INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

E.360.2

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

(05/2002)

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

International routing plan

QoS routing and related traffic engineering methods – Call routing and connection routing methods

ITU-T Recommendation E.360.2

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.360.2

QoS routing and related traffic engineering methods – Call routing and connection 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.360.2 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.

CONTENTS

Introduction

In this Recommendation we assume the separation of "call routing" and signalling for call establishment from "connection (or bearer-path) routing" and signalling for bearer-channel establishment. Call routing protocols primarily translate a number or a name, which is given to the network as part of a call setup, to a routing address needed for the connection (bearer-path) establishment. Call routing protocols are described for example in [Q.2761] for the Broadband ISDN Used Part (B-ISUP) call signalling, [ATM990048] for bearer-independent call control (BICC), or virtual trunking, call signalling, [H.323] for H.323 call signalling, [GR99] for the media gateway control [RFC2805] call signalling, and in [HSSR99] for the session initiation protocol (SIP) call signalling. Connection routing protocols include for example [Q.2761] for B-ISUP signalling, [ATM960055] for PNNI signalling, [ATM960061] for UNI signalling, [DN99] for switched virtual path (SVP) signalling, and [J00] for MPLS constraint-based routing label distribution protocol (CRLDP) signalling.

A specific connection or bearer-path routing method is characterized by the routing table used in the method. The routing table consists of a set of paths and rules to select one path from the route for a given connection request. When a connection request arrives at its originating node (ON), the ON implementing the routing method executes the path selection rules associated with the routing table for the connection to determine a selected path from among the path candidates in the route for the connection request. In a particular routing method, the path selected for the connection request is governed by the connection routing, or path selection, rules. Various path selection methods are discussed: fixed routing (FR) path selection, time-dependent routing (TDR) path selection, state-dependent routing (SDR) path selection, and event-dependent routing (EDR) path selection.

ITU-T Recommendation E.360.2

QoS routing and related traffic engineering methods – Call routing and connection 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, 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 Rec. ITU-T E.360.1.

3 Definitions

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

4 Abbreviations

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

5 Call routing methods

Call routing entails number (or name) translation to a routing address, which is then used for connection establishment. Routing addresses can consist, for example, of:

a) E.164 ATM end system addresses (AESAs) [E.191];

b) network routing addresses (NRAs) [E.353]; and/or

c) IP addresses [S94].

As discussed in ITU-T Rec. E.360.4, a TE requirement is the need for carrying E.164-AESA addresses, NRAs, and IP addresses in the connection-setup information element (IE). In that case, E.164-AESA addresses, NRAs, and IP addresses become the standard addressing method for interworking across IP-, ATM-, and TDM-based networks. Another TE requirement is that a call identification code (CIC) be carried in the call-control and bearer-control connection-setup IEs in order to correlate the call-control setup with the bearer-control setup ([Q.1901], [ATM990048]). Carrying these additional parameters in the Signalling System 7 (SS7) ISDN User Part (ISUP) connection-setup IEs is referred to as the bearer independent call control (BICC) protocol.

Number (or name) translation, then, should result in the E.164-AESA addresses, NRAs, and/or IP addresses. NRA formats are covered in [E.353], and IP-address formats in [S94]. The AESA address has a 20-byte format as shown in Figure 1a. [E.191].

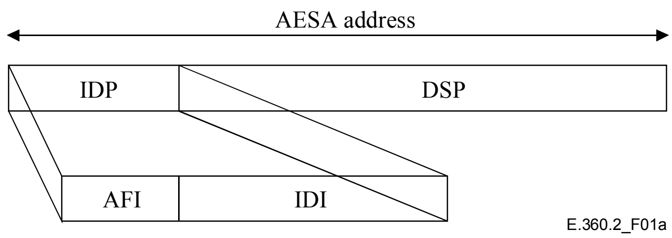

Figure 1a/E.360.2 – AESA address structure

The IDP is the initial domain part and the DSP is the domain specific part. The IDP is further subdivided into the AFI and IDI. The IDI is the initial domain identifier and can contain the 15-digit E.164 address if the AFI is set to 45. AFI is the authority and format identifier and determines what kind of addressing method is followed and, based on the 1 octet AFI value, the length of the IDI and DSP fields can change. The E.164-AESA address is used to determine the path to the destination endpoint. E.164-AESA addressing for B-ISDN services is supported in ATM networks using PNNI, through use of the above AESA format. In this case the E.164 part of the AESA address occupies the 8 octet IDI, and the 11 octet DSP can be used at the discretion of the network operator (perhaps for sub-addresses). The above AESA structure also supports AESA DCC (data country code) and AESA ICD (international code designator) addressing formats.

Within the IP network, routing is performed using IP addresses. Translation databases, such as based on domain name system (DNS) technology [F00], are used to translate the E.164 numbers/names for calls to IP addresses for routing over the IP network. The IP address is a 4-byte address structure as shown in Figure 1b.

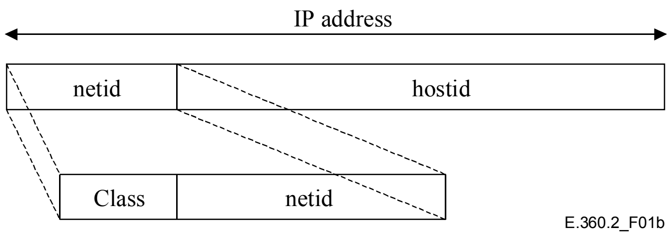

Figure 1b/E.360.2 – IP address structure

There are five classes of IP addresses. Different classes have different field lengths for the network identification field. Classless inter-domain routing (CIDR) allows blocks of addresses to be given to service providers in such a manner as to provide efficient address aggregation. This is accompanied by capabilities in the BGP4.0 protocol for efficient address advertisements [RL00], [S94].

6 Connection (bearer-path) routing methods

Connection routing is characterized by the routing table used in the method and rules to select one path from the route for a given connection or bandwidth-allocation request. When a connection/bandwidth-allocation request is initiated by an ON, the ON implementing the routing method executes the path selection rules associated with the routing table for the connection/bandwidth-allocation to find an admissible path from among the paths in the route that satisfies the connection/bandwidth-allocation request. In a particular routing method, the selected path is determined according to the rules associated with the routing table. In a network with originating connection/bandwidth-allocation control, the ON maintains control of the connection/bandwidth-allocation request. If crankback/bandwidth-not-available is used, for example, at a via node (VN), the preceding node maintains control of the connection/bandwidth-allocation request even if the request is blocked on all the links outgoing from the VN.

Here we are discussing network-layer connection routing (sometimes referred to as "layer-3" routing), as opposed to the link-layer logical-link ("layer-2") routing or physical-layer ("layer-1") routing. In the Recommendation the term "link" will normally mean "logical-link". In ITU-T Rec. E.360.5 we address logical-link routing.

The network-layer (layer-3) connection routing methods addressed include those discussed in:

- Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and Multiprotocol Label Switching (MPLS) for IP-based routing methods;
- User-to-Network Interface (UNI), Private Network-to-Network Interface (PNNI), ATM Inter-Network Interface (AINI), and Bandwidth Modify for ATM-based routing methods; and
- ITU-T Recs E.170, E.350, and E.351 for TDM-based routing methods.

In an IP network, logical links called traffic trunks can be defined which consist of MPLS label switched paths (LSPs) between the IP nodes. Traffic trunks are used to allocate the bandwidth of the logical links to various node pairs. In an ATM network, logical links called virtual paths (VPs) (the equivalent of traffic trunks) can be defined between the ATM nodes, and VPs can be used to allocate the bandwidth of the logical links to various node pairs. In a TDM network, the logical links consist of trunk groups between the TDM nodes.

A sparse logical link network is typically used with IP and ATM technology, as illustrated in Figure 2, and FR, TDR, SDR, and EDR can be used in combination with multilink shortest path selection.

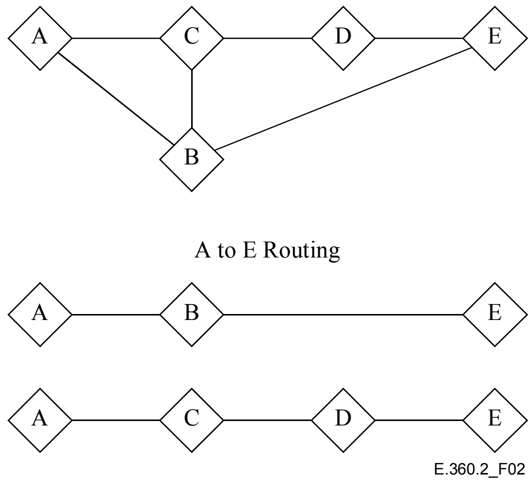

Figure 2/E.360.2 – Sparse logical network topology with connections routed on multilink paths

A meshed logical-link network is typically used with TDM technology, but can be used also with IP or ATM technology as well, and selected paths are normally limited to 1 or 2 logical links, or trunk groups, as illustrated in Figure 3.

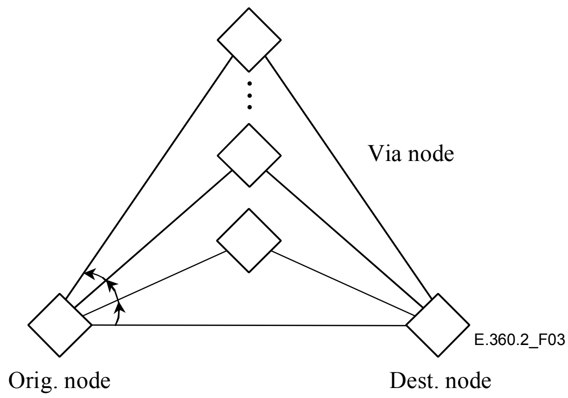

Figure 3/E.360.2 – Mesh logical network topology with connections routed on 1- and 2-link paths

Paths may be set up on individual connections (or "per-flow") for each call request, such as on a switched virtual circuits (SVC). Paths may also be set up for bandwidth-allocation requests associated with "bandwidth pipes" or traffic trunks, such as on switched virtual paths (SVPs) in ATM-based networks or constraint-based routing label switched paths (CRLSPs) in IP-based networks. Paths are determined by (normally proprietary) algorithms based on the network topology and reachable address information. These paths can cross multiple peer groups in ATM-based networks, and multiple autonomous systems (ASs) in IP-based networks. An ON may select a path from the routing table based on the routing rules and the QoS resource management criteria, described in ITU-T Rec. E.360.3, which must be satisfied on each logical-link in the path. If a link is not allowed based on the QoS criteria, then a release with crankback/bandwidth-not-available parameter is used to signal that condition to the ON in order to return the connection/bandwidth-allocation request to the ON, which may then select an alternate path. In addition to controlling bandwidth allocation, the QoS resource management procedures can check end-to-end transfer delay, delay variation, and transmission quality considerations such as loss, echo, and noise.

