of 90000 cells/sec (about 25% of link capacity). VPC11 is set to class 4 with an initial bandwidth allocation also of 90000 cells/sec.

The results are shown in Figure 12.47, Figure 12.48 and Figure 12.49.

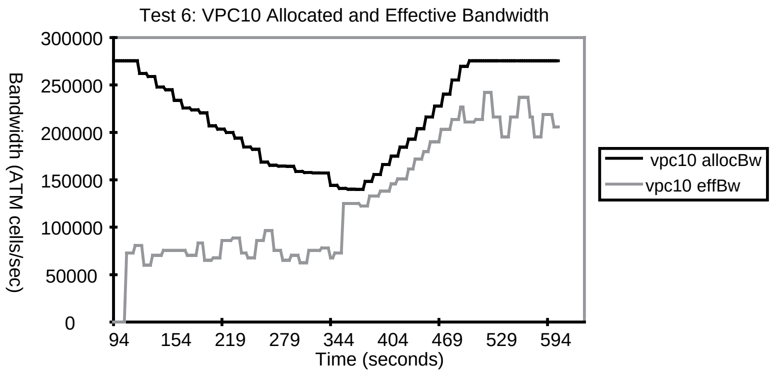

Figure 12.47

It can be seen in Figure 12.47 at time 100 seconds that all the spare bandwidth on the link has been allocated to VPC10, the class 3 VPC. At the same time (Figure 12.48) the class 4 VPC (VPC11) has been allocated only the initial 90000 cells/sec.

From time 100 to 350 seconds BD allocates more bandwidth to VPC11 (Figure 12.48) because the effective bandwidth on this VPC exceeds the upper threshold. This reduces the pool of spare bandwidth on the link, so the bandwidth allocated to VPC10 drops (Figure 12.47). During the period when BD is raising the bandwidth on VPC11 some calls are rejected on this VPC (Figure 12.49) whilst none are rejected on VPC10 due to the generous bandwidth provision here. At time 350 seconds the traffic on VPC10 rapidly increases, and BD allocates more bandwidth to it (Figure 12.47). This is done at the expense of the class 4 VPC, VPC11, so the bandwidth allocated to VPC11 is reduced, eventually back to the initial value of 90000 cells/sec

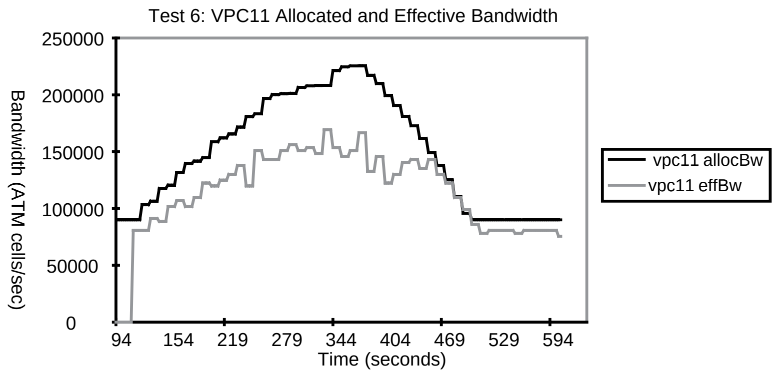

Figure 12.48

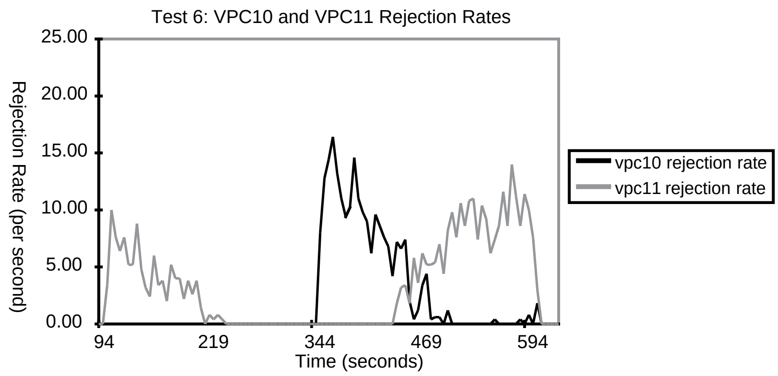

Figure 12.49

(Figure 12.48). The provision on VPC11 is now insufficient so that the call rejection rate rises (Figure 12.49).

