# V5 Interfaces and Architecture

# 9

> "There is no need for housekeeping at all, since after four years the mess doesn't get any worse."
>
> —adapted from Quentin Crisp

The V5 interface is the interface between an access network and a host exchange for the support of narrowband telecommunications services. It has two forms, V5.1 and V5.2, and the more sophisticated V5.2 form supports the concentration of traffic. The apparent complexity of the V5 interface is a result of meeting a number of simple basic requirements, and the avoidance of unnecessary functionality has been a driving consideration. One consequence of this is that in-band tones are passed transparently across the V5 interface and the responsibility for their generation and detection remains with the host exchange, not the access network.

In this chapter the architectural model associated with the V5 interface is described, together with the physical structure of the interfaces and the services they support. Subsequent chapters describe the general multiplexing and format of V5 messages, the multiplexing and handling of ISDN signaling, and the various V5-specific protocols. The chapters on the V5 control protocol and the protocol for PSTN signaling are relevant to both V5.1 and V5.2 interfaces. The later chapters on the BCC protocol, the link control protocol, and the protection protocol are only relevant to V5.2 interfaces. Broadband services are supported by VB5 interfaces, which are introduced more fully towards the end of the book.

The V5 interface is not limited to any specific access technology or medium, although much of the motivation for its development was the anticipated deployment of optical access networks. There has been considerable interest in its use for radio access networks, which, with hindsight, should have been obvious because of the technical and commercial developments in mobile radio access. Also with hindsight, the potential use of the V5 interface for interconnection between telecommunications networks belonging to different operators should have been more obvious.

Because it ignores the details of access technology, the model of an access network appropriate for a V5 interface can appear out of proportion to the access network as a whole. From the point of view of a V5 interface, an access network is a black box with borders but no internal structure. Features that are optional and viewed as just outside the borders are more important than those that are essential but not relevant to the V5 interface. For example, the optional transmission systems at the borders are of more concern than the internal transmission system of the access network.

To understand the V5 interface, it is necessary to first understand the nature of the boundaries of the access network. These include the physical links to the exchange, the physical links to the remotely located user ports, and the services supported at the user ports. Once the nature of the boundaries are clear, the details of the V5 protocols and the ways they are multiplexed together can be examined in context.

## 9.1 THE V5 ACCESS MODEL

The architectural model for the V5 specification differs from other models sometimes used for access networks because it is focused on the aspects relevant to the V5 interfaces and ignores those details that are only relevant to the specific access technologies. The basic model is shown in Figure 9.1.

It is necessary to distinguish between the *local line distribution network* (LLDN), which stretches from the host exchange in the core network to the *customer premises equipment* (CPE), and the access network as it is defined for the V5 interface. The difference between the two is that the LLDN also includes any *feeder transmission systems* (FTS) and *remote digital sections* (rDS) if they are present.

An FTS allows the headend of an access network to be located remotely from its host exchange. The FTS could take the form of an SDH ring with

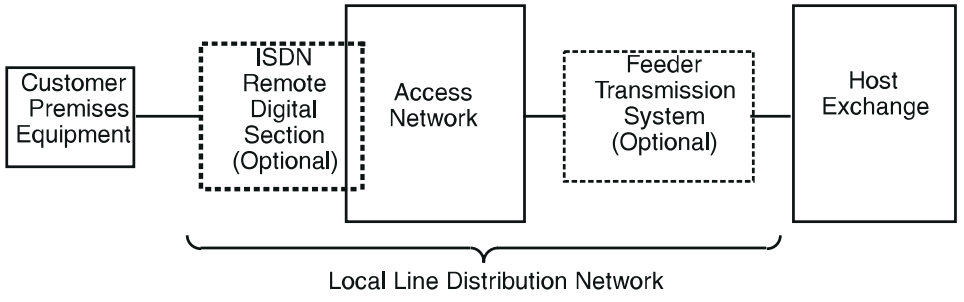


**Figure 9.1** The V5 access model.

add-drop multiplexers at the host exchange (or host exchanges) and also at the locations of the headend of the access network. An FTS extends the range of the access network. Transmission delays permitting, an FTS can allow an access network to be located in one country or continent and the host exchange to be located in another.