When source routing is used, setup of a connection/bandwidth-allocation request is achieved by having the ON identify the entire selected path including all VNs and DN in the path in a designated-transit-list (DTL) or explicit-route (ER) parameter in the connection-setup IE. If the QoS or traffic parameters cannot be realized at any of the VNs in the connection setup request, then the VN generates a crankback (CBK)/bandwidth-not-available (BNA) parameter in the connection-release IE which allows a VN to return control of the connection request to the ON for further alternate routing. In ITU-T Rec. E.360.4, the DTL/ER and CBK/BNA elements are identified as being required for interworking across IP-, ATM-, and TDM-based networks.

As noted earlier, connection routing, or path selection, methods are categorized into the following four types: fixed routing (FR), time-dependent routing (TDR), state-dependent routing (SDR), and event-dependent routing (EDR). We discuss each of these methods in the following paragraphs. Examples of each of these path selection methods are illustrated in Figures 4a and 4b and discussed in the following clauses.

Dynamic routing allows routing tables to be changed dynamically, either in an off-line, preplanned, time-varying manner, as in TDR, or on-line, in real time, as in SDR or EDR. With off-line, pre-planned TDR path selection methods, routing patterns contained in routing tables might change every hour or at least several times a day to respond to measured hourly shifts in traffic loads, and in general, TDR routing tables change with a time constant normally greater than a call/traffic-flow holding time. A typical TDR routing method may change routing tables every hour, which is longer than a typical voice call/traffic-flow holding time of a few minutes. Three implementations of TDR dynamic path selection are illustrated in Figure 4a, which shows multilink path routing, 2-link path routing, and progressive routing.

Dynamic multilink path routing	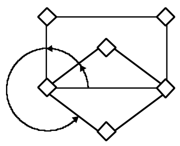	• One-link, two-link and multilink paths allowed • Paths hunted according to various rules: cyclic, cyclic block (CGH), skip-one-path, sequential • Path order changed with time; bandwidth broker predetermines routes based on traffic • Crankback from via node to originating node if blocked
Dynamic two-link path routing	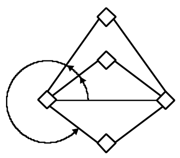	• One-link and two-link paths allowed • Paths hunted according to various rules: cyclic, cyclic block (CGH), skip-one-path, sequential • Path order changed with time; bandwidth broker predetermines routes based on traffic • Crankback from via node to originating node if blocked
Dynamic progressive routing	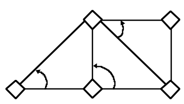	• One-link, two-link, and multilink paths allowed • Destination based link sequence hunted at each node • Link hunt order changed with time at originating node; bandwidth broker predetermines routes based on traffic • Progressive toward destination; no crankback or return to previous node

E.360.2_F04a

Figure 4a/E.360.2 – TDR dynamic path selection methods

TDR routing tables are preplanned, preconfigured, and recalculated perhaps each week within the capacity management network design function. Real-time dynamic path selection does not depend on precalculated routing tables. Rather, the node or centralized bandwidth broker senses the immediate traffic load and if necessary searches out new paths through the network possibly on a per-traffic-flow basis. With real-time path selection methods, routing tables change with a time constant on the order of or less than a call/traffic-flow holding time. As illustrated in Figure 4b, on-line, real-time path selection methods include EDR and SDR.

Event dependent routing	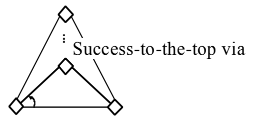	• Select primary path, then currently successful via path • Randomly select new via path when call is blocked • STT method: allow up to N crankbacks to find successful via path
State dependent routing (centralized, periodic)	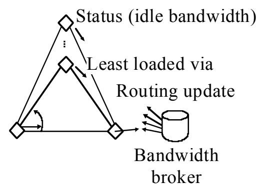	• Select primary path, then least loaded via path recommended by bandwidth broker • Each node sends periodic link status to bandwidth broker • Each node receives periodic routing update from bandwidth broker
State dependent routing (distributed, periodic)	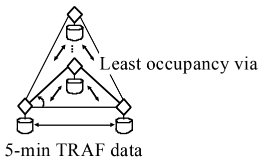	• Periodic (e.g. 5-minute) traffic data exchange • Select primary path, then least occupancy via path
State dependent routing (distributed, call-by-call)	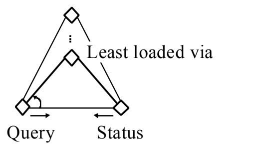	• Select primary path, then least loaded via path based on real-time status • Originating node queries terminating node for link status • Routing changes call-by-call

E.360.2_F04b

Figure 4b/E.360.2 – EDR and SDR dynamic path selection methods

7 Hierarchical Fixed Routing (FR) path selection

Hierarchical fixed routing (FR) is an important routing topology employed in all types of networks, including IP-, ATM-, and TDM-based networks. In IP-based networks, there is often a hierarchical relationship among different "areas", or sub-networks. Hierarchical multi-domain (or multi-area or multi-autonomous-system) topologies are normally used with IP routing protocols (OSPF, BGP) and ATM routing protocols (PNNI), as well as within almost all TDM-based network routing topologies.

For example, in Figure 4c, BB1 and BB2 could be backbone nodes in a "backbone area", and AN1 and AN2 could be access nodes in separate "access areas" distinct from the backbone area. Routing between the areas follows a hierarchical routing pattern, while routing within an area follows an interior gateway protocol (IGP), such as OSPF plus MPLS. Similarly, in ATM-based networks the same concept exists, but here the "areas" are called "peer-groups", and for example, the IGP used within peer-groups could be PNNI. In TDM-based networks, the routing between sub-networks, for example, metropolitan-area-networks and long-distance networks, is normally hierarchical, as in IP- and ATM-based networks, and the IGP in TDM-based networks could be either hierarchical or dynamic routing. We now discuss more specific attributes and methods for hierarchical FR path selection.

In a FR method, a routing pattern is fixed for a connection request. A typical example of fixed routing is a conventional, TDM-based, hierarchical alternate routing pattern where the route and route selection sequence are determined on a preplanned basis and maintained over a long period of time. Hierarchical FR is illustrated in Figure 4c. FR is more efficiently applied, however, when the network is nonhierarchical, or flat, as compared to the hierarchical structure [A98].

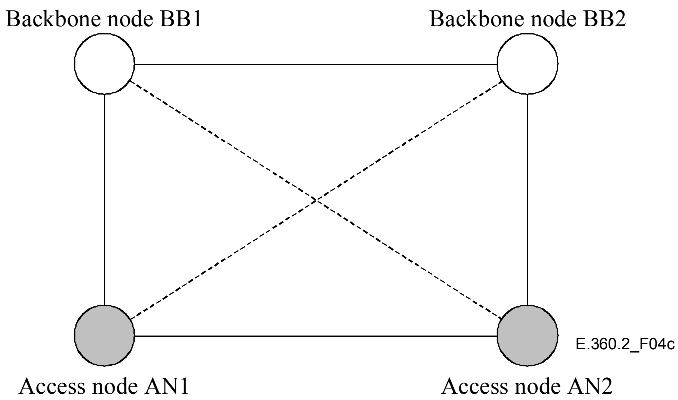

Figure 4c/E.360.2 – Hierarchical fixed routing path selection methods (2-level hierarchical network)

The aim of hierarchical fixed routing is to carry as much traffic as is economically feasible over direct links between pairs of nodes low in the hierarchy. This is accomplished by application of routing procedures to determine where sufficient load exists to justify high-usage logical-links, and then by application of alternate-routing principles that effectively pool the capacities of high-usage links with those of final links, to the end that all traffic is carried efficiently.

The routing of connection requests in a hierarchical network involves an originating ladder, a terminating ladder, and links interconnecting the two ladders. In a two-level network, for example, the originating ladder is the final link from lower level-1 node to the upper level-2 node, and the terminating ladder is the final link from upper level-2 node to the lower level-1 node. Links AN1-BB2, AN2-BB1, and BB1-BB2 in Figure 4c are examples of interladder links.

The identification of the proper interladder link for the routing of a given connection request identifies the originating ladder "exit" point and the terminating ladder "entry" point. Once these exit and entry points are identified and the intraladder links are known, a first-choice path from originating to terminating location can be determined.

Various levels of traffic concentration are used to achieve an appropriate balance between transport and switching. The generally preferred routing sequence for the AN1 to AN2 connections is:

1) A connection request involving no via nodes: path AN1-AN2 (if the link existed).
2) A connection request involving one via node: path AN1-BB2-AN2, AN1-BB1-AN2, in that order.
3) A connection request involving two via nodes: path AN1-BB1-BB2-AN2.

This procedure provides only the first-choice interladder link from AN1 to AN2. Connection requests from AN2 to AN1 often route differently. To determine the AN2-to-AN1 route requires reversing the diagram, making AN2-BB2 the originating ladder and AN1-BB1 the terminating ladder. In Figure 4c the preferred path from AN2 to AN1 is AN2-AN1, AN2-BB1-AN1, AN2-BB2-AN1, and AN2-BB2-BB1-AN1, in that order. The alternate path for any high-usage link is the path the node-to-node traffic load between the nodes would follow if the high-usage link did not exist. In Figure 4c, this is AN2-BB1-AN1.

