INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

E.360.5

(05/2002)

SERIES E: OVERALL NETWORK OPERATION, TELEPHONE SERVICE, SERVICE OPERATION AND HUMAN FACTORS

International routing plan

QoS routing and related traffic engineering methods – Transport routing methods

ITU-T Recommendation E.360.5

ITU-T E-SERIES RECOMMENDATIONS

OVERALL NETWORK OPERATION, TELEPHONE SERVICE, SERVICE OPERATION AND HUMAN FACTORS

INTERNATIONAL OPERATION	
Definitions	E.100–E.103
General provisions concerning Administrations	E.104–E.119
General provisions concerning users	E.120–E.139
Operation of international telephone services	E.140–E.159
Numbering plan of the international telephone service	E.160–E.169
International routing plan	E.170–E.179
Tones in national signalling systems	E.180–E.189
Numbering plan of the international telephone service	E.190–E.199
Maritime mobile service and public land mobile service	E.200–E.229
OPERATIONAL PROVISIONS RELATING TO CHARGING AND ACCOUNTING IN THE INTERNATIONAL TELEPHONE SERVICE	
Charging in the international telephone service	E.230–E.249
Measuring and recording call durations for accounting purposes	E.260–E.269
UTILIZATION OF THE INTERNATIONAL TELEPHONE NETWORK FOR NON-TELEPHONY APPLICATIONS	
General	E.300–E.319
Phototelegraphy	E.320–E.329
ISDN PROVISIONS CONCERNING USERS	E.330–E.349
INTERNATIONAL ROUTING PLAN	**E.350–E.399**
NETWORK MANAGEMENT	
International service statistics	E.400–E.409
International network management	E.410–E.419
Checking the quality of the international telephone service	E.420–E.489
TRAFFIC ENGINEERING	
Measurement and recording of traffic	E.490–E.505
Forecasting of traffic	E.506–E.509
Determination of the number of circuits in manual operation	E.510–E.519
Determination of the number of circuits in automatic and semi-automatic operation	E.520–E.539
Grade of service	E.540–E.599
Definitions	E.600–E.649
Traffic engineering for IP-networks	E.650–E.699
ISDN traffic engineering	E.700–E.749
Mobile network traffic engineering	E.750–E.799
QUALITY OF TELECOMMUNICATION SERVICES: CONCEPTS, MODELS, OBJECTIVES AND DEPENDABILITY PLANNING	
Terms and definitions related to the quality of telecommunication services	E.800–E.809
Models for telecommunication services	E.810–E.844
Objectives for quality of service and related concepts of telecommunication services	E.845–E.859
Use of quality of service objectives for planning of telecommunication networks	E.860–E.879
Field data collection and evaluation on the performance of equipment, networks and services	E.880–E.899

For further details, please refer to the list of ITU-T Recommendations.

ITU-T Recommendation E.360.5

QoS routing and related traffic engineering methods –Transport routing methods

Summary

The E.360.x series of Recommendations describes, analyzes, and recommends methods which control a network's response to traffic demands and other stimuli, such as link failures or node failures. The functions discussed, and recommendations made, related to traffic engineering (TE), are consistent with the definition given in the Framework document of the Traffic Engineering Working Group (TEWG) within the Internet Engineering Task Force (IETF):

Internet Traffic Engineering is concerned with the performance optimization of operational networks. It encompasses the measurement, modelling, characterization, and control of Internet traffic, and the application of techniques to achieve specific performance objectives, including the reliable and expeditious movement of traffic through the network, the efficient utilization of network resources, and the planning of network capacity.

The methods addressed in the E.360.x series include call and connection routing, QoS resource management, routing table management, dynamic transport routing, capacity management, and operational requirements. Some of the methods proposed herein are also addressed in, or are closely related to, those proposed in ITU-T Recs E.170 to E.179 and E.350 to E.353 for routing, E.410 to E.419 for network management and E.490 to E.780 for other traffic engineering issues.

The recommended methods are meant to apply to IP-based, ATM-based, and TDM-based networks, as well as the interworking between these network technologies. Essentially, all of the methods recommended are already widely applied in operational networks worldwide, particularly in PSTN networks employing TDM-based technology. However, these methods are shown to be extensible to packet-based technologies, that is, to IP-based and ATM-based technologies, and it is important that networks which evolve to employ these packet technologies have a sound foundation of methods to apply. Hence, it is the intent that the methods recommended in this series of Recommendations be used as a basis for requirements for specific methods, and, as needed, for protocol development in IP-based, ATM-based, and TDM-based networks to implement the methods.

The methods encompassed in this Recommendation include traffic management through control of routing functions, which include QoS resource management. Results of analysis models are presented which illustrate the tradeoffs between various approaches. Based on the results of these studies as well as established practice and experience, methods are recommended for consideration in network evolution to IP-based, ATM-based, and/or TDM-based technologies.

Source

ITU-T Recommendation E.360.5 was prepared by ITU-T Study Group 2 (2001-2004) and approved under the WTSA Resolution 1 procedure on 16 May 2002.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

© ITU 2003

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

CONTENTS

Introduction

This Recommendation describes and analyzes transport network architectures in light of evolving technology for integrated broadband networks. Dynamic transport routing offers advantages of simplicity of design and robustness to load variations and network failures. Dynamic transport routing can combine with dynamic traffic routing to shift transport bandwidth among node pairs and services through use of flexible transport switching technology. Dynamic transport routing can provide automatic link provisioning, diverse link routing, and rapid link restoration for improved transport capacity utilization and performance under stress.

We present reliable transport routing models to achieve reliable network design, so as to provide service for predefined restoration objectives for any transport link or node failure in the network and continue to provide connections to customers with essentially no perceived interruption of service. We show that robust routing techniques such as dynamic traffic routing, multiple ingress/egress routing, and logical link diversity routing improve response to node or transport failures.

Cross-connect devices, such as optical cross-connects (OXCs), are able to node transport channels, for example OC48 channels, onto different higher-capacity transport links such as an individual WDM channel on a fiberoptic cable. Transport paths can be rearranged at high speed using OXCs, typically within tens of milliseconds switching times. These OXCs can reconfigure logical transport capacity on demand, such as for peak day traffic, weekly redesign of link capacity, or emergency restoration of capacity under node or transport failure. Rearrangement of logical link capacity involves reallocating both transport bandwidth and node terminations to different links. OXC technology is amenable to centralized traffic management.

There is recent work in extending MPLS control capabilities to the setup of layer 2 logical links through OXCs, this effort dubbed multiprotocol lambda switching, after the switching of wavelengths in dense wavelength division multiplexing (DWDM) technology [ARDC99].

ITU-T Recommendation E.360.5

QoS routing and related traffic engineering methods –Transport routing methods

1 Scope

The E.360.x series of Recommendations describes, analyzes, and recommends methods which control a network's response to traffic demands and other stimuli, such as link failures or node failures. The functions discussed, and recommendations made, related to traffic engineering (TE) are consistent with the definitions given in the Framework document of the Traffic Engineering Working Group (TEWG) within the Internet Engineering Task Force (IETF):

Internet Traffic Engineering is concerned with the performance optimization of operational networks. It encompasses the measurement, modelling, characterization, and control of Internet traffic, and the application of techniques to achieve specific performance objectives, including the reliable and expeditious movement of traffic through the network, the efficient utilization of network resources, and the planning of network capacity.

The methods addressed in the E.360.x series include call and connection routing, QoS resource management, routing table management, dynamic transport routing, capacity management, and operational requirements. Some of the methods proposed herein are also addressed in, or are closely related to, those proposed in ITU-T Recs E.170 to E.179 and E.350 to E.353 for routing, E.410 to E.419 for network management and E.490 to E.780 for other traffic engineering issues.

The recommended methods are meant to apply to IP-based, ATM-based, and TDM-based networks, as well as the interworking between these network technologies. Essentially all of the methods recommended are already widely applied in operational networks worldwide, particularly in PSTN networks employing TDM-based technology. However, these methods are shown to be extensible to packet-based technologies, that is, to IP-based and ATM-based technologies, and it is important that networks which evolve to employ these packet technologies have a sound foundation of methods to apply. Hence, it is the intent that the methods recommended in this series of Recommendations be used as a basis for requirements for specific methods, and, as needed, for protocol development in IP-based, ATM-based, and TDM-based networks to implement the methods.

Hence the methods encompassed in this series of Recommendations include:

- traffic management through control of routing functions, which include call routing (number/name translation to routing address), connection routing, QoS resource management, routing table management, and dynamic transport routing.
- capacity management through control of network design, including routing design.
- operational requirements for traffic management and capacity management, including forecasting, performance monitoring, and short-term network adjustment.

Results of analysis models are presented which illustrate the tradeoffs between various approaches. Based on the results of these studies as well as established practice and experience, TE methods are recommended for consideration in network evolution to IP-based, ATM-based, and/or TDM-based technologies.

2 References

See clause 2 of ITU-T Rec. E.360.1.

3 Definitions

See clause 3 of ITU-T Rec. E.360.1.

4 Abbreviations

See clause 4 of ITU-T Rec. E.360.1.

5 Dynamic transport routing principles

An important element of network architecture is the relationship between the transport network and the traffic network. An illustration of a transport network is shown in Figures 1 and 2 illustrates the mapping of layer-2 logical links in the traffic network onto the layer-1 physical transport network of Figure 1. Some logical links overlay two or more fiber-backbone links. For example, in Figure 1, logical link AD traverses fiber-backbone links AB, BC, and CD.