There is no reason why the functionality of an FTS should not be included in the access network itself. This could be achieved by creating a two-stage access network. The first stage would transport the traffic payloads between the host exchange and various remote locations. The second stage would transport smaller payloads between the remote locations and the final destinations. It would also be possible to carry out the entire transmission in a single step with a very sophisticated optical transmission system, which would have to allow diversity of routing to the remote locations to ensure security and be able to operate with a wide range of distances to the remote ends. If the functionality of the FTS is included in the access network, then there is no FTS functional block in that particular case of the access model, since the FTS has no separate existence. If an FTS is present, then it must be managed as a separate entity, distinct from the access network. If the functionality of an FTS exists, but it is managed as part of the access network, then there is no genuine FTS present, only a more sophisticated access network.

It is also possible to have an rDS between the CPE belonging to an ISDN user and the true access network. In many cases, there is no need for an rDS, because this functionality is included within the access network. The rDS is a legacy of the *digital section* (DS) initially used for ISDN deployment, which consists of an ISDN line termination in the exchange with some form of digital transmission to a remote NT1. A DS is under the control of its associated exchange and is not managed independently of it as part of the access network.

An rDS exists if the line termination of a DS is physically removed into the access network, but the associated control remains in the exchange. In particular, the exchange must be able to activate and deactivate the transmission in the rDS and to monitor the performance of the rDS. If the functionality of an rDS is present, but it is managed as part of the access network, then there is no genuine rDS present, because it is not controlled and monitored by the exchange.

## 9.2 SERVICES AND USER PORTS

It is important to distinguish between the types of user port and the types of services supported at user ports, particularly for ISDN user ports, because physical ports can support different services.

There are four generic service types that can be supported at a user port associated with a V5 interface, although only up to three of these can be supported simultaneously. The first generic type of service is on-demand serv-

ice, for either ISDN or PSTN, where the connection is set up by the host exchange at the start of each call. In addition to on-demand service, there are two types of leased service. A leased service differs from an on-demand service in that the connection is created by the configuration of the network and not set up for the individual calls.

The first type of leased service is permanently leased service. This is handled by a leased-line network that is separate from the host exchange. User ports that only support permanently leased services are independent of V5 interfaces, because there is no association with a host exchange. The second class of leased service is semipermanent service, where the traffic is routed through a host exchange via a V5 interface, but the connection is set up by network configuration and not for each call. The V5 interface only allows 64-Kbps B-channels to be used for semipermanent services, because narrowband host exchanges are designed to connect 64-Kbps channels. Figure 9.2 summarizes the different generic service types.

User ports that are associated with a V5 interface and that support on-demand services are classified as either PSTN or ISDN user ports, regardless of any leased services supported at the user port. This distinction is only significant for ISDN user ports, because a user port can only be classified as a PSTN port if it supports on-demand PSTN service, and this leaves no channel available for any leased services. The D-channel of an ISDN user port is always connected to the host exchange over the V5 interface because the D-channel contains the call control for the on-demand services. A V5.1 interface can only support basic rate ISDN because it does not have sufficient capacity for standard primary rate ISDN. A V5.2 interface can support both basic rate and primary rate ISDN. Figure 9.3 summarizes the relationship between user ports and services.

User ports that are associated with a V5 interface and that do not support on-demand services are classified as leased ports. These must support either semipermanent service or a combination of semipermanent and permanently leased services, because a leased port that only supports permanently leased service is not connected to an exchange over a V5 interface. Leased ports that only require a single 64-Kbps B-channel on the V5 interface are handled in a similar way to PSTN ports. Leased ports that require more than one B-channel are handled in a similar way to ISDN ports.

| On-Demand | Leased |
|---|---|
| PSTN | Semipermanent |
| ISDN | Permanently Leased |

**Figure 9.2** Generic service types.

| SERVICES | PORTS | | |
|---|---|---|---|
| | PSTN | ISDN | Leased |
| PSTN | Mandatory | n/a | n/a |
| ISDN | n/a | Mandatory | n/a |
| Semipermanent | n/a | Optional | Mandatory |
| Permanently leased | n/a | Optional | Optional |

**Figure 9.3** Services supported at different ports.

The architectural model for services and ports is shown in Figure 9.4. This model ignores the possible FTS, because a FTS is transparent and so has no effect on ports and services. A distinction is made between ISDN ports with and without an rDS, because they have different interfaces at the access network. The nature of the interface to remote NT1s for ISDN is not specified, because it is not internationally standardized. The interface to leased-line equipment is not specified because it also is not standardized.