8 Time-Dependent Routing (TDR) path selection

TDR methods are a type of dynamic routing in which the routing tables are altered at a fixed point in time during the day or week. TDR routing tables are determined on an off-line, preplanned basis and are implemented consistently over a time period. The TDR routing tables are determined considering the time variation of traffic load in the network, for example based on measured hourly load patterns. Several TDR time periods are used to divide up the hours on an average business day and weekend into contiguous routing intervals sometimes called load set periods. Typically, the TDR routing tables used in the network are coordinated by taking advantage of noncoincidence of busy hours among the traffic loads.

In TDR, the routing tables are preplanned and designed off-line using a centralized bandwidth broker, which employs a TDR network design model. Such models are discussed in ITU-T Rec. E.360.6. The off-line computation determines the optimal routes from a very large number of possible alternatives, in order to maximize network throughput and/or minimize the network cost. The designed routing tables are loaded and stored in the various nodes in the TDR network, and periodically recomputed and updated (e.g. every week) by the bandwidth broker. In this way, an ON does not require additional network information to construct TDR routing tables, once the routing tables have been loaded. This is in contrast to the design of routing tables on-line in real time, such as in the SDR and EDR methods described below. Paths in the TDR routing table may consist of time varying routing choices and use a subset of the available paths. Paths used in various time periods need not be the same.

Paths in the TDR routing table may consist of the direct link, a 2-link path through a single VN, or a multiple-link path through multiple VNs. Path routing implies selection of an entire path between originating and terminating nodes before a connection is actually attempted on that path. If a connection on one link in a path is blocked (e.g. because of insufficient bandwidth), the connection request then attempts another complete path. Implementation of such a routing method can be done through control from the originating node, plus a multiple-link crankback capability to allow paths of two, three, or more links to be used. Crankback is an information-exchange message capability that allows a connection request blocked on a link in a path to return to the originating node for further alternate routing on other paths. Path-to-path routing is nonhierarchical and allows the choice of the most economical paths rather than being restricted to hierarchical paths.

Path selection rules employed in TDR routing tables, for example, may consist of simple sequential routing. In the sequential method, all traffic in a given time period is offered to a single route, and lets the first path in the route overflow to the second path which overflows to the third path, and so on. Thus, traffic is routed sequentially from path to path, and the route is allowed to change from hour to hour to achieve the preplanned dynamic, or time varying, nature of the TDR method.

Other TDR path selection rules can employ probabilistic techniques to select each path in the route and thus influence the realized flows. One such method of implementing TDR multilink path selection is to allocate fractions of the traffic to routes and to allow the fractions to vary as a function of time. One approach is cyclic path selection, illustrated in Figure 4a, which has as its first route (1, 2, ..., M), where the notation (i, j, k) means that all traffic is offered first to path i, which overflows to path j, which overflows to path k. The second route of a cyclic route choice is a cyclic permutation of the first route: (2, 3, ..., M, 1). The third route is likewise (3, 4, ..., M, 1, 2), and so on. This approach has computational advantages because its cyclic structure requires considerably fewer calculations in the design model than does a general collection of paths. The route congestion level of cyclic routes are identical; what varies from route to route is the proportion of flow on the various links.

Two-link TDR path selection is illustrated in Figure 4a. An example implementation is 2-link sequential TDR (2S-TDR) path selection. By using the crankback signal, 2S-TDR limits path connections to at most two links, and, in meshed network topologies, such TDR 2-link sequential path selection allows nearly as much network utilization and performance improvement as TDR multilink path selection. This is because in the design of multilink path routing in meshed networks, about 98 percent of the traffic is routed on one- and 2-link paths, even though paths of greater length are allowed. Because of switching costs, paths with one or two links are usually less expensive than paths with more links. Therefore, as illustrated in Figure 4a, 2-link path routing uses the simplifying restriction that paths can have only one or two links, which requires only single-link crankback to implement and uses no common links as is possible with multilink path routing. Alternative 2-link path selection methods include the cyclic routing method described above, and sequential routing.

In sequential routing, all traffic in a given hour is offered to a single route, and the first path is allowed to overflow to the second path, which overflows to the third path, and so on. Thus, traffic is routed sequentially from path to path with no probabilistic methods being used to influence the realized flows. The reason that sequential routing works well is that permuting path order provides sufficient flexibility to achieve desired flows without the need for probabilistic routing. In 2S-TDR, the sequential route is allowed to change from hour to hour. The TDR nature of the dynamic path selection method is achieved by introducing several route choices which consist of different sequences of paths, and each path has one or, at most, two links in tandem.

Paths in the routing table are subject to depth-of-search (DoS) restrictions for QoS resource management, which is discussed in ITU-T Rec. E.360.3. DoS requires that the bandwidth capacity available on each link in the path be sufficient to meet a DoS bandwidth threshold level, which is passed to each node in the path in the setup message. DoS restrictions prevent connections that path on the first-choice or primary (often the shortest) ON-DN path, for example, from being swamped by alternate routed multiple-link connections.

A TDR connection set-up example is now given. The first step is for the ON to identify the DN and routing table information to the DN. The ON then tests for spare capacity on the first or shortest path, and, in doing this, supplies the VNs and DN on this path, along with the DoS parameter, to all nodes in the path. Each VN tests the available bandwidth capacity on each link in the path against the DoS threshold. If there is sufficient capacity, the VN forwards the connection setup to the next node, which performs a similar function. If there is insufficient capacity, the VN sends a release message with crankback/bandwidth-not-available parameter back to the ON, at which point the ON tries the next path in the route as determined by the routing table rules. As described above, the TDR routes are preplanned off-line, and then loaded and stored in each ON.

Allocating traffic to the optimum path choice during each time period leads to design benefits due to the noncoincidence of loads. Since in many network applications traffic demands change with time in a reasonably predictable manner, the routing also changes with time to achieve maximum link utilization and minimum network cost. Several TDR routing time periods are used to divide up the hours on an average business day and weekend into contiguous routing intervals. The network design is performed in an off-line, centralized computation in the bandwidth broker that determines the optimal routing tables from a very large number of possible alternatives in order to minimize the network cost. In TDR path selection, rather than determine the optimal routing tables based on real-time information, a centralized bandwidth broker design system employs a design model, such as described in ITU-T Rec. E.360.6. The effectiveness of the design depends on how accurately we can estimate the traffic load on the network. Forecast errors are corrected in the short-term capacity management process, which allows routing table updates to replace link augments whenever possible, as described in ITU-T Rec. E.360.7.

9 State-Dependent Routing (SDR) path selection

In SDR, the routing tables are altered automatically according to the state of the network. For a given SDR method, the routing table rules are implemented to determine the path choices in response to changing network status, and are used over a relatively short time period. Information on network status may be collected at a central bandwidth broker processor or distributed to nodes in the network. The information exchange may be performed on a periodic or on-demand basis. SDR methods use the principle of routing connections on the best available path on the basis of network state information. For example, in the least loaded routing (LLR) method, the residual capacity of candidate paths is calculated, and the path having the largest residual capacity is selected for the connection. Various relative levels of link occupancy can be used to define link load states, such as lightly-loaded, heavily-loaded, or bandwidth-not-available states. Methods of defining these link load states are discussed in ITU-T Rec. E.360.3. In general, SDR methods calculate a path cost for each connection request based on various factors such as the load-state or congestion state of the links in the network.

In SDR, the routing tables are designed on-line by the ON or a central bandwidth broker processor (BBP) through the use of network status and topology information obtained through information exchange with other nodes and/or a centralized BBP. There are various implementations of SDR distinguished by:

a) whether the computation of the routing tables is distributed among the network nodes or centralized and done in a centralized BBP; and

b) whether the computation of the routing tables is done periodically or connection by connection.

This leads to three different implementations of SDR:

a) centralized periodic SDR (CP-SDR) – here the centralized BBP obtains link status and traffic status information from the various nodes on a periodic basis (e.g. every 10 seconds) and performs a computation of the optimal routing table on a periodic basis. To determine the optimal routing table, the BBP executes a particular routing table optimization procedure such as LLR and transmits the routing tables to the network nodes on a periodic basis (e.g. every 10 seconds);

b) distributed periodic SDR (DP-SDR) – here each node in the SDR network obtains link status and traffic status information from all the other nodes on a periodic basis (e.g. every 5 minutes) and performs a computation of the optimal routing table on a periodic basis (e.g. every 5 minutes). To determine the optimal routing table, the ON executes a particular routing table optimization procedure such as LLR;

c) distributed connection-by-connection SDR (DC-SDR) – here an ON in the SDR network obtains link status and traffic status information from the DN, and perhaps from selected VNs, on a connection by connection basis and performs a computation of the optimal routing table for each connection. To determine the optimal routing table, the ON executes a particular routing table optimization procedure such as LLR.

In DP-SDR path selection, nodes may exchange status and traffic data, for example, every five minutes, between traffic management processors, and based on analysis of this data, the traffic management processors can dynamically select alternate paths to optimize network performance. This method is illustrated in Figure 4b. Flooding is a common technique for distributing the status and traffic data, however other techniques with less overhead are also available, such as a query-for-status method, as discussed in ITU-T Rec. E.360.4.

Figure 4b illustrates a CP-SDR path selection method with periodic updates based on periodic network status. CP-SDR path selection provides near-real-time routing decisions by having an update of the idle bandwidth in each link sent to a network database every five seconds. Routing tables are determined from analysis of the status data using a path selection method which provides that the shortest path choice is used if the bandwidth is available. If the shortest path is busy (e.g. bandwidth is unavailable on one or more links), the second path is selected from the list of feasible paths on the basis of having the greatest level of idle bandwidth at the time; the current second path choice becomes the third, and so on. This path update is performed, for example, every five seconds. The CP-SDR model uses dynamically activated bandwidth reservation and other controls to automatically modify routing tables during network overloads and failures. CP-SDR requires the use of network status and routing recommendation information-exchange messages.