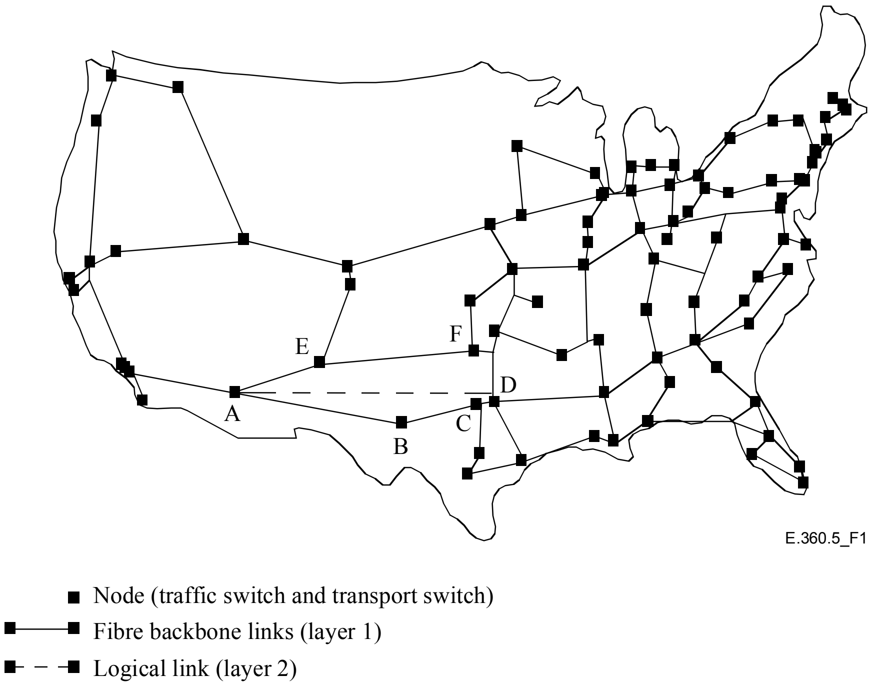

Figure 1/E.360.5 – Transport network model

Figure 2 further illustrates the difference between the physical transport network (layer 1) and the logical transport network (layer 2). Logical links are individual logical connections between network nodes, which make up the logical link connections and are routed on the physical transport network. Logical links can be provisioned at given rates, such as T1, OC3, OC12, OC48, OC192, etc., and is dependent on the level of traffic demand between nodes.

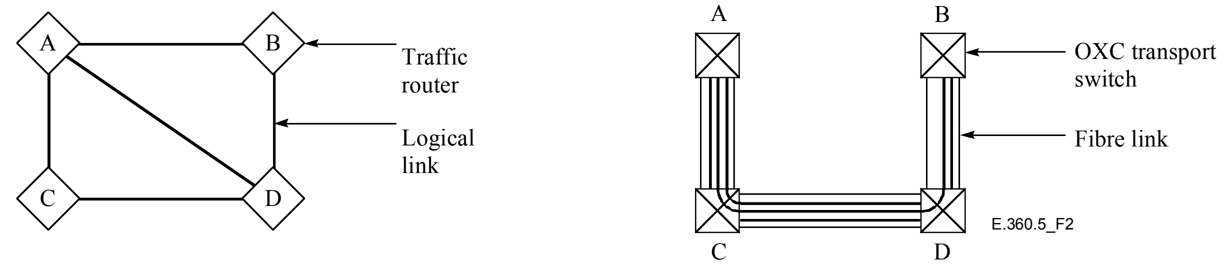
Logical (layer 2) transport network view

Physical (layer 1) transport network view

Figure 2/E.360.5 – Logical (layer 2) and physical (layer 1) transport networks

Figure 2 indicates that a direct logical link is obtained by cross-connecting through a transport switching location. Thus, the traffic network is a logical network overlaid on a sparse physical one. A cross-connect device is traversed at each network node on a given logical link path over layer-1 physical transport links, as illustrated in Figure 2. This is particularly promising when such a device has low cost.

It is clear from Figures 1 and 2 that in a highly interconnected traffic network, or logical transport network, many node pairs may have a "direct" logical link connection where none exists in the physical transport network. In this case, a direct logical link is obtained by cross-connecting through a transport switching location, such as an OXC. This is distinct from the traffic routing situation in which a bearer connection is actually switched at an intermediate location. This distinction between cross-connecting and switching is a bit subtle, but it is fundamental to traffic routing of calls and transport routing of logical links. Referring to Figure 2, we illustrate one of the logical inconsistencies we encounter when we design the traffic network to be essentially separate from the transport network. On the alternative traffic path from node B to node D through A, the physical path is, in fact, up and back from B to A (a phenomenon known as "backhauling") and then across from B to D. The sharing of capacity by various traffic loads in this way actually increases the efficiency of the network because the backhauled capacity to and from B and A is only used when no direct A-to-B or A-to-D traffic wants to use it. It is conceivable that under certain conditions, capacity could be put to more efficient use, and this is studied in this Recommendation.

Hence, a logical link connection is obtained by cross-connecting through transport switching devices, such as OXCs, and this is distinct from per-flow routing, which switches a call on the logical links at each node in the call path. In this way, the logical transport network is overlaid on a sparser physical transport network. In ITU-T Rec. E.360.2 we discussed a wide variety of dynamic traffic routing methods. Dynamic transport routing methods incorporate dynamic path selection which seeks out and uses idle network capacity by using frequent, perhaps call-by-call, traffic and transport routing table update decisions. The trend in both traffic and transport routing architecture is toward greater flexibility in resource allocation, which includes transport and switching resource allocation. A fixed transport routing architecture may have dynamic traffic routing but fixed transport routing of logical link capacity. In a dynamic transport routing architecture, however, the logical link capacities can be rapidly rearranged that is, they are not fixed.

With dynamic transport routing, the logical transport bandwidth is shifted rapidly at layer 2 among node pairs and services through the use of dynamic cross-connect devices. In this case, the layer-1 physical fiber-link bandwidth is allocated among the layer-2 logical links. Bandwidth allocation at layer 3 also creates the equivalent of direct links, and we refer to these links as traffic trunks, which in turn comprise virtual networks (VNETs) as described in ITU-T Rec. E.360.3. Traffic trunks can be implemented, for example, by using MPLS label switched paths (LSPs). Bandwidth is allocated to traffic trunks in accordance with traffic demands, and normally not all logical link bandwidth is assigned; thus, there is a pool of unassigned bandwidth. In cases of traffic overload for a given node pair, the node first sets up calls on the traffic trunk that connects the node pair. If that is not possible, the node then sets up calls on the available pool of bandwidth. If there is available bandwidth, then the bandwidth is allocated to the traffic trunk and used to set up the call. If bandwidth is not available, then the layer-2 logical link bandwidth might be dynamically increased by the bandwidth broker, and then allocated to the traffic trunk and finally, the call. In a similar manner, in the event that bandwidth is underutilized in a traffic trunk, excess bandwidth is released to the available pool of bandwidth and then becomes available for assignment to other node pairs. If logical link bandwidth is sufficiently underutilized, the bandwidth might be returned to the available pool of layer-1 fiber-link bandwidth. The bandwidth broker reassigns network resources on a dynamic basis, through analysis of traffic data collected from the individual nodes.

In the dynamic transport architecture, we allow logical link between the various nodes to be rearranged rapidly, such as by hour of the day, or perhaps in real time. Dynamic transport routing capability enables rearrangement of the logical link capacities on demand. This capability appears most desirable for use in relatively slow rearrangement of capacity, such as for busy-hour traffic, weekend traffic, peak-day traffic, weekly redesign of logical link capacities, or for emergency restoration of capacity under node or transport failure. At various times, the demands for node and transport capacity by the various node pairs and services that ride on the same optical fibers will differ. In this network, if a given demand for logical link capacity between a certain node pair decreases and a second goes up, we allow the logical link capacity to be reassigned to the second node pair. The ability to rearrange logical link capacity dynamically and automatically results in cost savings. Large segments of bandwidth can be provided on fiber routes, and then the transport capacity can be allocated at will with the rearrangement mechanism. This ability for simplified capacity management is discussed further in ITU-T Rec. E.360.6.

Figure 3 illustrates the concept of dynamic traffic (layer 3) and transport routing (layer 2) from a generalized switching node point of view.

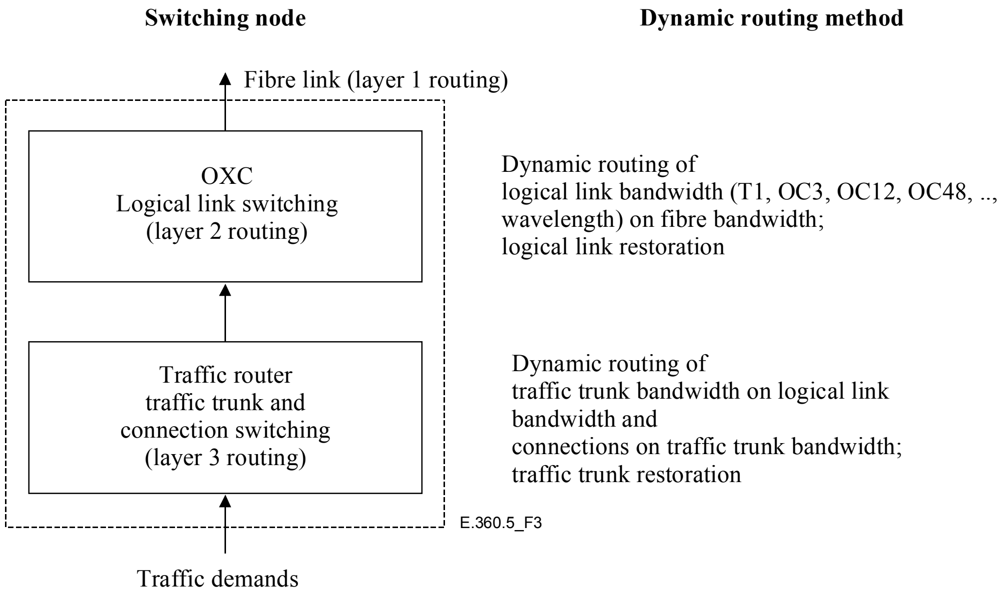

Figure 3/E.360.5 – Dynamic tranport (layer 2) routing and dynamic connection (layer 3) routing

Figure 3 illustrates the relationship of the call-level and transport-level dynamic routing methods used in the dynamic transport routing network. Dynamic connection routing, such as discussed in ITU-T Rec. E.360.2, is used to route calls comprising the underlying traffic demand. Traffic trunk capacity allocations are made for each VNET on the transport link capacity. For each call the originating node analyzes the called number and determines the terminating node, class-of-service, and virtual network. The originating node tries to set up the call on the traffic trunk to the terminating node and, if unavailable, dynamic routing is used at to rearrange the traffic trunk capacity as required to match the traffic demands and to achieve inter-node diversity, access diversity, and traffic trunk restoration following node, OXC, or fiber transport failures. The traffic trunk capacities are allocated by the traffic router to the logical link bandwidth, and the logical link bandwidth allocated by the bandwidth broker to the fiber-link bandwidth, such that the bandwidth is efficiently used according to the level of traffic demand between the nodes.