The host exchange supports the on-demand services (PSTN and ISDN) and the semipermanent leased services. The leased-line network supports the permanently leased services.

## 9.3 V5 LINKS AND TIME SLOT STRUCTURE

The V5 interface can take two forms, V5.1 and V5.2. A V5.1 interface consists of a single 2.048-Mbps link. A V5.2 interface consists of between one and

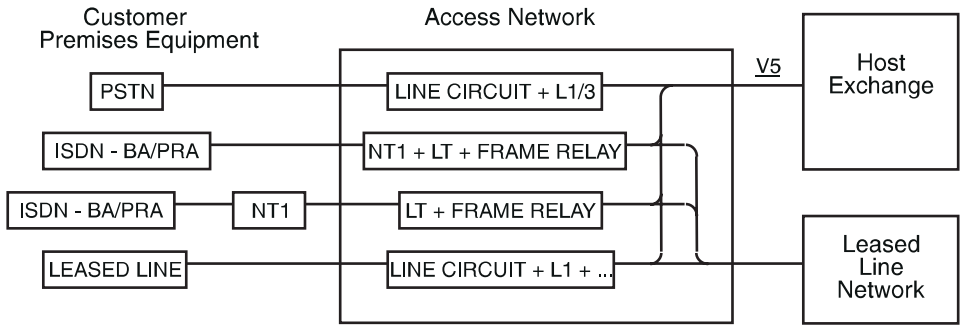


**Figure 9.4** Service architecture.

sixteen 2.048-Mbps links, although in practice a single-link V5.2 interface would be exceptional, because the increased amount of traffic a V5.2 interface supports makes additional links sensible for greater security. In addition to the functions of the V5.1 interface, the V5.2 interface supports both the concentration of traffic and the dynamic assignment of time slots. The 2.048-Mbps links for both interfaces are formatted in the normal way into 32 time slots, with time slot 0 used for frame alignment (see Figure 9.5). A single V5.1 interface can support up to 30 PSTN ports (or 15 ISDN basic access ports), while a single V5.2 interface can support several thousand ports. In both cases, both PSTN and ISDN ports may be supported on the same 2.048-Mbps link.

The V5 interface contains a number of different communications protocols. These are divided into housekeeping communications protocols (control, link control, bearer channel connection, and protection) and call control communications protocols (both for PSTN and ISDN). The call control protocols and the V5 control protocol are relevant to both the V5.1 and the V5.2 interfaces, but the other housekeeping protocols are only relevant to V5.2 interfaces.

ISDN communications are grouped into P-type, F-type, and S-type communications paths. These paths correspond to packet data (SAPI 16), frame data (SAPI 32 to 62), and D-channel signaling (other SAPIs), respectively. Each type of ISDN communication from a single user port is mapped onto a common communications path for that type of information, and each path has an associated V5 communications channel corresponding to a V5 time slot. No two communications paths of the same type can share the same V5 time slot, since communication paths of the same type are only differentiated because they use different time slots. A single ISDN user port always uses the same V5 time slot for each of the three types, but it can use different V5 time slots for different types. Different ISDN user ports may use different communications paths on different V5 time slots for the same type of communication.

Unlike the ISDN communication paths, the housekeeping protocols always share the same V5 time slot. This is time slot 16 of the first 2.048-Mbps link. The PSTN call control protocol also only uses a single time slot, but neither it nor the ISDN communications paths are forced to share the time slot used by the housekeeping protocols to allow extra bandwidth to be allocated to

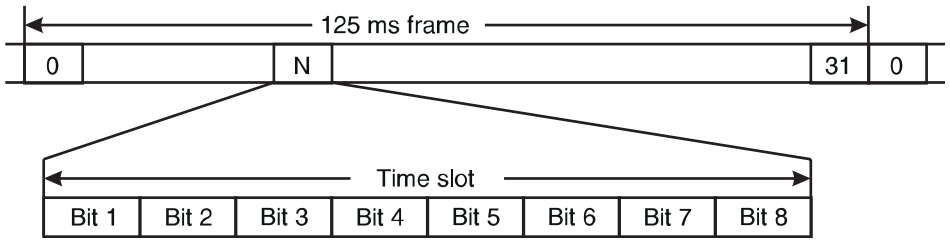


**Figure 9.5** Format of 2.048-Mbps links at V5 interfaces.

call control as the number of user ports increases or as ISDN D-channel traffic increases.