Figure 4b also illustrates an example of a DC-SDR path selection method. In DC-SDR, the routing computations are distributed among all the nodes in the network. DC-SDR uses real-time exchange of network status information, such as with query and status messages, to determine an optimal path from a very large number of possible choices. With DC-SDR, the originating node first tries the primary path and, if it is not available, finds an optimal alternate path by querying the destination node and perhaps several via nodes through query-for-status network signalling for the busy-idle load status of all links connected on the alternate paths to the destination node. The originating node then finds the least loaded alternate path to route the connection request. DC-SDR computes required bandwidth allocations by virtual network from node-measured traffic flows and uses this capacity allocation to reserve capacity when needed for each virtual network. Any excess traffic above the expected flow is routed to temporarily idle capacity borrowed from capacity reserved for other loads that happen to be below their expected levels. Idle link capacity is communicated to other nodes via the query-status information-exchange messages, as illustrated in Figure 4b, and the excess traffic is dynamically allocated to the set of allowed paths that are identified as having temporarily idle capacity. DC-SDR controls the sharing of available capacity by using dynamic bandwidth reservation, as described in ITU-T Rec. E.360.3, to protect the capacity required to meet expected loads and to minimize the loss of traffic for classes-of-service which exceed their expected load and allocated capacity.

Paths in the SDR routing table may consist of the direct link, a 2-link path through a single VN, or a multiple-link path through multiple VNs. Paths in the routing table are subject to DoS restrictions on each link.

10 Event-Dependent Routing (EDR) path selection

In EDR, the routing tables are updated locally on the basis of whether connections succeed or fail on a given path choice. In the EDR learning approaches, the path last tried, which is also successful, is tried again until blocked, at which time another path is selected at random and tried on the next connection request. EDR path choices can also be changed with time in accordance with changes in traffic load patterns. Success-to-the-top (STT) EDR path selection, illustrated in Figure 4b, is a decentralized, on-line path selection method with update based on random routing. STT-EDR uses a simplified decentralized learning method to achieve flexible adaptive routing. The primary path path-p is used first if available, and a currently successful alternate path path-s is used until it is blocked. In the case that path-s is blocked (e.g. bandwidth is not available on one or more links), a new alternate path path-n is selected at random as the alternate path choice for the next connection request overflow from the primary path. As described in ITU-T Rec. E.360.3, dynamically activated bandwidth reservation is used under congestion conditions to protect traffic on the primary path. STT-EDR uses crankback when an alternate path is blocked at a via node, and the connection request advances to a new random path choice. In STT-EDR, many path choices can be tried by a given connection request before the request is blocked.

In the EDR learning approaches, the current alternate path choice can be updated randomly, cyclically, or by some other means, and may be maintained as long as a connection can be established successfully on the path. Hence the routing table is constructed with the information determined during connection setup, and no additional information is required by the ON. Paths in the EDR routing table may consist of the direct link, a 2-link path through a single VN, or a multiple-link path through multiple VNs. Paths in the routing table are subject to DoS restrictions on each link. Note that for either SDR or EDR, as in TDR, the alternate path for a connection request may be changed in a time-dependent manner considering the time-variation of the traffic load.

11 Interdomain routing

In current practice, interdomain routing protocols generally do not incorporate standardized path selection or per class-of-service resource management. For example, in IP-based networks BGP [RL00] is used for interdomain routing but does not incorporate per class-of-service resource allocation as described in this clause. Also, MPLS techniques have not yet been addressed for interdomain applications. Extensions to interdomain routing methods discussed in this clause therefore can be considered to extend the call routing and connection routing concepts to routing between network domains.

Many of the principles described for intradomain routing can be extended to interdomain routing. As illustrated in Figure 5, interdomain routing paths can be divided into three types:

- a primary shortest path between the originating domain and destination domain;
- alternate paths with all nodes in the origination domain and destination domain; and
- alternate or transit paths through other transit domains.

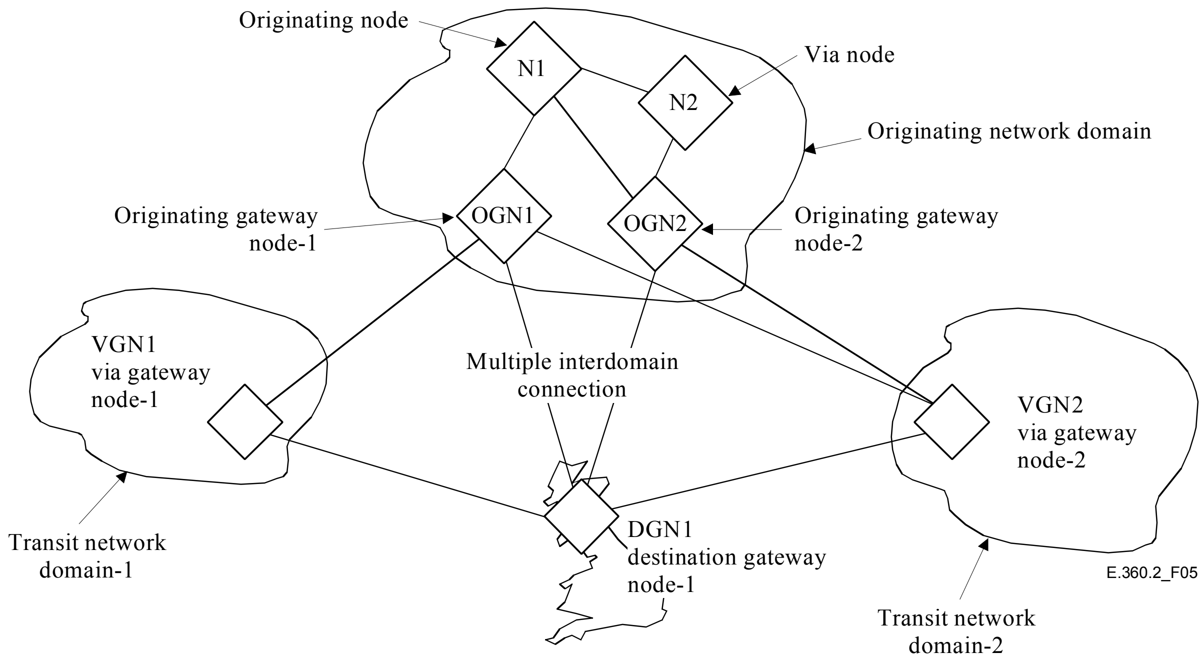

Figure 5/E.360.2 – Multiple ingress/egress interdomain routing

Interdomain routing can support a multiple ingress/egress capability, as illustrated in Figure 5 in which a connection request is routed either on the shortest path or, if not available, via an alternate path through any one of the other nodes from an originating node to a gateway node.

Within an originating network, a destination network could be served by more than one gateway node, such as OGN1 and OGN2 in Figure 5, in which case multiple ingress/egress routing is used. As illustrated in Figure 5, with multiple ingress/egress routing, a connection request from the originating node N1 destined for the destination gateway node DGN1 tries first to access the links from originating gateway node OGN2 to DGN1. In doing this, it is possible that the connection request could be routed from N1 to OGN2 directly or via N2. If no bandwidth is available from OGN2 to DGN1, the control of the connection request can be returned to N1 with a crankback/bandwidth-not-available indicator, after which the connection request is routed to OGN1 to access the OGN1-to-DGN1 bandwidth. If the connection request cannot be completed on the link connecting gateway node OGN1 to DGN1, the connection request can return to the originating node N1 through use of a crankback/band-not-available indicator for possible further routing to another gateway node (not shown). In this manner, all ingress/egress connectivity is utilized to a connecting network, maximizing connection request completion and reliability.

Once the connection request reaches an originating gateway node (such as OGN1 or OGN2), this node determines the routing to the destination gateway node DGN1 and routes the connection request accordingly. In completing the connection request to DGN1, an originating gateway node can dynamically select a direct shortest path, an alternate path through an alternate node in the destination network, or perhaps an alternate path through an alternate node in another network domain. Hence, with interdomain routing, connection requests are routed first to a shortest primary path between the originating and destination domain, then to a list of alternate paths through alternate nodes in the terminating network domain, then to a list of alternate paths through alternate nodes in the originating network domain (e.g. OGN1 and OGN2 in Figure 5), and finally to a list of alternate paths through nodes in other transit network domains.

Examples of alternate paths which might be selected through a transit network domain are N1-OGN1-VGN1-DGN1, N1-OGN1-VGN2-DGN1, or N1-N2-OGN2-VGN2-DGN1 in Figure 5. Such paths through transit network domains may be tried last in the example network configuration in the Figure. For example, flexible interdomain routing may try to find an available alternate path based on link load states, where known, and connection request completion performance, where it can be inferred. That is, the originating gateway node (e.g. node OGN1 in Figure 5) may use its link status to a via node in a transit domain (e.g. links OGN1-VGN1 and OGN1-VGN2) in combination with the connection request completion performance from the candidate via node to the destination node in the destination network domain, in order to find the most available path to route the connection request over. For each path, a load state and a completion state are tracked. The load state indicates whether the link bandwidth from the gateway node to the via node is lightly loaded, heavily loaded, reserved, or busy. The completion state indicates whether a path is achieving above-average completion, average completion, or below-average completion. The selection of a via path, then, is based on the load state and completion state. Alternate paths in the same destination network domain and in a transit network domain are each considered separately. During times of congestion, the link bandwidth to a candidate via node may be in a reserved state, in which case the remaining link bandwidth is reserved for traffic routing directly to the candidate via node. During periods of no congestion, capacity not needed by one virtual network is made available to other virtual networks that are experiencing loads above their allocation.

Similar to intradomain routing, interdomain routing can use discrete load states for interdomain links terminating in the originating domain (e.g. links OGN1-VGN1, OGN1-DGN1, OGN2-DGN1). As described in ITU-T Rec. E.360.3, these link load states could may include lightly-loaded, heavily-loaded, reserved, and busy/bandwidth-not-available, in which the idle link bandwidth is compared with the load state thresholds for the link to determine its load condition. Completion rate is tracked on the various via paths (such as the path through via node VGN1 or VGN2 to destination node DGN1 in Figure 5) by taking account of the information relating either the successful completion or non-completion of a connection request through the via node. A non-completion, or failure, is scored for the connection request if a signalling release message is received from the far end after the connection request seizes an egress link, indicating a network in-completion cause value. If no such signalling release message is received after the connection request seizes capacity on the egress link, then the connection request is scored as a success. Each gateway node keeps a connection request completion history of the success or failure, for example, of the last 10 connection requests using a particular via path, and it drops the oldest record and adds the connection request completion for the newest connection request on that path. Based on the number of connection request completions relative to the total number of connection requests, a completion state is computed.