At the traffic demand level in the transmission hierarchy, flow requests are switched using dynamic traffic routing on the logical link network by node routing logic. At the OC3 and higher demand levels in the transmission hierarchy, logical link demands are switched using OXC systems, which allow dynamic transport routing to route transport demands in accordance with traffic levels. Real-time logical link and real-time response to traffic congestion can be provided by OXC dynamic transport routing to improve network performance.

As illustrated in Figure 4, the dynamic transport routing network concept includes backbone routers (BRs), access routers (ARs), and OXCs. Access routers could route traffic from local offices, access tandems, customer premises equipment, and overseas international switching centers. Here a logical link transmission channel could consist, for example, of OC3-, OC12-, OC48-, or OCx-level bandwidth allocation. An OXC can cross-connect (or "switch") a logical link transmission channel within one terminating fiber wavelength channel in a dense wavelength division multiplexing (DWDM) system to a like-channel within another fiber DWDM system. In the example illustrated, access routers connect to the OXC by means of transport links such as link AX1, and BRs connect to OXCs by means of transport links such as BX1. A number of backbone fiber/DWDM transport links interconnect the OXC network elements, such as links XX1 and XX2. Backbone logical links are terminated at each end by OXCs and are routed over fiber/DWDM spans on the physical transport network on the shortest physical paths. Inter-BR logical links are formed by cross-connecting the bandwidth channels through OXCs between a pair of BRs.

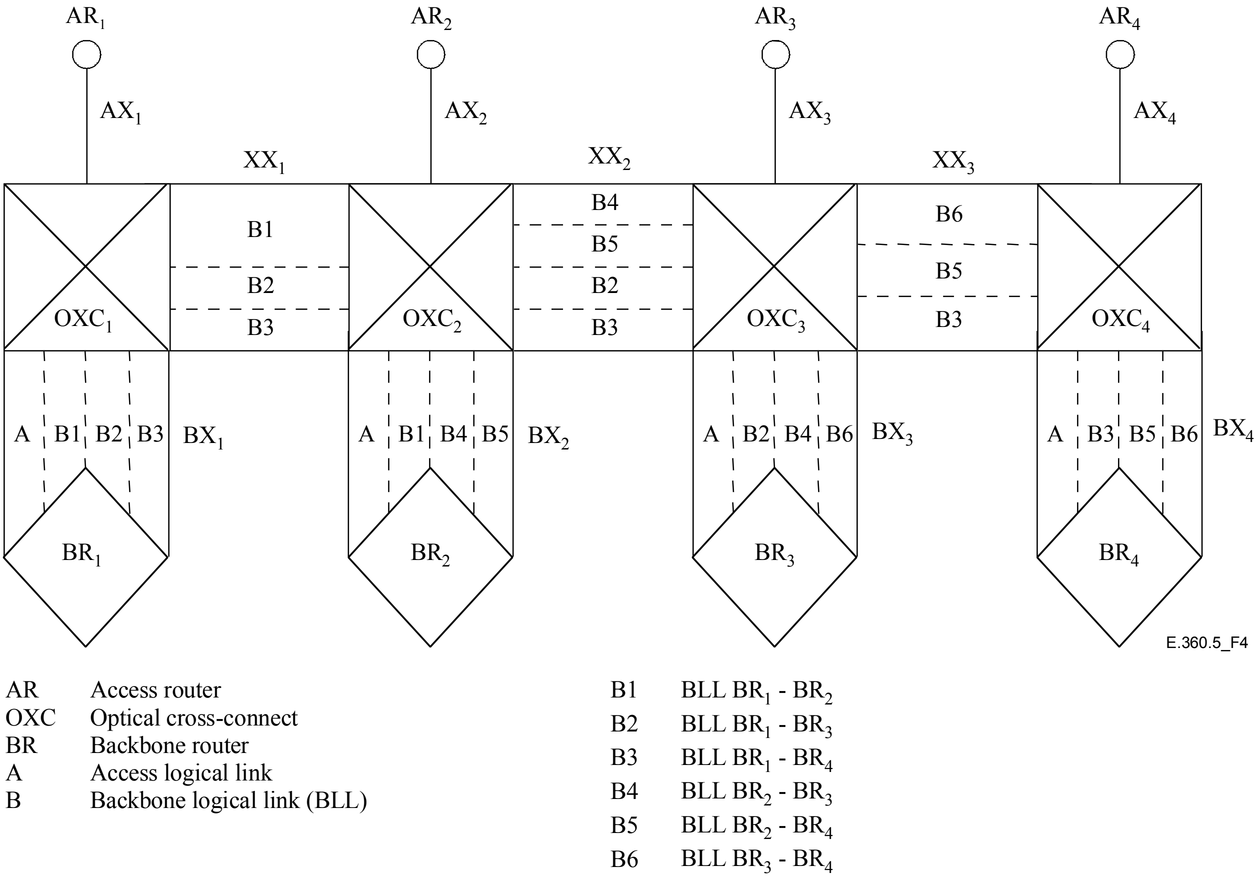

Figure 4/E.360.5 – Dynamic transport routing network

For example, the backbone logical link B2 from BR1 to BR3 is formed by connecting between BR1 and BR3 through fiber/DWDM links BX1, XX1, XX2, and BX3 by making appropriate cross-connects through OXC1, OXC2, and OXC3. Logical links have variable bandwidth capacity controlled by the bandwidth broker implementing the dynamic transport routing network. Access logical links are formed by cross-connecting between ARs and BRs, for example, access router AR1 connected on fiber/DWDM links AX1 and BX1 through OXC1 to BR1 or, alternatively, access router AR1 connected on fiber/DWDM links AX1, XX1, and BX2 cross-connected through OXC1 and OXC2 to BR2. For additional network reliability, backbone routers and access routers may be dual-homed to two OXCs, possibly in different building locations.

6 Dynamic transport routing examples

There are significant network design opportunities with dynamic transport routing and, in this clause, we give examples of dynamic transport routing over different time scales. These examples illustrate the network efficiency and performance improvements possible with seasonal, weekly, daily, and real-time transport rearrangement.

An illustration of dynamic transport routing for varying seasonal traffic demands is given in Figure 5. As seasonal demands shift, the dynamic transport network is better able to match demands to routed transport capacity, thus gaining efficiencies in transport requirements. Figure 5 illustrates how dynamic transport routing achieves network capacity reductions, and shows how transport demand is routed according to varying seasonal requirements. As seasonal demands shift, the dynamic transport network is better able to match demands to routed transport capacity, thus gaining efficiencies in transport requirements. Figure 5 illustrates the variation of winter and summer capacity demands. With fixed transport routing, the maximum termination capacity and transport capacity are provided across the seasonal variations because, in a manual environment without dynamic transport rearrangement, it is not possible to disconnect and reconnect capacity on such short cycle times. When transport rearrangement is automated with dynamic transport routing, however, the termination and transport design can be changed on a weekly, daily, or, with high-speed packet switching, real-time basis to exactly match the termination and transport design with the actual network demands. Notice that in the fixed transport network, there is unused termination and transport capacity that cannot be used by any demands; sometimes this is called "trapped capacity" because it is available, but cannot be accessed by any actual demand. The dynamic transport network, in contrast, follows the capacity demand with flexible transport routing and, together with transport network design, it reduces the trapped capacity.