### 9.3.1 The V5.1 Interface

For V5.1, there is only a single S-ISDN (signaling) communications path with a corresponding unique V5 time slot, which may or may not be shared by other communications protocols or different types of ISDN communications paths. There may be a number of different P-ISDN (packet data) and F-ISDN (frame relay) communications paths making use of up to three time slots. If there is only one time slot used for all communications channels, then it must be time slot 16, because the control protocol is located there.

If two time slots are used for communications, then they must be time slots 16 and 15. The control protocol must use time slot 16 and at least one of the other communications paths must obviously use time slot 15 (see Figure 9.6). For instance, there may be F-ISDN communications paths on both time slots. Likewise, there may be P-ISDN communications paths on both time slots. The PSTN protocol and the S-ISDN communications path may each use either time slot 16 or 15.

If three time slots are used for communications, then these are time slots 16, 15, and 31 (see Figure 9.7). Once again the control protocol must use time slot 16. Since the PSTN protocol can only use a single time slot, there must also be ISDN communications present if three time slots are used. There may be F-ISDN and P-ISDN communication paths on any of the time slots. The PSTN protocol and the S-ISDN communications path may each use either time slot 16, 15, or 31.

### 9.3.2 The V5.2 Interface

In addition to the difference in the number of 2.048-Mbps links, the V5.2 interface differs from the V5.1 interface in two major ways. First, the V5.2

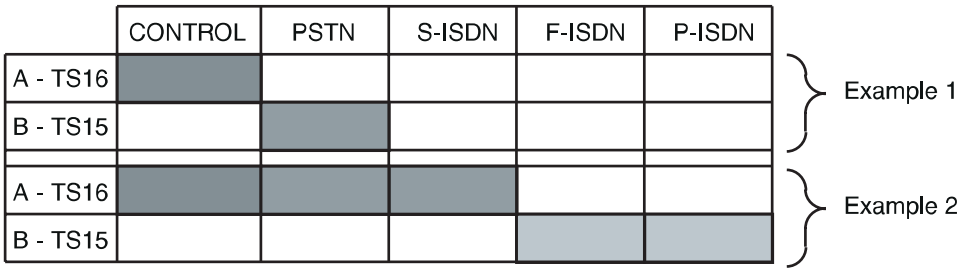


**Figure 9.6** Possible assignments for V5.1 with two communications time slots.

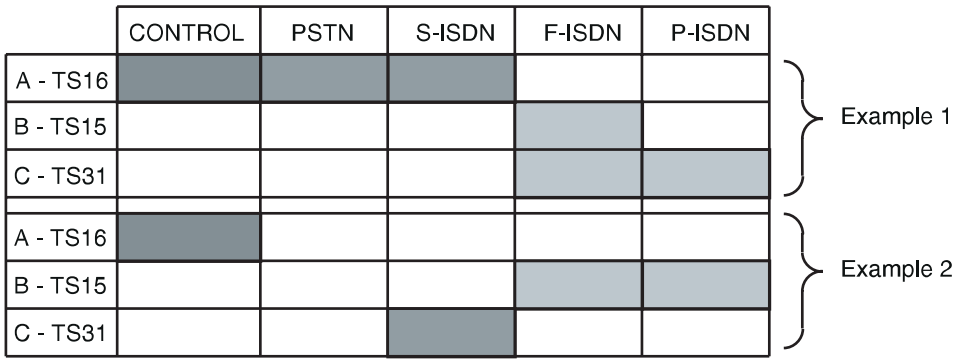


**Figure 9.7** Possible assignments for V5.1 with three communications time slots.

interface supports additional housekeeping protocols, which share the same time slot as the control protocol. The second major way in which the V5.2 interface differs from V5.1 interface is that it has additional backup time slots to improve the security of the communications (see Figure 9.8). Apart from these major differences, V5.2 differs from V5.1 in that there may be more than one S-ISDN communications path, so that ISDN call control is not constrained to a single V5 time slot. This is necessary to allow extra bandwidth to be allocated to call control for the additional ISDN ports the V5.2 interface can support. Extra bandwidth for the additional PSTN ports was not thought to be necessary, since PSTN signaling should be less demanding on bandwidth.