Based on the completion states, connection requests are normally routed on the first path with a high completion state with a lightly loaded egress link. If such a path does not exist, then a path having an average completion state with a lightly loaded egress link is selected, followed by a path having a low completion state with a lightly loaded egress link. If no path with a lightly loaded egress link is available, and if the search depth permits the use of a heavily loaded egress link, the paths with heavily loaded egress links are searched in the order of high completion, average completion, and low completion. If no such paths are available, paths with reserved egress links are searched in the same order, based on the connection request completion state, if the search depth permits the use of a reserved egress link.

The rules for selecting primary shortest paths and alternate paths for a connection request are governed by the availability of shortest path bandwidth and node-to-node congestion. The path sequence consists of the primary shortest path, lightly loaded alternate paths, heavily loaded alternate paths, and reserved alternate paths. Alternate paths are first selected which include nodes only in the originating and destination domains, and then selected through transit domains if necessary.

Thus we have illustrated that interdomain routing methods can be considered to extend the intradomain call routing and connection routing concepts, such as flexible path selection and per-class-of-service bandwidth selection, to routing between network domains.

12 Conclusions/recommendations

We have discussed call routing and connection routing methods employed in TE functions. Several connection routing alternatives were discussed, which include FR, TDR, EDR, and SDR methods. Models were presented to illustrate the network design and performance tradeoffs between the many TE approaches explained in the Recommendation, and conclusions were drawn on the advantages of various routing and topology options in network operation. Overall the packet-based (e.g. MPLS/TE) multilink, sparse-topology routing strategies were found to offer several advantages.

The following conclusions/recommendations are reached in the Recommendation:

- TE methods are recommended to be applied, and in all cases of the TE methods being applied, network performance is always better and usually substantially better than when no TE methods are applied.
- Sparse-topology multilink-routing networks are recommended and provide better overall performance under overload than meshed-topology networks, but performance under failure may favor the 2-link STT-EDR/DC-SDR meshed-topology options with more alternate routing choices.
- Single-area flat topologies are recommended and exhibit better network performance and, as discussed and modeled in ITU-T Rec. E.360.6, greater design efficiencies in comparison with multi-area hierarchical topologies. As illustrated in ITU-T Rec. E.360.4, larger administrative areas can be achieved through use of EDR-based TE methods as compared to SDR-based TE methods.
- Event-dependent-routing (EDR) TE path selection methods are recommended and exhibit comparable, or better, network performance compared to state-dependent-routing (SDR) methods.
  a) EDR TE methods are shown to an important class of TE algorithms. EDR TE methods are distinct from the TDR and SDR TE methods in how the paths (e.g. MPLS label switched paths, or LSPs) are selected. In the SDR TE case, the available link bandwidth (based on LSA flooding of ALB information) is typically used to compute the path. In the EDR TE case, the ALB information is not needed to compute the path, therefore the ALB flooding does not need to take place (reducing the overhead).

b) EDR TE algorithms are adaptive and distributed in nature and typically use learning models to find good paths for TE in a network. For example, in a success-to-the-top (STT) EDR TE method, if the LSR-A to LSR-B bandwidth needs to be modified, say increased by delta-BW, the primary LSP-p is tried first. If delta-BW is not available on one or more links of LSP-p, then the currently successful LSP-s is tried next. If delta-BW is not available on one or more links of LSP-s, then a new LSP is searched by trying additional candidate paths until a new successful LSP-n is found or the candidate paths are exhausted. LSP-n is then marked as the currently successful path for the next time bandwidth needs to be modified. The performance of distributed EDR TE methods is shown to be equal to, or better than, SDR methods, centralized or distributed.

c) While SDR TE models typically use available-link-bandwidth (ALB) flooding for TE path selection, EDR TE methods do not require ALB flooding. Rather, EDR TE methods typically search out capacity by learning models, as in the STT method above. ALB flooding can be very resource intensive, since it requires link bandwidth to carry LSAs, processor capacity to process LSAs, and the overhead can limit area/autonomous system (AS) size. Modelling results show EDR TE methods can lead to a large reduction in ALB flooding overhead without loss of network throughput performance (as shown in ITU-T Rec. E.360.4).

d) State information as used by the SDR options (such as with link-state flooding) provides essentially equivalent performance to the EDR options, which typically used distributed routing with crankback and no flooding.

e) Various path selection methods can interwork with each other in the same network, as required for multi-vendor network operation.

• Interdomain routing methods are recommended which extend the intradomain call routing and connection routing concepts, such as flexible path selection and per-class-of-service bandwidth selection, to routing between network domains.

Annex A

Modelling of traffic engineering methods

In the Recommendation, a full-scale national network node model is used, together with a multiservice traffic demand model, to study various TE scenarios and tradeoffs. The 135-node national model is illustrated in Figure A.1.

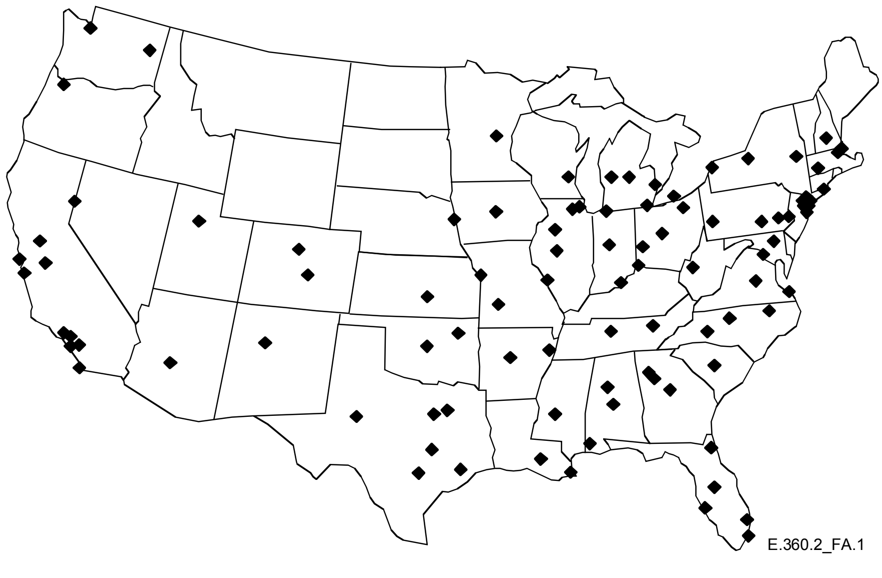

Figure A.1/E.360.2 – 135-node national network model

Typical voice/ISDN traffic loads are used to model the various network alternatives, which are based on 72 hours of a full-scale national network loading. Tables A.1a, A.1b, and A.1c summarize the multiservice traffic model used for the TE studies. Three levels of traffic priority – key, normal, and best-effort – are given to the various class-of-service categories, or virtual networks (VNETs), illustrated in Tables A.1a-A.1c. Class-of-service, traffic priority, and QoS resource management are all discussed further in ITU-T Rec. E.360.3.

The voice/ISDN loads are further segmented in the model into eight constant-bit-rate (CBR) VNETs, including business voice, consumer voice, international voice in and out, key-service voice, normal and key-service 64-kbit/s ISDN data, and 384-kbit/s ISDN data. For the CBR voice services, the mean data rate is assumed to be 64 kbit/s for all VNETs except the 384-kbit/s ISDN data VNET-8, for which the mean data rate is 384 kbit/s.

Table A.1a/E.360.2 – Virtual Network (VNET) traffic model used for TE studies

Virtual network index	Virtual network name	Service identity examples	Virtual network traffic priority and traffic characteristics
VNET-1 (CBR)	Business voice	Virtual private network (VPN), direct connect 800, 800 service, 900 service	Normal priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-2 (CBR)	Consumer voice	Long distance service (LDS)	Normal priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-3 (CBR)	INTL voice outbound	INTL LDS outbound, INTL 800 outbound, global VPN outbound, INTL transit	Normal priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-4 (CBR)	INTL voice inbound (key)	INTL LDS inbound, INTL 800 inbound, global VPN inbound, INTL transit inbound	Key priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-5 (CBR)	800-gold (key)	Direct connect 800 gold, 800 gold, VPN-key	Key priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-6 (CBR)	64 kbit/s ISDN	64 kbit/s switched digital service (SDS), 64 kbit/s switched digital INTL (SDI)	Normal priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-7 (CBR)	64 kbit/s ISDN (key)	64 kbit/s SDS and SDI (key)	Key priority; 64 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-8 (CBR)	384 kbit/s ISDN	384 kbit/s SDS, 384 kbit/s SDI	Normal priority; 384 kbit/s CBR; 72 hours traffic load data (Saturday, Sunday, Monday, 1998)
VNET-9 (VBR-RT)	IP telephony variable rate, equiv-BW allocation, interactive and delay sensitive	IP telephony, compressed voice	Normal priority; variable rate, equiv-BW allocation, interactive and delay sensitive; VBR-RT: 10% of VNET1+VNET2+VNET3+VNET4+VNET5 traffic load, call data rate varies from 6.4 kbit/s to 51.2 kbit/s (25.6 kbit/s mean)

Table A.1a/E.360.2 – Virtual Network (VNET) traffic model used for TE studies

Virtual network index	Virtual network name	Service identity examples	Virtual network traffic priority and traffic characteristics
VNET-10 (VBR-NRT)	IP multimedia variable rate, equiv-BW allocation, non-interactive and not delay sensitive	IP multimedia, WWW, credit card check	Normal priority; variable rate, equiv-BW allocation, non-interactive and not delay sensitive; VBR-NRT: 30% of VNET2 traffic load, call data rate varies from 38.4 kbit/s to 64 kbit/s (51.2 kbit/s mean)
VNET-11 (UBR)	UBR best effort variable rate, no BW allocation, non-interactive and not delay sensitive	Voice mail, email, file transfer	Best-effort priority; variable rate, no BW allocation, non-interactive and not delay sensitive; UBR: 30% of VNET1 traffic load, call data rate varies from 6.4 kbit/s to 3072 kbit/s (1536 kbit/s mean)