This test shows that BD correctly allocates spare bandwidth to class 3 VPCs in preference to class 4 VPCs, after taking into account extra requirements of class 4 VPCs due to high threshold crossing of effective bandwidth. It also shows that when more bandwidth is required by a class 3 VPC, BD correctly allocates this at the expense of class 4 VPCs.

12.4.3.4 Conclusions on the Bandwidth Distribution tests

For relatively simple scenarios involving mixes of classes of VPCs over two physical links, BD has been shown to behave as specified in the design document. It was not the purpose of these tests to investigate the efficacy of the algorithms used in BD, although some conclusions can be drawn from the results presented above. Further tests would be required involving more complicated interactions between VPCs of mixed classes traversing multiple links and a wider range of traffic conditions to refine the algorithms further.

12.5 Conclusions

This chapter has presented the ICM validation rationale, followed by a summary of the testing performed and details of two of the test sets performed to validate the ICM TMN functional implementations. The experimental results and their evaluation aimed to show that the TMN system performed as specified, and to show the TMN applications to be able to improve the performance of the managed network.

The following points outline the main conclusions that can be drawn from the results of the testing work undertaken by ICM:

- The results verify that the developed components of the ICM testbed operate according to their functional specifications. Various functional validation tests have been carried out for each of the testbed components, validating their interfaces, their intelligence and their inter-operation. Functional validation has been asserted from a number of test cases covering a representative set of network configurations and traffic conditions.
- The results indicate that the management services and components developed by the project can improve network performance in a number of test experiments covering representative cases of network configuration and traffic scenarios. These preliminary results justify the role of performance management systems and advocate research in the area of efficient network management schemes and algorithms.
- The successful completion of the validation and verification tests not only validates the specific management algorithms employed but also validates the overall system architecture. Furthermore the ICM approach and methodology to management systems specification, design and implementation (see Chapter 3) is validated.
- Throughout the testing process, the use of network simulation tools was validated and their usefulness and utility in carrying out complex testing tasks was verified. Their flexibility (in terms of network configurations, traffic scenarios, network functions), portability and scaleability make to possible to define and execute many different test scenarios which would be infeasible to achieve on real networks given realistic constraints of time, effort and cost.
- The facilities offered by the ICM management platform were proved to be very useful throughout the testing process. In particular, the MIB browser, the generic management applications and the CMIS scripting language were invaluable for testing and debugging purposes.
- The problem of testing and evaluating the performance of management functionality is of enormous complexity, beyond that of software system testing. This is due to the complexity and uncertainty of the network environment and is compounded by the critical role that management plays in network operation. ICM, taking this into account, formulated methods for integrating and testing management functionality. The procedures adopted were validated through the successful completion of the testing activities and although they may not be complete in several aspects, they outline a general approach to management systems testing.

12.6 References

[12.1] RACE R2059 ICM Deliverable 12, “Integration and Validation Principles,” R2059/NTUA/EECS/DR/R/P/012/b1, Dionysis Reisis, editor, October 1994.
[12.2] RACE R2059 ICM Deliverable 18, “ICM TMN Testbed Description,” R2059/CRA/ATG/DS/018/b1, Richard Lewis, editor, December 1995.
[12.3] RACE R2059 ICM Deliverable 21, “Integration, Performance Verification and Evaluation of Test Results,” R2059/GNN/MIH/DS/P/021/b1, Michael Hansen, editor, December 1995.