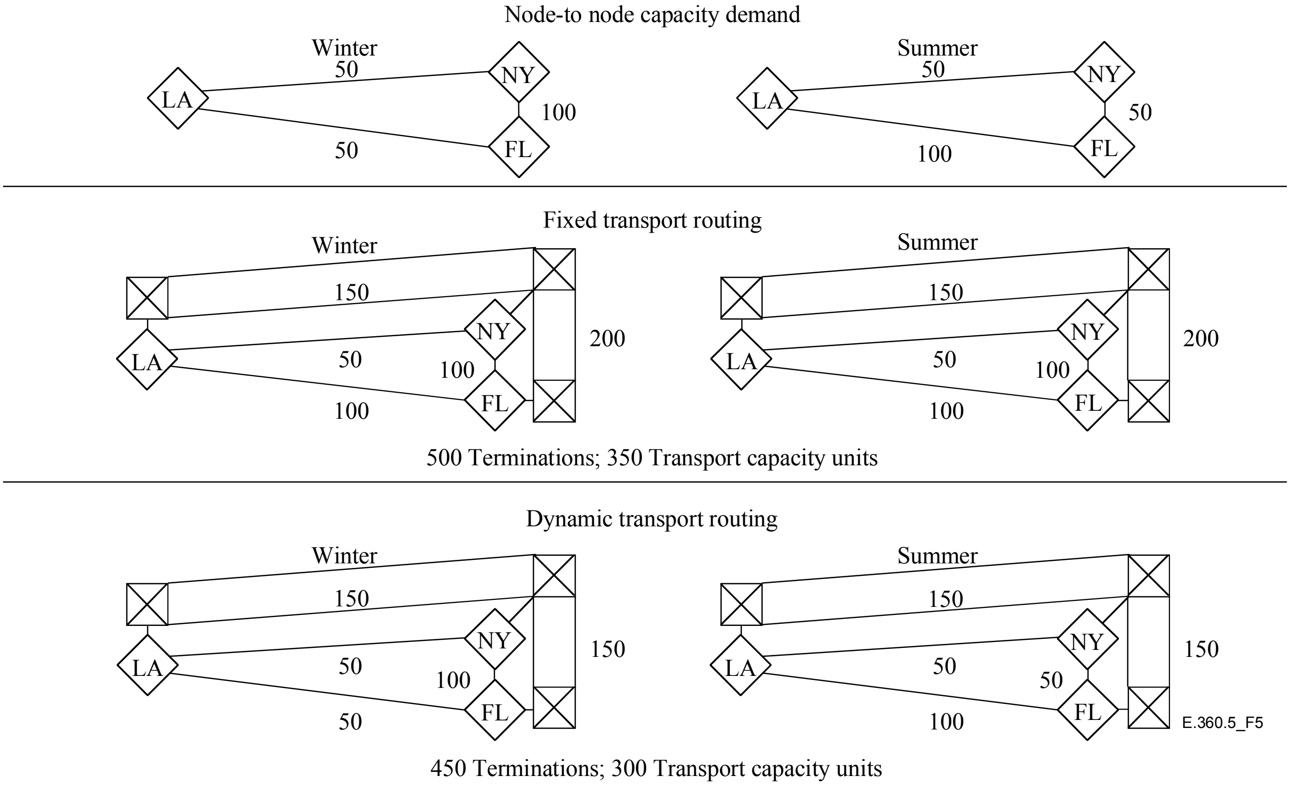

Figure 5/E.360.5 – Dynamic transport routing vs. fixed transport routing

Therefore, the variation of demands leads to capacity-sharing efficiencies which, in the example of Figure 5, reduce termination capacity requirements by 50 node terminations, or approximately 10 percent compared with the fixed transport network, and by 50 transport capacity requirements, or approximately 14 percent. Therefore, with dynamic transport routing capacity utilization can be made more efficient in comparison with fixed transport routing because, with dynamic transport network design, the link sizes can be matched to the network load.

With dynamic traffic routing and dynamic transport routing design models, reserve capacity can be reduced in comparison with fixed transport routing. In-place capacity that exceeds the capacity required to exactly meet the design loads with the objective performance is called reserve capacity. Reserve capacity comes about because load uncertainties, such as forecast errors, tend to cause capacity buildup in excess of the network design that exactly matches the forecast loads. Reluctance to disconnect and rearrange traffic trunk and transport capacity contributes to this reserve capacity buildup. Typical ranges for reserve capacity are from 15 to 25 percent or more of network cost. Models show that dynamic traffic routing compared with fixed traffic routing provides a potential 5 percent reduction in reserve capacity while retaining a low level of short-term capacity design [A98].

With dynamic transport network design, the link sizes can be matched to the network load. With dynamic transport routing, the link capacity disconnect policy becomes, in effect, one in which link capacity is always disconnected when not needed for the current traffic loads. Models given in [FHH79] predict reserve capacity reductions of 10% or more under this policy, and the results presented in Annex A, based on weekly dynamic transport design, substantiate this conclusion.

Weekly design and rearrangement of logical link capacity can approach zero reserve capacity designs. Figures 6 and 7 illustrate the changing of routed transport capacity on a weekly basis between node pairs A-B, C-D, and B-E, as demands between these node pairs change on a weekly basis.

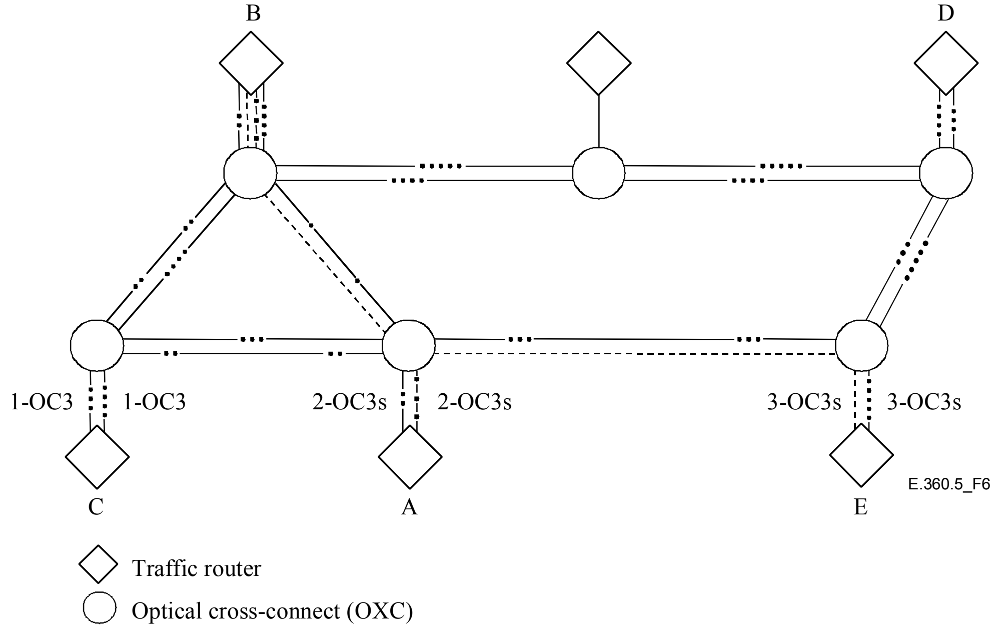

Figure 6/E.360.5 – Dynamic transport routing network weekly arrangement (week 1 load pattern)

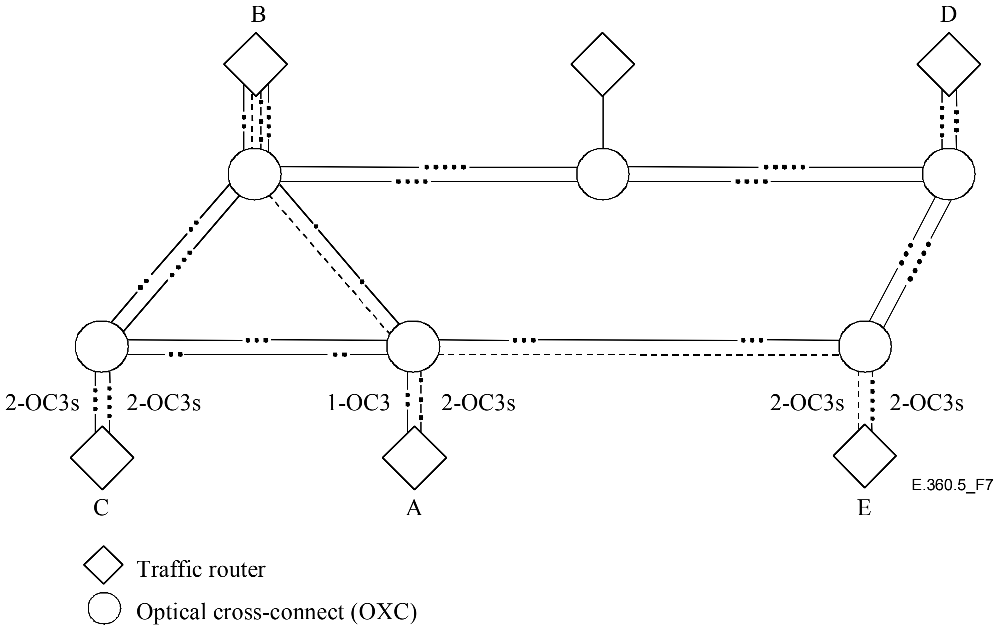

Figure 7/E.360.5 – Dynamic transport routing network weekly arrangement (week 2 load pattern)

These transport routing and capacity changes are made automatically in the dynamic transport network in which diverse transport routing of logical links A-B and C-D is maintained by the dynamic transport routing network. Logical link diversity achieves additional network reliability.

Daily design and rearrangement of transport link capacity can achieve performance improvements for similar reasons, due to noncoincidence of transport capacity demands that can change daily. An example is given in Figures 8 and 9 for traffic noncoincidence experienced on peak days such as Christmas Day. In Figure 8, we illustrate the normal business-day routing of access demands and inter-BR demands. On Christmas Day, however, there are many busy nodes and many idle nodes. For example, node BR2 may be relatively idle on Christmas Day (for example, if it were a downtown business node), while BR1 may be very busy. Therefore, on Christmas Day, BR2 demands to everywhere else in the network are reduced, and through dynamic transport routing these transport capacity reductions can be made automatically. Similarly, BR1 demands are increased on Christmas Day. Access demands such as those from AR1 can be redirected to freed-up termination capacity on BR2, as illustrated in Figure 9, which also frees up termination capacity on BR1 to be used for inter-BR demand increases. By this kind of access demand and inter-BR demand rearrangement, based on noncoincident traffic shifts, more traffic to and from BR1 can be completed because inter-BR logical link capacity is increased, now using freed-up transport capacity from the reduction in the transport capacity needed by BR2. On a peak day, such as Christmas Day, the busy nodes are often limited by inter-BR logical link capacity; this rearrangement reduces or eliminates this bottleneck, as is illustrated in the Christmas Day dynamic transport network design example in Annex A.

The balancing of access and inter-BR capacity throughout the network can lead to robustness to unexpected load surges. This load-balancing design is illustrated in Annex A with an example based on a Hurricane-caused focused overload in the northeastern United States. Capacity addition rearrangements based on instantaneous reaction to unforeseen events such as earthquakes could be made in the dynamic transport network.