Table A.1b/E.360.2 – Virtual Network (VNET) traffic model used for TE studies

Average number of flows by network busy hours (CST)									
Virtual network index	**Virtual network name**	**Sunday 8:00 pm**	**Monday 10:00 am**	**Monday 11:00 am**	**Monday 2:00 pm**	**Monday 3:00 pm**	**Monday 4:00 pm**	**Monday 8:00 pm**	**Monday 9:00 pm**
VNET-1 (CBR)	Business voice	108 459.3	616 190.8	678 423.2	672 853.4	676 348.1	661 489.9	232 997.4	193 837.5
VNET-2 (CBR)	Consumer voice	457 580.8	247 198.4	269 968.7	258 178.2	263 387.9	280 522.8	465 911.6	484 810.9
VNET-3 (CBR)	INTL voice outbound	28 124.5	25 976.3	27 276.2	22 417.6	23 079.2	23 053.9	21 939.3	22 064.3
VNET-4 (CBR)	INTL voice inbound (key)	11 725.8	23 969.9	25 098.4	18 491.8	18 034.8	17 382.3	1 2112.0	12 239.6
VNET-5 (CBR)	800-gold (key)	1 506.5	6 672.7	7 489.9	7 457.3	7 611.5	7 408.6	3 211.4	2 741.6
VNET-6 (CBR)	64 kbit/s ISDN	908.1	3 306.7	3 587.7	3 922.3	3 515.7	3 161.6	1 677.5	1 390.6
VNET-7 (CBR)	64 kbit/s ISDN (key)	77.2	454.8	419.2	181.2	168.6	168.8	162.5	116.8
VNET-8 (CBR)	384 kbit/s ISDN	1.0	21.0	18.0	29.2	33.2	26.8	2.2	2.0
VNET-9(VBR-RT)	IP telephony	60 739.8	92 000.8	100 825.9	97 940.1	9 8846.3	98 986.0	7 3616.9	71 567.2
VNET-10 (VBR-NRT)	IP multimedia	137 274.5	74 159.5	80 990.6	77 453.5	7 9016.2	84 156.8	139 773.5	145 443.2
VNET-11(UBR)	UBR best effort	27 154.7	166 574.9	184 626.4	183 204.2	183 602.5	179 601.3	60 477.9	49 866.7
TOTAL		833 532.0	1 256 501.1	1 378 697.2	1 342 099.4	1 353 611.5	1 355 925.9	1 011 880.0	984 066.5

Table A.1c/E.360.2 – Virtual Network (VNET) Traffic Model used for TE Studies

Average Data Volume (Mbit/s) by Network Busy Hours (CST)									
Virtual network index	**Virtual network name**	**Sunday 8:00 pm**	**Monday 10:00 am**	**Monday 11:00 am**	**Monday 2:00 pm**	**Monday 3:00 pm**	**Monday 4:00 pm**	**Monday 8:00 pm**	**Monday 9:00 pm**
VNET-1 (CBR)	Business voice	6 941.3	39 436.3	43 419.0	43 062.7	43 286.3	42 335.4	14 911.7	12 405.5
VNET-2 (CBR)	Consumer voice	29 285.2	15 820.7	17 278.0	16 523.5	16 856.8	17 953.4	29 818.3	31 028.0
VNET-3 (CBR)	INTL voice outbound	1 800.0	1 662.5	1 745.7	1 434.7	1 477.1	1 475.4	1 404.1	1 412.1
VNET-4 (CBR)	INTL voice inbound (key)	750.4	1 534.1	1 606.3	1 434.7	1 477.1	11 112.5	775.2	783.3
VNET-5 (CBR)	800-gold (key)	96.4	427.1	479.4	477.3	487.1	474.2	205.5	175.5
VNET-6 (CBR)	64 kbit/s ISDN	58.1	211.6	229.6	251.0	225.0	202.3	107.4	89.0
VNET-7 (CBR)	64 kbit/s ISDN (key)	4.9	29.1	26.8	11.6	10.8	10.8	10.4	7.5
VNET-8 (CBR)	384 kbit/s ISDN	0.4	8.1	6.9	11.2	12.8	10.3	0.9	0.8
VNET-9 (VBR-RT)	IP telephony	1 554.9	2 355.2	2 581.1	2 507.3	2 530.5	2 534.0	1 884.6	1 832.2
VNET-10 (VBR-NRT)	IP multimedia	7 028.5	3 797.0	4 146.7	3 965.6	4 045.6	4 308.8	7 156.4	7 446.7
VNET-11 (UBR)	UBR best effort	41 709.7	255 858.8	283 585.9	281 401.8	282 012.7	275 867.7	92 894.2	76 595.1
TOTAL		89 226.2	321 138.0	355 103.2	350 827.0	352 092.8	346 280.8	149 165.4	131 774.0

The data services traffic model incorporates typical traffic load patterns and comprises three additional VNET load patterns. These data services VNETs include:

- variable bit rate real-time (VBR-RT) VNET-9, representing services such as IP-telephony and compressed voice;
- variable bit rate non-real-time (VBR-NRT) VNET-10, representing services such as WWW multimedia and credit card check; and
- unassigned bit rate (UBR) VNET-11, representing services such as email, voice mail, and file transfer multimedia applications.

For the VBR-RT connections, the data rate varies from 6.4 to 51.2 kbit/s with a mean of 25.6 kbit/s. The VBR-RT connections are assumed to be interactive and delay sensitive. For the VBR-NRT connections, the data rate varies from 38.4 to 64 kbit/s with a mean of 51.2 kbit/s, and the VBR-NRT flows are assumed to be non-delay sensitive. For the UBR connections, the data rate varies from 6.4 to 3072 kbit/s with a mean of 1536 kbit/s. The UBR flows are assumed to be best-effort priority and non-delay sensitive. For modelling purposes, the service and link bandwidth is segmented into 6.4 kbit/s slots, that is, 10 slots per 64 kbit/s channel.

Here the traffic loads are dynamically varying and tracked by the exponential smoothing models discussed in ITU-T Rec. E.360.3. Table A.1b gives the average number of flows for each class of service (VNET) in various network busy hours, and Table A.1c gives the average data volume in Mbit/s for each class of service (VNET) in various network busy hours. We can see that the voice/ISDN traffic (i.e. VNETs 1-8 in Tables A.1a-A.1c) has a majority of the flows (approximately 75%) of the total in Monday busy hours, compared to the various "data" traffic sources (i.e. VNETs 9-11 in Tables A.1a-A.1c). However, the voice/ISDN traffic has a minority of the total traffic data volume approximately 70-80%) of the total Mbit/s demand, compared to the various "data" traffic sources (i.e. VNETs 9-11 in Tables A.1a-A.1c). The model is based on traffic projections for the "data" traffic and actual voice/ISDN traffic levels, wherein the data traffic dominating the voice/ISDN traffic is a realistic scenario under many traffic projections.

The cost model represents typical switching and transport costs, and illustrates the economies-of-scale for costs projected for high capacity network elements in the future. Table A.2 gives the model used for average switching and transport costs allocated per 64 kbit/s unit of bandwidth, as follows:

Table A.2/E.360.2 – Cost assumptions (average cost per equivalent 64 kbit/s bandwidth)

Data rate	Average transport cost	Average switching/ cross-connect cost
DS3	0.19 × miles + 8.81	26.12
OC3	0.17 × miles + 9.76	19.28
OC12	0.15 × miles + 7.03	9.64
OC48	0.05 × miles + 2.77	3.92

A discrete event network design model, described in ITU-T Rec. E.360.6, is used in the design and analysis of 5 connection routing methods with TE methods applied: 2-STT-EDR path routing in a meshed logical network, 2-link DC-SDR routing in a meshed logical network, and multilink STT-EDR, DC-SDR, and DP-SDR routing, as might be supported for example by MPLS TE in a sparse logical network. We also model the case where no TE call and connection routing methods are applied.

The network models for the 2-link STT-EDR/DC-SDR, and multilink STT-EDR/DC-SDR/DP-SDR networks are now described. In the 2-link STT-EDR and DC-SDR models, we assume 135 packet-switched-nodes (MPLS- or PNNI-based). Synchronous to asynchronous conversion (SAC) is assumed to occur at the packet-switched-nodes for link connections from circuit-switched-nodes. Links in these 2-link STT-EDR/DC-SDR models are assumed to provide more fine-grained (1.536 Mbit/sT1-level) logical link bandwidth allocation, and a meshed network topology design results among the nodes, that is, links exist between most (90 percent or more) of the nodes. In the 2-link STT-EDR/DC-SDR models, one and 2-link routing with crankback is used throughout the network. Two-link path selection is modeled both with both STT path selection and distributed connection-by-connection SDR (DC-SDR) path selection. Packet-switched-nodes use 2-link STT-EDR or 2-link DC-SDR routing to all other nodes. Quality-of-Service priority queuing is modeled in the performance analyses, in which the key-services are given the highest priority, normal services the middle priority, and best-effort services the lowest priority in the queuing model. This queuing model quantifies the level of delayed traffic for each virtual network. In routing a connection with 2-link STT-EDR routing, the ON checks the equivalent bandwidth and allowed DoS first on the direct path, then on the current successful 2-link via path, and then sequentially on all candidate 2-link paths. In routing a connection with 2-link DC-SDR, the ON checks the equivalent bandwidth and allowed DoS first on the direct path, and then on the least-loaded path that meets the equivalent bandwidth and DoS requirements. Each VN checks the equivalent bandwidth and allowed DoS provided in the setup message, and uses crankback to the ON if the equivalent bandwidth or DoS are not met.

In the multilink STT-EDR/DC-SDR/DP-SDR model, we assume 135 packet-switched-nodes. Because high rate OC3/12/48 links provide highly aggregated link bandwidth allocation, a sparse network topology design results among the packet-switched-nodes, that is, high rate OC3/12/48 links exist between relatively few (10 to 20 percent) of the packet-switched-nodes. Secondly, multilink shortest path selection with crankback is used throughout the network. For example, the STT EDR TE algorithm used is adaptive and distributed in nature and uses learning models to find good paths for TE in a network. With STT EDR, if the LSR-A to LSR-B bandwidth needs to be modified, say increased by delta-BW, the primary LSP-p is tried first. If delta-BW is not available on one or more links of LSP-p, then the currently successful LSP-s is tried next. If delta-BW is not available on one or more links of LSP-s, then a new LSP is searched by trying additional candidate paths until a new successful LSP-n is found or the candidate paths are exhausted. LSP-n is then marked as the currently successful path for the next time bandwidth needs to be modified.