The introduction of the additional housekeeping protocols may have an indirect impact on the allocation of communications paths to time slots, because they reduce the spare capacity on the time slot used by the control protocol. Because of the presence of the additional protocols, call control

| Timeslots | Primary Link | Secondary Link | Other Links |
|---|---|---|---|
| 15 | (optional) | (optional) | (optional) |
| 16 | Housekeeping Protocols | Housekeeping Protection | (optional) |
| 31 | (optional) | (optional) | (optional) |

**Figure 9.8** V5.2 communications time slots.

communications are less likely to also share the same time slot, particularly if they are heavily used. There is no direct impact on the allocation of protocols to time slots due to the introduction of the additional housekeeping protocols, because all housekeeping protocols are effectively allocated to a single communications time slot as a single compound protocol.

The use of backup channels has an effect similar to the allocation of additional time slots to a communications path. The net effect is to increase the number of time slots associated with the communications. The only change to principles involved in the V5.1 case is that backup makes it necessary for a communications path to be dynamically associated with more than one time slot, preferably on different links in case there is a failure on a specific link.

#### *9.3.2.1 Protection of Communications Channels*

The V5.2 interface has the ability to automatically protect the logical communications channels that carry signaling and housekeeping protocols between the access network and the host exchange. This feature allows the V5.2 interface to survive the failure of one of its component links, because the communications on the failed link can be automatically switched to another link. This assumes, of course, that a V5.2 interface consists of at least two links.

Protection is performed for specified 64-Kbps logical communications channels and includes all communications paths assigned to that channel. Protected communications channels belong to either protection group 1 or to protection group 2.

Protection group 1 handles the logical communications channel containing the housekeeping protocols and uses time slot 16 on both the primary and the secondary V5.2 links (see Figure 9.8). This logical communication channel is the main V5.2 communications channel and the two time slots are the physical communications time slots with which it can be associated. Initially the main communications channel is associated with time slot 16 on the primary V5.2 link.

The V5.2 protection protocol monitors time slot 16 on both the primary and the secondary V5.2 links. This ensures that degradation on the primary link is detected and that the availability of the secondary link is ensured. If the performance on the primary link drops too far, then the main logical communication channel is switched to time slot 16 of the secondary link. It is possible that a few messages will be corrupted during this switch-over, but this is not significant, because the corruption will be detected and the messages retransmitted.

Logical communications channels other than the main channel can be protected by including them in protection group 2. Protection group 2 differs from protection group 1 in that protection group 2 does not have a standby time slot for each active time slot and the V5.2 protection protocol is not transmitted on the backup time slots. There are no more than three standby time slots in

protection group 2, because three are sufficient to provide protection in the event of a single-link failure. Any number of active communications time slots, which are not otherwise protected, may be assigned to protection group 2. Logical communications channels in protection group 2 will be switched to a free standby time slot in the event of a failure of their original time slots. As for protection group 1, any messages that are corrupted as a result of the switch-over will be retransmitted as part of the normal error correction process.

The protection protocol coordinates the switch-over for both protection groups, so that both sides of the interface switch in the same way. Logical communications channels other than the main communications channel can be left unprotected.

#### *9.3.2.2 Control of the 2.048-Mbps Links*

The multiple links of a V5.2 interface are managed across the V5.2 interface by its link control protocol. This protocol allows links to be identified and blocked or unblocked.

Link identification is required to verify the integrity of the physical connections of the V5.2 interface. It operates by tagging the link to be identified, in the same way that a tone may be applied to a specific copper pair in cable. In this case a digital signal is applied instead of a tone.

Link blocking and unblocking are required to allow links to be maintained with the minimum disruption of traffic and to allow the interface to evolve as the traffic evolves. This is almost identical to the blocking and unblocking of user ports on the control protocol.

## 9.4 BEARER TIME SLOTS AND V5 TRAFFIC CAPACITY

The bearer time slots on a V5 interface are used to carry 64-Kbps circuit-switched traffic from the user ports to the host exchange. These time slots must be allocated to the user ports in a way that is clearly agreed-on, so that both the access network and the host exchange know which time slots are used for a particular user port.