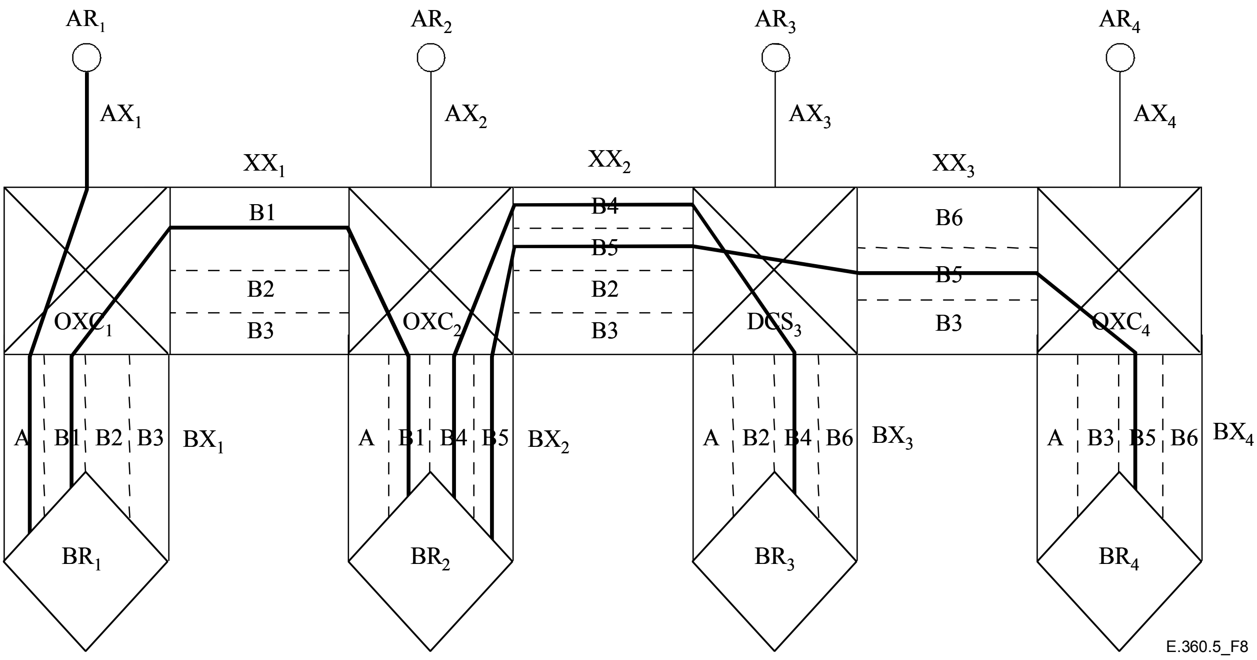

AR	Access router
OXC	Optical cross-connect
DCS	Digital cross-connect
BR	Backbone router
A	Access logical link
B	Backbone logical link (BLL)

B1	BLL BR_1 - BR_2
B2	BLL BR_1 - BR_3
B3	BLL BR_1 - BR_4
B4	BLL BR_2 - BR_3
B5	BLL BR_2 - BR_4
B6	BLL BR_3 - BR_4

Figure 8/E.360.5 – Dynamic transport routing peak day design

Dynamic transport routing can provide dynamic restoration of failed capacity, such as that due to fiber cuts, onto spare or backup transport capacity. Dynamic transport routing provides a self-healing network capability to ensure a networkwide path selection and immediate adaptation to failure.

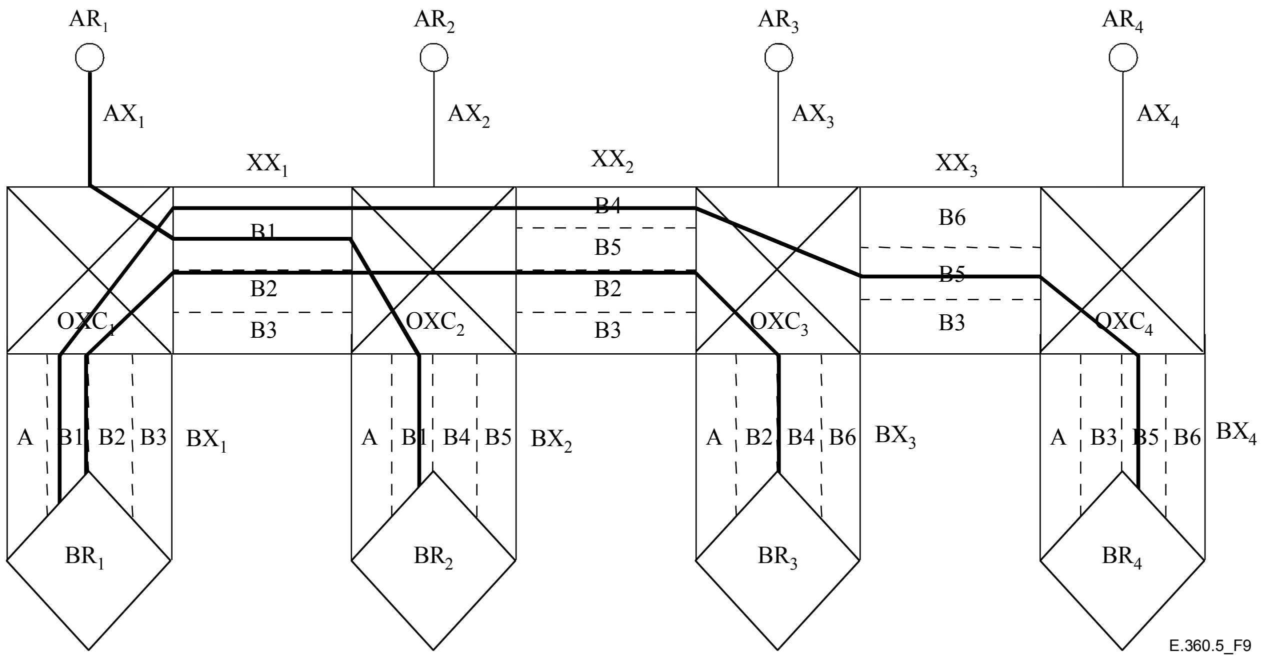

AR	Access router
OXC	Optical cross-connect
BR	Backbone router
A	Access logical link
B	Backbone logical link (BLL)

B1	BLL BR_1 - BR_2
B2	BLL BR_1 - BR_3
B3	BLL BR_1 - BR_4
B4	BLL BR_2 - BR_3
B5	BLL BR_2 - BR_4
B6	BLL BR_3 - BR_4

Figure 9/E.360.5 – Dynamic transport routing peak day design

FASTAR [CED91], for example, implements central automatic control of transport switching devices to quickly restore service following a transport failure. As illustrated in Figure 10, a fiber cut can disrupt large traffic trunk capacities, and dynamic transport restoration can quickly restore transport capacity. Dynamic transport routing provides a self-healing network capability to ensure a networkwide path selection and immediate adaptation to failure. As illustrated in Figure 10, a fiber cut near the Nashville node severed 8.576 Gbit/s of traffic trunk capacity of switched-network traffic (there was also private-line traffic), and after dynamic transport restoration a total of 3.84 Gbit/s of traffic trunk capacity was still out of service in the switched network. In the example, dynamic transport restoration is implemented by centralized automatic control of transport cross-connect devices to quickly restore service following a transport failure, such as caused by a cable cut. Over the duration of this event, more than 12 000 calls were blocked in the switched network, almost all of them originating or terminating at the Nashville node, and it is noteworthy that the blocking in the network returned to zero after the 4.736 Gbit/s of traffic trunk capacity was restored in the first 11 minutes, even though there was still 3.84 Gbit/s of traffic trunk capacity still out of service.

Dynamic traffic routing was able to find paths on which to complete traffic even though there was far less logical link capacity than normal, even after the dynamic transport restoration. Hence, dynamic traffic routing, in combination with dynamic transport restoration, provides a self-healing network capability, and even though the cable was repaired two hours after the cable cut, degradation of service was minimal. In this example, dynamic traffic routing also provided priority routing for selected customers and services, as described in ITU-T Rec. E.360.3, which permits priority calls to be routed in preference to other calls, and blocking of the priority services is essentially zero throughout the whole event.

Over the duration of an event, calls are blocked until sufficient capacity is restored for the network to return to zero blocking. That is, both dynamic transport routing and dynamic traffic routing are able to find available paths on which to restore the failed traffic. Hence, this example clearly illustrates how real-time dynamic traffic routing in combination with real-time dynamic transport routing can provide a self-healing network capability, and even if the cable is repaired two hours after the cut, degradation of service is minimal. This improved network performance provides additional service revenues as formerly blocked calls are completed, and it improves service quality to the customer.

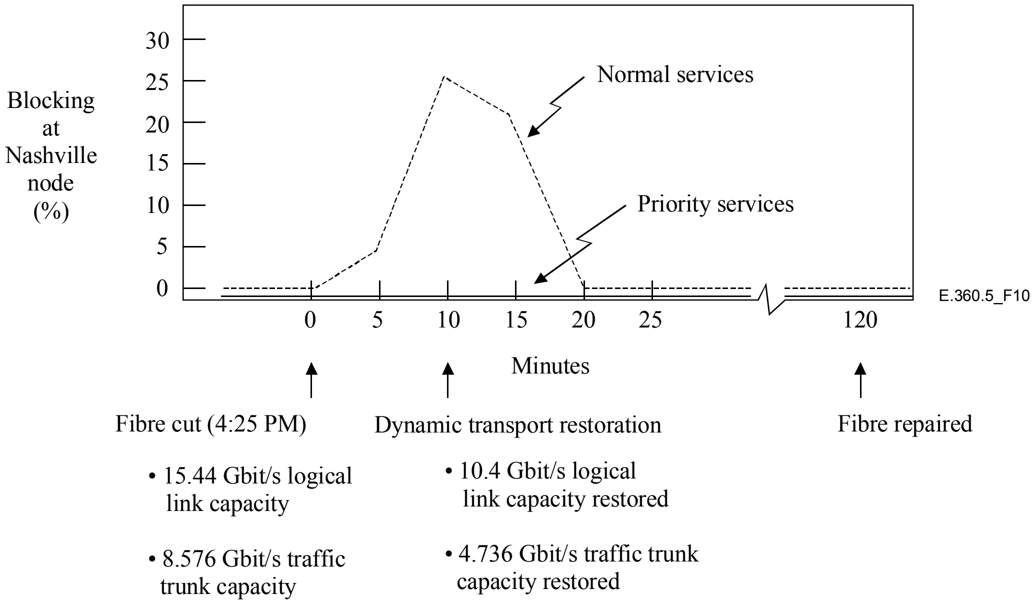

Figure 10/E.360.5 – Fibre cut example with dynamic traffic routing and dynamic transport routing

These examples illustrate that implementation of dynamic transport routing provides better network performance at reduced cost. These benefits are similar to those achieved by dynamic traffic routing and, as shown, the combination of dynamic traffic and transport routing provides synergistic reinforcement to achieve these network improvements.