Quality-of-Service priority queuing is modeled in the performance analyses, in which the key-services are given the highest priority, normal services the middle priority, and best-effort services the lowest priority in the queuing model. This queuing model quantifies the level of delayed traffic for each virtual network. The multilink path selection options are modeled with STT path selection, DC-SDR path selection, and distributed periodic path selection (DP-SDR). In the model of DP-SDR, the status updates, which are modeled with flooding link status updates every 10 seconds. Note that the multilink DP-SDR performance results should also be comparable to the performance of multilink centralized-periodic SDR (CP-SDR), in which status updates and path selection updates are made every 10 seconds, respectively, to and from a bandwidth-broker processor.

In routing a connection with multilink shortest path selection with 2-link STT-EDR routing, for example, the ON checks the equivalent bandwidth and allowed DoS first on the first choice path, then on current successful alternate path, and then sequentially on all candidate alternate paths. Again, each VN checks the equivalent bandwidth and allowed DoS provided in the setup message, and uses crankback to the ON if the equivalent bandwidth or DoS are not met.

In the models, the logical network design is optimized for each routing alternative, while the physical transport links and node locations are held fixed. We examine the performance and network design tradeoffs of:

- logical topology design (sparse or mesh); and
- routing method (2-link, multilink, fixed, dynamic, SDR, EDR, hierarchical, nonhierarchical, etc.).

Generally, the meshed logical topologies are optimized by 1- and 2-link routing, while the sparse logical topologies are optimized by multilink shortest path routing. Modelling results include:

- designs for dynamic 2-link routing (SDR, EDR) and multilink routing (SDR, EDR);
- designs for voice/ISDN-only traffic (VNETs 1-8 in Table A.1) and data-only traffic (VNETs 9-11);
- designs for integrated voice/ISDN and data traffic (VNETs 1-11);
- designs for fixed hierarchical routing;
- designs where all voice traffic is compressed (VNETs 1-5 and VNET 9 all use the IP-telephony traffic characteristics of VNET 9);
- performance analyses for overloads and failures.

A.1 Network design comparisons

Illustrative network design costs for the voice/ISDN-only designs (VNETs 1-8 in Table A.1), for the data-only designs (VNETs 9-11 in Table A.1), and for the integrated voice/ISDN and data designs (VNETs 1-11 in Table A.1), are given in Figures A.2, A.3 and A.4, respectively. These design costs and details are discussed further in ITU-T Rec. E.360.6.

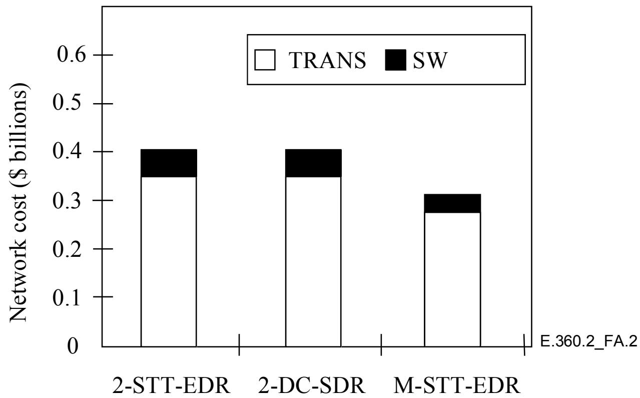

Figure A.2/E.360.2 – Voice/ISDN network design cost (includes traffic for VNET-1 to VNET-8 in Table A.1)

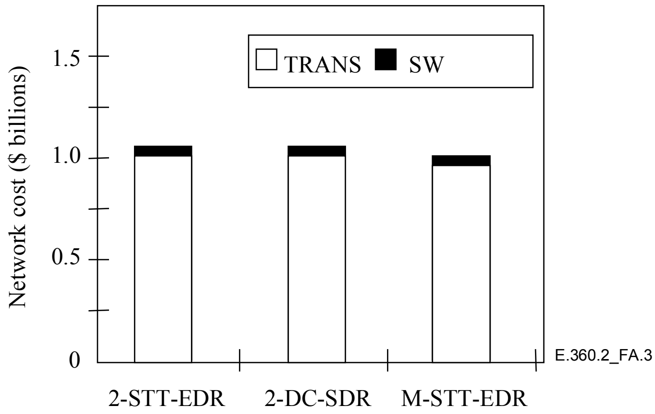

Figure A.3/E.360.2 – Data network design cost (includes traffic for VNET-9 to VNET-11 in Table A.1)

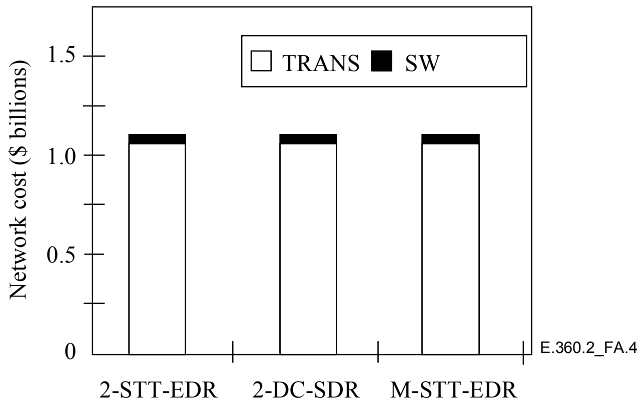

Figure A.4/E.360.2 – Integrated voice/ISDN and data network design cost (includes traffic for VNET-1 to VNET-11 in Table A.1)

The design results show that the 2-link STT-EDR and 2-link DC-SDR logical mesh networks are highly connected (90%+), while the multilink MPLS-based and PNNI-based networks are sparsely connected (10-20%). The network cost comparisons illustrate that the sparse MPLS and PNNI networks achieve a small cost advantage, since they take advantage of the greater cost efficiencies of high bandwidth logical links (up to OC48), as reflected in Table A.2. However, these differences in cost may not be significant, and can change as equipment costs evolve and as the relative cost of switching and transport equipment changes. Sensitivities of the results to different cost assumptions were investigated. For example, if the relative cost of transport increases relative to switching, then the 2-link meshed networks can appear to be more efficient than the sparse multilink networks. These results are consistent with those presented in other studies of meshed and sparse logical networks, as a function of relative switching and transport costs, see for example [A98].

Comparing the results of the separate voice/ISDN and data designs and the integrated voice/ISDN and data designs shows that integration does achieve some capital cost advantage of about 5 to 20 percent. The larger range of integration design efficiencies is achieved as a result of the economies of scale of larger capacity network elements, as reflected in cost assumptions given in Table A.2. However, probably more significant are the operational savings of integration which result from operating a single network rather than two or more networks. In addition, the performance of an integrated voice and data network leads to advantages in capacity sharing, especially when different traffic classes having different routing priorities, such as key service and best-effort service, are integrated and share capacity on the same network. These performance results are reported below. A study of voice compression for all voice traffic, such as might occur if IP-telephony is widely deployed, shows that network capital costs might be reduced by as much as 10% if this evolutionary direction is followed. An analysis of fixed hierarchical routing versus dynamic routing illustrates that more than 20% reduction in network capital costs can be achieved with dynamic routing. In addition, operation savings also result from simpler provisioning of the dynamic routing options.

A.2 Network performance comparisons

The performance analyses for overloads and failures include connection request admission control with QoS resource management. As discussed in ITU-T Rec. E.360.3, in the example QoS resource management approach, we distinguish the key services, normal services, and best-effort services. Performance comparisons are presented in Tables A.3, A.4 and A.5 for various TE methods, including 2-link and multilink EDR and SDR approaches, and a baseline case of no TE methods applied. Table A.3 gives performance results for a 30% general overload, Table A.4 gives performance results for a six-times overload on a single network node, and Table A.5 gives performance results for a single logical-link failure.

Table A.3/E.360.2 – Performance comparison for various connection-routing TE methods and no TE methods – 30% general overload (% lot/delayed traffic) (135-node multiservice network model)

Virtual network	2-link STT-EDR	2-link DC-SDR	Multi-link STT-EDR	Multi-link DC-SDR	Multi-link DP-SDR	No TE methods applied
Business-voice	0.00	0.00	0.00	0.00	0.00	3.18
Consumer-voice	0.03	0.02	0.00	0.00	0.00	2.61
INTL-out	5.40	4.82	0.00	0.00	0.00	3.62
INTL-in (key)	0.00	0.00	0.00	0.00	0.00	3.63
Key voice	0.00	0.00	0.00	0.00	0.00	3.27
64-kbit/S ISDN data	1.27	1.21	0.00	0.00	0.00	3.18
64-kbit/s ISDN data (key)	0.00	0.00	0.00	0.00	0.00	2.58
384-kbit/S ISDN data	0.00	0.00	0.00	0.00	0.00	6.51
VBR-RT voice	0.28	0.20	0.00	0.00	0.00	3.07
VBR-NRT MM	0.04	0.02	0.00	0.00	0.00	2.54
UBR MM	21.8	23.2	4.16	4.16	4.15	3.37

Table A.4/E.360.2 – Performance comparison for various connection-routing TE methods and No TE methods – 6X focused overload on OKBK (% lost/delayed traffic) (135-node multiservice network model)

Virtual network	2-link STT-EDR	2-link DC-SDR	Multi-link STT-EDR	Multi-link DC-SDR	Multi-link DP-SDR	No TE methods applied
Business-voice	5.27	2.28	0.00	0.06	0.08	9.42
Consumer-voice	7.29	3.50	0.00	0.20	0.23	13.21
INTL-out	3.43	3.36	0.00	0.00	0.04	6.03
INTL-in (key)	2.19	4.21	0.00	0.00	0.00	6.55
Key voice	0.81	1.77	0.00	0.00	0.00	8.47
64-kbit/s ISDN data	0.84	0.33	0.00	0.00	0.00	2.33
64-kbit/S ISDN data (key)	0.00	0.00	0.00	0.00	0.00	0.46
384-kbit/s ISDN data	0.00	0.00	0.00	0.00	0.00	0.00
VBR-RT voice	5.42	2.59	0.00	0.39	0.49	9.87
VBR-NRT MM	7.12	3.49	0.00	2.75	3.18	12.88
UBR MM	14.07	14.68	12.46	12.39	12.32	9.75