The implementation of a dynamic transport routing network allows significant reductions in capital costs and network management and design expense with rearrangeable transport capacity design methods. Automated logical link provisioning and rearrangement lead to annual operations expense savings. Other network management and design impacts, leading to additional reduction in operations expense, are to simplify logical link provisioning systems; automate preservice logical-link testing, and simplify maintenance systems; integrate logical-link capacity forecasting, administration, and bandwidth allocation into capacity planning and delivery; simplify node and transport planning; and automate inventory tracking.

7 Reliable transport network design

In the event of link, node, or other network failure, the network design needs to provide sufficient surviving capacity to meet the required performance levels. For example, if a major fiber link fails, it could have a catastrophic effect on the network because traffic for many node pairs could not use the failed link. Similarly, if one of the nodes fails, it could isolate a whole geographic area until the node is restored to service.

With these two kinds of major failures in mind, we present here reliable transport routing models to achieve reliable network design, so as to provide service for predefined restoration objectives for any transport link or node failure in the network, and continue to provide connections to customers with essentially no perceived interruption of service. This approach tries to integrate capabilities in both the traffic and transport networks to make the network robust, or insensitive to failure. The basic aims of these models are to provide link diversity and protective capacity augmentation where needed so that specific "network robustness" objectives, such as traffic restoration level objectives, are met under failure events. This means that the network is designed so that it carries at least the fraction of traffic known as the traffic restoration level (TRL) under the failure event. For example, a traffic restoration level objective of 70% means that under any single transport link failure in the transport network, at least 70% of the original traffic for any affected node pair is still carried after the failure; for the unaffected node pairs, the traffic is carried at the normal blocking probability grade-of-service objective. These design models provide designs that address the network response immediately after a network event. It is also desirable to have transport restoration respond after the occurrence of the network event to bring service back to normal. Transport restoration is also addressed in this Recommendation.

Reliable network performance objectives may require, for example, the network to carry 50% of its busy-hour load on each link within five minutes after a major network failure, in order to eliminate isolations among node pairs. Such performance may be provided through traffic restoration techniques, which include link diversity, traffic restoration capacity and dynamic traffic routing. Reliable network performance objectives might also require a further reduction of connection setup blocking level to less than 5% within, say, 30 minutes to limit the duration of degraded service. This is possible through transport restoration methods that utilize transport nodes along with centralized transport restoration control. A further objective may be to restore at least 50% of severed trunks in affected links within this time period.

The transport restoration process restores capacity for switched, as well as dedicated ("private-line"), services in the event of link failures. In one implementation, transport restoration is conducted via a centralized system that restores the affected transport capacity until all available restoration capacity is exhausted. Optimization of the total cost of transport restoration capacity is possible through a design that increases sharing opportunities of the restoration capacity among different failure scenarios. Real-time transport restoration may also require the use of dedicated restoration capacity for each link and, thus, a lesser opportunity for sharing the restoration capacity. For the purpose of this analysis, we assume that all network transport may be protected with an objective level of restoration capacity. Transport restoration level (TPRL) is the term used to specify the minimal percentage of capacity on each transport link that is restorable. A transport restoration level is implemented in the model by restoring each affected link in a failure to a specified level. Here, we further distinguish between transport restoration level for switched circuits and dedicated circuits and designate them by TPRLs and TPRLp, respectively.

We now describe logical transport routing design models for survivable networks. Before we describe the models, we discuss the distinction between the traffic and transport networks and the concept of link diversity.

To distinguish between the traffic and transport networks, consider the example of a three-node network in Figure 2. The physical transport network is depicted at the bottom of the figure, and the corresponding logical transport (traffic) network at the top. For example, the direct logical link for connecting nodes A and B may ride the path A-C-D-B in the physical transport network. There is not a logical link between nodes B and C, which means there is no direct traffic trunk capacity from node B to C. A single physical transport link failure may affect more than one logical link. For example, in Figure 2, the failure of the physical transport link C-D affects logical links A-D, C-D, and A-B. Logical link diversity refers to a logical link design in which direct capacity for a logical link is split on two or more different physical transport paths. For example, in Figure 1, the direct logical link capacity for the link A-D may be split on to the two physical transport paths A-B-C-D and a physically diverse path A-E-F-D. A link diversity policy, say, of 70/30 corresponds to the fact that no more than 70% of the direct logical link (traffic trunk) capacity is routed on a single physical transport link for the different transport paths for that logical link. The advantage of logical link diversity is that if a physical transport link fails, the traffic for a particular node pair can still use the direct logical link capacity that survived on the physical transport path not on the failed link.

We now present models for transport routing design for physical transport link failure and node failure.

7.1 Transport link design models

We assume that we have two distinct transport paths for direct logical link capacity for each node pair. In the model, traffic demand is converted to virtual trunk demand, such as based on direct and overflow logical transport capacity, as illustrated in Figure 1 of E.360.6. Let v be the virtual trunk requirement for the traffic demand for a particular node pair. Let d be the virtual trunk capacity to be put on the primary physical transport path and s be the virtual trunk capacity to be put on the alternate physical transport paths for the direct logical link of the given node pair. Let b be the number of trunks for this traffic link that are designed by the network design model. Let t be the traffic restoration level (TRL) objective under a link failure scenario. Let Δ_b be the link capacity augmentation that may be needed for this logical link.

What we would like in a failure event is to carry a portion tv of the total virtual trunk demand for the affected node pairs. Thus, if $tv \le b/2$, we set $\Delta_b = 0$ (no augmentation) with $d = b - tv$ and $s = tv$. In this way, if either transport path fails, we can carry at least tv of the virtual trunk demand. On the other hand, if $tv > b/2$, then we want:

$$(b + \Delta_b)/2 = tv$$

which implies: $\Delta_b = 2tv - b$

In this case, we set $d = s = (b + \Delta_b)/2$. The above procedure is repeated for every demand pair in the network. The incremental cost of the network is the cost of trunk augmentation, and routing the direct logical link capacity for each node pair, if any, on two transport paths.

The above procedure can be extended to the general case of k distinct physical transport paths. So, for k distinct physical transport paths, if:

$$tv \le \frac{k-1}{k} b$$

then $\Delta_b = 0$ with $d = b - tv$ on the first physical transport path, and $tv/(k - 1)$ on each of the other $(k - 1)$ transport paths. If:

$$tv > \frac{k-1}{k} b$$

then: $\Delta_b = \frac{k}{k-1} tv - b$

with each of the k transport paths having $(b + \Delta_b)/k$ trunks.

A transport link model is illustrated in Figure 11, where each transport cross section is assumed on the average to contain certain fractions of switched (N) and dedicated (M) circuits. A portion of the switched and dedicated circuits is further presumed to be restorable in real time, such as with ring or dual-feed transmission arrangements (lowercase n and m values).

Circuits that are not restored in real time are restored with transport restoration to a specified transport restoration level (TPRL) value. The lower part of Figure 11 demonstrates the interaction between the switched and dedicated capacity in the restoration process. The restoration process is assumed to restore the first unit of transport capacity (e.g. OC3) after x seconds, with y seconds to restore each additional unit of transport capacity. The restoration times are illustrative and are not critical to the reliable network design principles being discussed. SONET ring restoration can occur in about 50-200 milliseconds, and such real-time restoration is included in the model. A prioritization method is assumed, whereby transport links that carry higher-priority dedicated services are restored first. Because switched and dedicated capacity is often mixed at the logical link level, some switched capacity is also restored.

Different levels of transport restoration may also be assigned for the dedicated (TPRLp) and switched (TPRLs) networks. Each type of circuit demand is then restored to the corresponding level of restoration. Figure 11 also shows how the restoration level for switched circuits varies as a function of time. Some level of traffic circuits is restored in real time (n). After x seconds, transport restoration is initiated with one unit of transport capacity being restored in each y seconds, and with a smaller fraction of each transport capacity unit being switched traffic. The switched portion in each transport capacity unit subsequently increases to a larger fraction after dedicated traffic is restored to its designated level TPRLp. Transport restoration stops after both the TPRLp and TPRLs objectives are met.

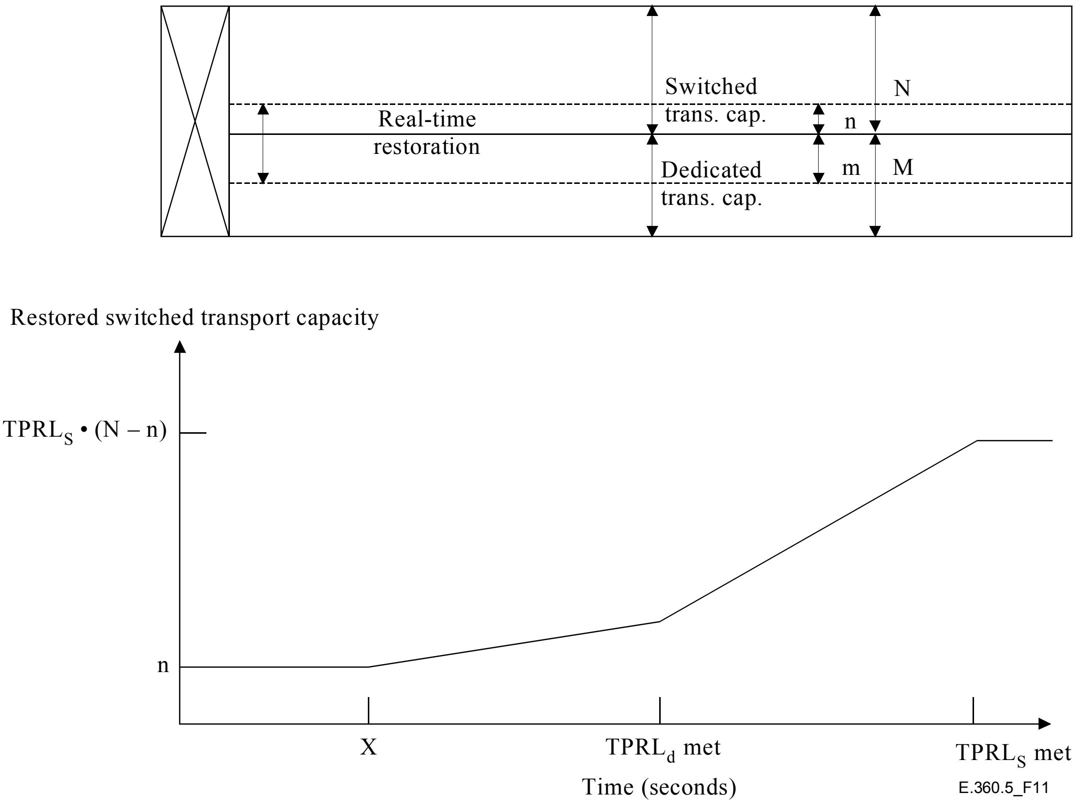

$TPRL_S$ Transport restoration level for switched transport capacity

$TPRL_d$ Transport restoration level for dedicated transport capacity

Figure 11/E.360.5 – Transport restoration model

7.2 Node design models

Node failure restoration design incorporates the concept of dual homing, as discussed in ITU-T Rec. E.360.2, along with multiple ingress/egress routing. With single homing, the traffic from a particular geographical area normally goes to the single node nearest to it in order to carry the ingress and egress traffic. For example, in Figure 12 the traffic from the served area A1 enters the network through node B, and, similarly, the served area A2 traffic enters the network through node C. Now, if node B fails, then area A1 gets isolated. To protect against such an event, areas A1 and A2 are homed to more than one node (to nodes A and B in this case (Figure 12, bottom)). This is the concept of dual homing, in which we address the issue of designing a reliable network when one of the nodes may fail.

For every node to be protected, we assign a dual-homed node. Before failure, we assume that any load from a node to the protected node and its dual-homed node is equally divided; that is, if the original load between area A1 and node A is a1, and between A and the dual-homed node, B, is a2, then we assume that under normal network conditions, the load between nodes A and B is (a1 + a2)/2, and the same for the load between nodes A and C. We refer to this concept as balanced load. Then, under a failure event such as a node B failure, we carry load equal to (a1 + a2)/2 between nodes A and C. (See Figure 12, bottom.) We call this design objective a 50% traffic restoration level objective in a manner quite similar to the link failure event. As we can see from the lower part of Figure 12, this restoration level of traffic from or to area A1 is then still carried.

In the restoration model, the node pairs that are going to be considered for the node failure design scenarios are determined. Next, the dual-homing nodes are determined for each node and the balanced-load traffic routed accordingly. Then the vk virtual trunks are computed for the balanced loads. For each node pair, one of the nodes is assumed to fail (say, node B). Then this node cannot have any incoming or outgoing traffic and, also, cannot be a via node for any two-link traffic between other node pairs. Using these constraints, we solve a linear programming model that minimizes the incremental augmentation cost. Then, we reverse the roles of the nodes for this pair and solve the linear programming model again with the above-mentioned constraints. This design process is repeated for every pair of candidate nodes for each node failure scenario.

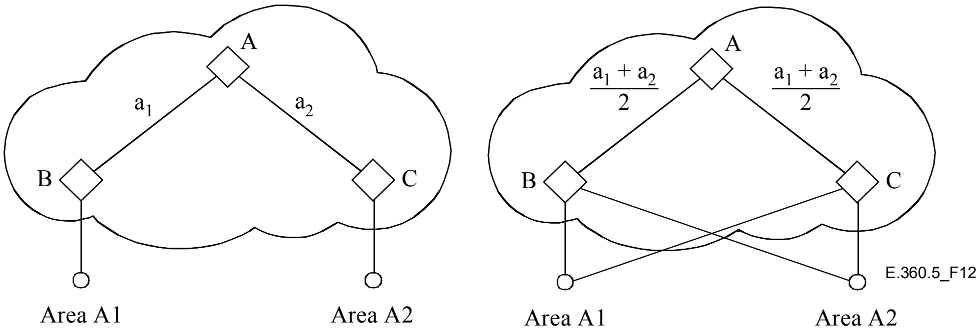

Figure 12/E.360.5 – Illustration of dual homing

8 Conclusions/recommendations

In this Recommendation, we present and analyze dynamic transport network architectures. Dynamic transport routing is a routing and bandwidth allocation method, which combines dynamic traffic routing with dynamic transport routing and for which we provide associated network design methods. We find that networks benefit more in efficiency and performance as the ability to reassign transport bandwidth is increased, and can simplify network management and design. We present results of a number of analysis, design, and simulation studies related to dynamic transport network architectures.

Models are used to measure the performance of the network for dynamic transport routing network design in comparison with the fixed transport network design, under a variety of network conditions including normal daily load patterns, unpredictable traffic load patterns such as caused by a hurricane, known traffic overload patterns such as occur on Christmas Day, and a network failure conditions such as a large fiber cut.

The conclusions/recommendations reached in this Recommendation are as follows:

- Dynamic transport routing is recommended and provides greater network throughput and, consequently, enhanced revenue, and at the same time capital savings should result, as discussed in ITU-T Rec. E.360.6.
  a) Dynamic transport routing network design enhances network performance under failure, which arises from automatic inter-backbone-router and access logical-link diversity in combination with the dynamic traffic routing and transport restoration of logical links.
  b) Dynamic transport routing network design is recommended and improves network performance in comparison with fixed transport routing for all network conditions simulated, which include abnormal and unpredictable traffic load patterns.
- Traffic and transport restoration level design is recommended and allows for link diversity to ensure a minimum level of performance under failure.
- Robust routing techniques are recommended, which include dynamic traffic routing, multiple ingress/egress routing, and logical link diversity routing; these methods improve response to node or transport failures.

Annex A

Modelling of traffic engineering methods

In this annex we give modelling results for dynamic transport routing capacity design, performance under network failure, and performance under various network overload scenarios.

A.1 Dynamic transport routing capacity design

Design for traffic loads with week-to-week traffic variation. Dynamic transport routing network design allows more efficient use of node capacity and transport capacity and can lead to a reduction of network reserve trunk capacity by about 10%, while improving network performance. Table A.1 illustrates a comparative forecast of a national intercity network's normalized logical-link capacity requirements for the base case without dynamic transport routing and the network requirements with dynamic transport routing network design. When week-to-week traffic variations, which reflect seasonal variations, are taken into account, as in this analysis, the dynamic transport routing design can provide a reduction in network reserve capacity. As shown in Table A.1, the traffic trunk savings always exceed 10%, which translates into a significant reduction in capital expenditures.

Table A.1/E.360.5 – Dynamic transport routing capacity savings with week-to-week seasonal traffic variations (normalized capacity)

Forecast period	Capacity fixed transport routing design	Capacity dynamic transport routing design	Capacity savings (%)
Year 1	1.000	0.873	12.7
Year 2	1.048	0.919	12.3
Year 3	1.087	0.968	11.0
Year 4	1.138	1.019	10.4

Dynamic transport routing network design for transport capacity achieves higher fiber link fill rates, which further reduces transport costs. The dynamic transport routing network implements automated inter-BR and access logical-link diversity, logical-link restoration, and node backup restoration to enhance the network survivability over a wide range of network failure conditions. We now illustrate dynamic transport routing network performance under design for normal traffic loads, fiber transport failure events, unpredictable traffic load patterns, and peak-day traffic load patterns.

A.2 Performance for network failures

Simulations are performed for the fixed transport and dynamic transport network performance for a fiber cut in Newark, New Jersey, in which approximately 8.96 Gbit/s of traffic trunk capacity was lost. The results are shown in Table A.2. Here, a threshold of 50 percent or more node-pair blocking is used to identify node pairs that are essentially isolated; hence, the rearrangeable transport network design eliminates all isolations during this network failure event.

Table A.2/E.360.5 – Network performance for fiber cut in Newark, NJ

	Traffic lost/delayed (%)	Number of node pairs with lost/delayed > 50%
Fixed transport routing	14.4	963
Dynamic transport routing	4.2	0

An analysis is also performed for the network performance after transport restoration, in which the fixed and dynamic transport network designs are simulated after 29 percent of the lost trunks are restored. The results are shown in Table A.3. Again, the dynamic transport network design eliminates all network isolations, some of which still exist in the base network after traffic trunk restoration.

Table A.3/E.360.5 – Network performance for fiber cut in Newark, NJ (after logical-link restoration)

	Traffic lost/delayed (%)	Number of node pairs with lost/delayed > 50%
Fixed transport routing	7.0	106
Dynamic transport routing	0.6	0

From this analysis we conclude that the combination of dynamic traffic routing, logical-link diversity design, and transport restoration provides synergistic network survivability benefits. Dynamic transport network design automates and maintains logical-link diversity, as well as access network diversity in an efficient manner, and provides automatic transport restoration after failure.