Table A.5/E.360.2 – Performance comparison for various connection-routing TE methods and No TE methods – failure on CHCG-NYCM link (% lost/delayed traffic) (135-node multiservice network model)

Virtual network	2-link STT-EDR	2-link DC-SDR	Multi-link STT-EDR	Multi-link DC-SDR	Multi-link DP-SDR	No TE methods applied
Business-voice	0.00	0.00	0.00	0.64	0.64	0.72
Consumer-voice	0.00	0.00	0.00	0.44	0.43	0.52
INTL-out	0.00	0.00	0.00	0.00	0.00	0.00
INTL-in (key)	0.00	0.00	0.00	0.18	0.19	0.23
Key voice	0.00	0.00	0.00	0.46	0.51	0.58
64-kbit/s ISDN data	0.00	0.00	0.00	0.95	0.89	0.94
64-kbit/s ISDN data (key)	0.00	0.00	0.00	0.00	0.00	0.00
384-kbit/s ISDN data	0.00	0.00	0.00	0.00	0.00	0.00
VBR-RT voice	0.00	0.00	0.00	0.55	0.55	0.62
VBR-NRT MM	0.00	0.00	0.00	0.44	0.42	0.51
UBR MM	2.06	1.65	0.17	0.64	0.64	0.72

In all cases of the TE methods being applied, the performance is always better and usually substantially better than when no TE methods are applied. The performance analysis results show that the multilink STT-EDR/DC-SDR/DP-SDR options (in sparse topologies) perform somewhat better under overloads than the 2-link STT-EDR/DC-SDR options (in meshed topologies), because of greater sharing of network capacity. Under failure, the 2-link STT-EDR/DC-SDR options perform better for many of the virtual network categories than the multilink STT-EDR/DC-SDR/CP-SDR options, because they have a richer choice of alternate routing paths and are much more highly connected than the multilink STT-EDR/DC-SDR/DP-SDR networks. Loss of a link in a sparely connected multilink STT-EDR/DC-SDR/DP-SDR network can have more serious consequences than in more highly connected logical networks. The performance results illustrate that capacity sharing of CBR, VBR, and UBR traffic classes, when combined with QoS resource management and priority queuing, leads to efficient use of bandwidth with minimal traffic delay and loss impact, even under overload and failure scenarios. These QoS resource management trends are further examined in ITU-T Rec. E.360.3.

The STT and SDR path selection methods are quite comparable for the 2-link, meshed-topology network scenarios. However, the STT path selection method performs somewhat better than the SDR options in the multilink, sparse-topology case. In addition, the DC-SDR path selection option performs somewhat better than the CP-DCR option in the multilink case, which is a result of the 10-second old status information causing misdirected paths in some cases. Hence, it can be concluded that frequently-updated, available-link-bandwidth (ALB) state information does not necessarily improve performance in all cases, and that if ALB state information is used, it is sometimes better that it is very recent status information.

A.3 Single-area flat topology vs. multi-area 2-level hierarchical network topology

We also investigate the performance of hierarchical network designs, which represent the topological configuration to be expected with multi-area (or multi-autonomous-system (multi-AS), or multi-domain) networks. In Figure A.5 we show the model considered, which consists of 135 edge nodes each homed onto one of 21 backbone nodes. Typically, the edge nodes may be grouped into separate areas or autonomous systems, and the backbone nodes into another area or autonomous system. Within each area, a flat routing topology exists. However, between edge areas and the backbone area, a hierarchical routing relationship exists. This routing hierarchy is modeled in ITU-T Rec. E.360.3 for both the per-flow and per-virtual-network bandwidth allocation examples, here the results are given for the per-flow allocation case in Tables A.6 to A.8 for the 30% general overload, 6-times focused overload, and link failure examples, respectively. We can see that the performance of the hierarchical network case is substantially worse than the flat network model, which models a single area or autonomous system consisting of 135 nodes.

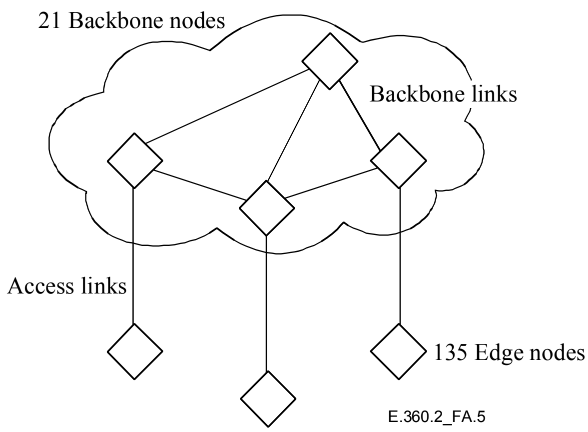

Figure A.5/E.360.2 – Hierarchical network model

Table A.6/E.360.2 – Performance of single-area flat vs. multi-area 2-level hierarchical network topology – percent lost/delayed traffic under 30% general overload (multilink STT-EDR routing; 135-node network model)

Virtual network	Single-area flat topology	Multi-area 2-level hierarchical topology
Business-voice	0.00	0.00
Consumer-voice	0.00	0.00
INTL-out	0.00	0.00
INTL-in (key)	0.00	0.00
Key voice	0.00	0.00
64-SDS	0.00	0.00
64-kbit/s ISDN data (key)	0.00	0.00
384-kbit/s ISDN data	0.00	0.00
VBR-RT voice	0.00	0.00
VBR-NRT MM	0.00	0.00
UBR MM	4.16	9.06

Table A.7/E.360.2 – Performance of single-area flat vs. Multi-area 2-level hierarchical network topology – percent lost/delayed traffic under 6X focused overload on OKBK (Multilink STT-EDR routing; 135-node network model)

Virtual network	Single-area flat topology	Multi-area 2-level hierarchical topology
Business-voice	0.00	1.70
Consumer-voice	0.00	2.22
INTL-out	0.00	0.89
INTL-in (key)	0.00	0.00
Key voice	0.00	0.00
64-kbit/s ISDN data	0.00	0.27
64-kbit/s ISDN data (key)	0.00	0.00
384-kbit/s ISDN data	0.00	0.00
VBR-RT voice	0.00	0.93
VBR-NRT MM	0.00	1.80
UBR MM	12.46	12.88

Table A.8/E.360.2 – Performance of single-area flat vs. Multi-area 2-level hierarchical network topology – percent lost/delayed traffic under failure on CHCG-NYCM link (Multilink STT-EDR routing; 135-node network model)

Virtual network	Single-area flat topology	Multi-area 2-level hierarchical topology
Business-voice	0.00	0.00
Consumer-voice	0.00	0.00
INTL-out	0.00	0.00
INTL-in (key)	0.00	0.00
Key voice	0.00	0.00
64-kbit/s ISDN data	0.00	0.00
64-kbit/s ISDN data (key)	0.00	0.00
384-kbit/s ISDN data	0.00	0.00
VBR-RT voice	0.00	0.00
VBR-NRT MM	0.00	0.00
UBR MM	0.17	1.38

A.4 Network modelling conclusions

The TE modelling conclusions are summarized as follows:

1) Capital cost advantages may be attributed to the sparse topology options, such as the multilink STT-EDR/DC-SDR/DP-SDR options, but may not be significant compared to operational costs, and are subject to the particular switching and transport cost assumptions. Capacity design models are further detailed in ITU-T Rec. E.360.6 and operational issues in ITU-T Rec. E.360.7.

2) In all cases of the TE methods being applied, the performance is always better and usually substantially better than when no TE methods are applied.

3) The sparse-topology multilink-routing networks provide better overall performance under overload, but performance under failure may favor the 2-link STT-EDR/DC-SDR options with more alternate routing choices. One item of concern in the sparse-topology multilink-routing networks, is with post dial delay in which perhaps 5 or more links may need to be connected for an individual connection request.

4) Single-area flat topologies exhibit better network performance and, as discussed and modeled in ITU-T Rec. E.360.6, greater design efficiencies in comparison with multi-area hierarchical topologies. As illustrated in ITU-T Rec. E.360.4, larger administrative areas can be achieved through use of EDR-based TE methods as compared to SDR-based TE methods.

5) State information as used by the 2-link and multilink SDR options provides only a small network capital cost advantage, and essentially equivalent performance to the 2-link STT-EDR options, as illustrated in the network performance results.

6) Various path selection methods can interwork with each other in the same network, which is required for multi-vendor network operation.

7) QoS resource management, as further described in ITU-T Rec. E.360.3, is shown to be effective in achieving key service, normal service, and best effort service differentiation.

8) Voice and data integration can provide capital cost advantages, but may be more important in achieving operational simplicity and cost reduction.

9) If IP-telephony takes hold, and a significant portion of voice calls use voice compression technology, this could lead to more efficient networks.

Overall the packet-based (e.g. MPLS/TE) multilink, sparse-topology routing strategies offer several advantages. The sparse logical topology with the high-speed switching and transport links may have economic benefit due to lower cost network designs achieved by the economies of scale of higher rate network elements. The sparse, high-bandwidth, logical-link networks have been shown to have better response to overload conditions than logical mesh networks, due to greater sharing of network capacity. The packet-based routing protocols have capabilities for automatic provisioning of links, nodes, and reachable addresses, which provide operational advantages for such networks. Because the sparse high-bandwidth-link network designs have dramatically fewer links to provision compared to mesh network designs (10-20% connected versus 90% or more connected for mesh networks), there is less provisioning work to perform. In addition to having fewer links to provision, sparse high-bandwidth-link network designs use larger increments of capacity on individual links and therefore capacity additions would need to occur less frequently than in highly connected mesh networks, which would have much smaller increments of capacity on the individual links. The sparse-topology, multilink-routing methods are synergistic with evolution of data network services which implement these protocols, and such routing methods have been in place for many years in data networks. Should a service provider pursue integration of the voice/ISDN and data services networks, these factors will help support such an integration direction.

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