A network reliability example is given for dual-homing transport demands on various OXC transport nodes. In one example, an OXC failure at the Littleton, MA node, in the model illustrated in Figure 1 is analyzed, and results given in Table A.4. Because transport demands are diversely routed between nodes and dual-homed between access nodes and OXC devices, this provides additional network robustness and resilience to traffic node and transport node failures. When the network is designed for load balancing between access and internode demands, and traffic trunk restoration is performed, the performance of the dynamic transport routing network is further improved.

Table A.4/E.360.5 – Dynamic transport network performance under OXC failure

	Traffic lost/delayed (%)	Number of node pairs with lost/delayed > 50%
Fixed transport routing	4.1	231
Dynamic transport routing, dual-homing	1.3	0
Dynamic transport routing, dual-homing, load balancing, logical-link restoration	0.6	0

Figure A.1 illustrates a typical instance of network design with traffic restoration level objectives and transport restoration level objectives, as compared with the base network with no TRL or TPRL design objectives. In the example, a fiber link failure occurs in the model network f seconds after the beginning of the simulation, severing a large amount of transport capacity in the network and cutting off thousands of existing connections. Therefore, in the simulation results shown in Figure A.1, we see a large jump in the blocking at the instant of the cut. A transient flurry of reattempts follows as cut-off customers redial and reestablish their calls. This call restoration process is aided by the traffic restoration level design which provides link diversity, and protective transport capacity, to meet the TRL objectives immediately following a failure. This TRL design, together with the ability of dynamic traffic routing to find surviving capacity wherever it exists, quickly reduces the transient blocking level which then remains roughly constant for about x seconds until the transport restoration process begins. At x seconds after the link failure, the transport restoration process begins to restore capacity that was lost due to the failure. Blocking then continues to drop during that period when transport restoration takes place until it reaches essentially a level of zero blocking. Figure A.1 illustrates the comparison between network performances with and without the traffic and transport restoration design techniques presented in this Recommendation.

This traffic restoration level design allows for varying levels of diversity on different links to ensure a minimum level of performance. Robust routing techniques such as dynamic traffic routing, multiple ingress/egress routing, and logical link diversity routing further improve response to node or transport failures. Transport restoration is necessary to reduce network blocking to low levels. Given, for example, a 50 percent traffic restoration level design, it is observed that this, combined with transport restoration of 50 percent of the failed transport capacity in affected links, is sufficient to restore the traffic to low blocking levels. Therefore, the combination of traffic restoration level design and transport restoration level design is seen both to be cost-effective and to provide fast and reliable performance. The traffic restoration level design eliminates isolations between node pairs, and transport restoration level is used to reduce the duration of poor service in the network. Traffic restoration techniques combined with transport restoration techniques provide the network with independent means to achieve reliability against multiple failures and other unexpected events and are perceived to be a valuable part of a reliable network design.

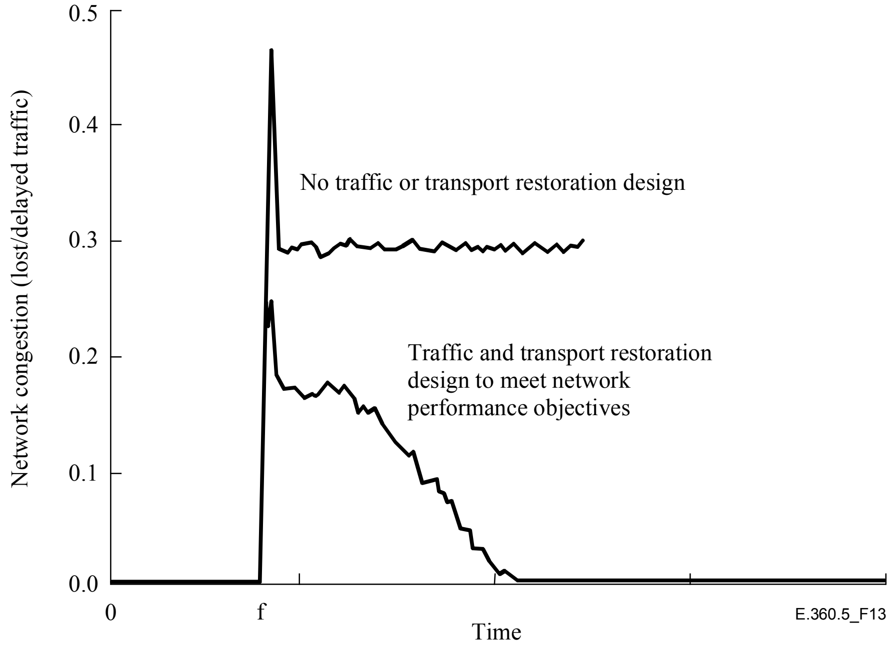

Figure A.1/E.360.5 – Network performance for link failure with traffic and transport restoration design

A.3 Performance for general traffic overloads

The national network model is designed for dynamic transport routing with normal engineered traffic loads using the discrete event flow optimization (DEFO) model described in ITU-T Rec. E.360.6, and it results in a 15 percent savings in reserve trunk capacity over the fixed transport routing model. In addition to this large savings in network capacity, the network performance under a 10 percent overload results in the performance comparison illustrated in Table A.5. Hence, dynamic transport routing network designs achieve significant capital savings while also achieving superior network performance.

Table A.5/E.360.5 – Network performance for 10% traffic overload

	Traffic lost/delayed (%)	**Node pair maximum lost/delayed (%)**
Fixed transport routing	0.11	17.3
Dynamic transport routing	0	0

A.4 Performance for unexpected overloads

Dynamic transport routing network design provides load balancing of node traffic load and logical-link capacity so that sufficient reserve capacity is provided throughout the network to meet unexpected demands on the network. The advantage of such design is illustrated in Table A.6, which compares the simulated network blocking for the fixed transport routing network design and dynamic transport routing network design during an hurricane-caused focused traffic overload in the northeastern United States. Such unexpected focused overloads are not unusual in a switched network, and the additional robustness provided by dynamic transport routing network design to the unexpected traffic overload patterns is clear from these results.

Table A.6/E.360.5 – Network performance for unexpected traffic overload (focused overload in Northeastern US caused by hurricane)

	Traffic lost/delayed (%)	Node pair maximum lost/delayed (%)
Fixed transport routing	0.43	22.7
Dynamic transport routing	0.28	13.3

Another illustration of the benefits of load balancing is given in Figure A.2, in which a 25% traffic overload is focused on a node in Jackson, Mississippi.

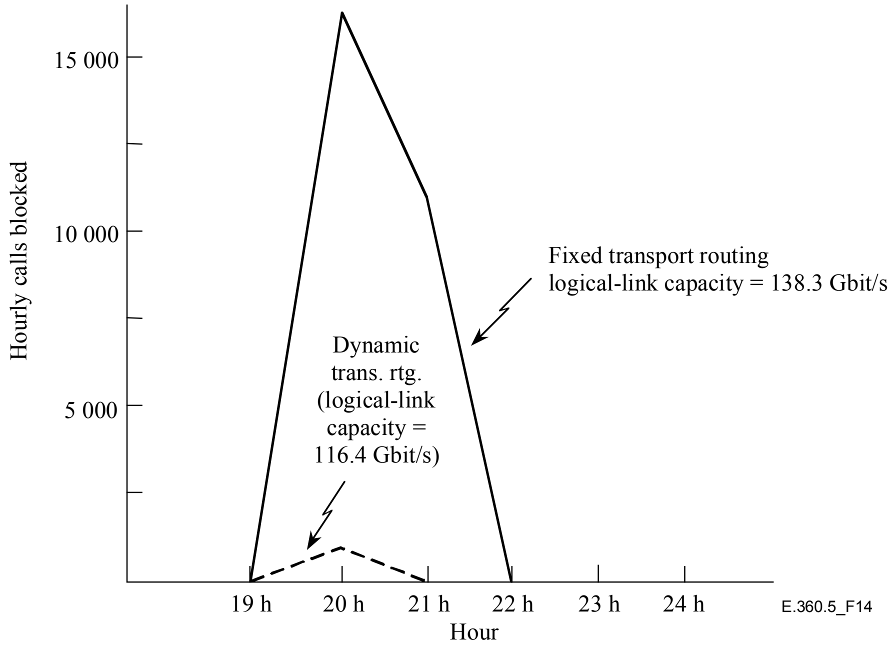

Figure A.2/E.360.5 – Dynamic transport routing performance for 25% overload on Jackson, Mississippi, node

Because the dynamic transport network is load balanced between access demands and inter-BR demands, this provides additional network robustness and resilience to unexpected traffic overloads, even though the dynamic transport routing network in this model has more than 15 percent less capacity than the fixed transport routing network. In this example, blocking-triggered rearrangement is allowed in the dynamic transport network. That is, as soon as node-pair blocking is detected, additional logical-link capacity is added to the affected links by cross-connecting spare node-termination capacity and spare logical-link capacity, which has been freed up as a result of the more efficient dynamic transport network design. As can be seen from Figure A.2, this greatly improves the network response to the overload.

A.5 Performance for peak-day traffic loads

A dynamic transport network design is performed for the Christmas traffic loads, and simulations performed for the base network and rearrangeable transport network design for the Christmas Day traffic. Results for the inter-BR blocking are summarized in Table A.7. Clearly, the rearrangeable transport network design eliminates the inter-BR network blocking, although the access node to BR blocking may still exist but is not quantified in the model. In addition to increased revenue, customer perception of network quality is also improved for these peak-day situations.

Table A.7/E.360.5 – Network performance for christmas day traffic overload

Hour of day	Fixed transport network traffic lost/delayed (%)	Dynamic transport network traffic lost/delayed (%)
9 to 10 am	17.2	0
10 to 11 am	22.2	0
11 to 12 am	29.7	0

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	**Overall network operation, telephone service, service operation and human factors**
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks and open system communications
Series Y	Global information infrastructure and Internet protocol aspects
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2003