For a V5.1 interface, the allocation of bearer time slots to user ports is static, but may be reconfigured over the management interfaces at the access network and at the host exchange. Here the term *static* is used in the sense that the allocation does not change from call to call. There is a one-to-one mapping between the appropriate bearer channels at the user ports and the bearer time slots on V5.1 interfaces. For a simple-access network, this mapping may be hard-wired.

For a V5.2 interface, the allocation of bearer time slots to user ports is dynamic and will normally change from call to call. The mapping between the

V5-related bearer channels at the user ports and the bearer time slots on the V5.2 interface is controlled by the V5.2 BCC protocol. Bearer time slots are allocated to user ports flexibly according to demand. This flexibility gives greater security and supports the concentration of traffic.

The dynamic allocation of bearer time slots on a V5.2 interface gives greater security, because service is maintained even if a link is lost. This requires, of course, that there is more that one link on the V5.2 interface. Individual calls may be lost if a V5.2 link fails, but these can be reestablished on a different link if the user redials. The quality of service after a failure will be less because the traffic is supported by fewer time slots. This increased security is impossible on a V5.1 interface, because its static allocation ties the service to the lost bearer time slots.

The dynamic allocation of bearer time slots on a V5.2 interface also supports the concentration of bearer traffic. The interface can support more bearer channels at the user ports than bearer time slots on the V5.2 interface. Concentration takes advantage of the fact that only a fraction of all user ports are likely to be active at any given time. For reasonably large systems, a concentration factor of about 8 is routinely applied, because it creates no perceived reduction in the quality of service. This allows an access network with about 1,000 PSTN ports to be supported by a single V5.2 interface with only four links.

A single V5.2 interface is capable of supporting about 4,000 PSTN ports, because it may have 16 links and concentrate the bearer traffic by a factor of 8. A single V5.1 interface is only capable of supporting 30 PSTN ports, because at least one time slot on the link is required for signaling and another is required for frame alignment.

Dynamic allocation of V5 bearer time slots is not identical to the concentration of bearer traffic, because dynamic allocation does not determine the ratio of user ports to bearer time slots on the V5 interface. In theory, dynamic allocation can be used with more time slots than those required if all user ports were busy, but this is unrealistic, except perhaps when a V5 interface is newly installed and only a few customers have been allocated to it. It is not necessary for a V5.2 interface to be concentrating, but it must have dynamic allocation of bearer time slots.

Concentration of traffic across the V5.2 interface is different from concentration of traffic within the access network itself. The transmission system of the access network need not concentrate traffic, even if the V5.2 interface is concentrating. A concentrating V5.2 interface may be used regardless of the details of the transmission within the access network, because the cost per user port of the interface can be less despite the greater complexity of the interface, since fewer links are required to support the traffic. Even if the access network is concentrating, the concentration can be hidden within the access network, which could use a nonconcentrating V5.1 interface. For example, an access

network using radio transmission may use concentration on the transmission because radio bandwidth is limited, but use a nonconcentrating V5.1 if the size of the system is too low to justify the added complexity of a V5.2 interface.

## 9.5 SUMMARY

The V5 interface is the interface between an access network and a host exchange for the support of narrowband telecommunications services. It has two forms, V5.1 and V5.2, and the more sophisticated V5.2 form supports the concentration of traffic. The architectural model of the access network used in the definition of V5 interfaces allows an access network to be linked to its host exchange by an FTS and it allows ISDN ports to be connected to the access network by an rDS. Although both FTSs and rDSs are permitted, neither is required.

Both V5.1 and V5.2 interfaces support both PSTN and ISDN user ports, but a V5.2 interface is required if primary rate ISDN is to be supported. V5 interfaces also support leased-line ports with semipermanent leased lines connected via the exchange by treating them like PSTN or ISDN lines. The V5 interface also allows an access network to have ISDN ports with certain B-channels nailed up for leased service. The traffic on these nailed-up B-channels may be routed through a service interface that is not a V5 interface, for instance, to a leased-line network.

V5.1 interfaces consist of a single 2.048-Mbps link, while V5.2 interfaces may consist of up to 16 such links. Both types of V5 interfaces may use time slots 15, 16, and 31 for signaling, but there are detailed constraints on the allocation of signaling to these time slots.

V5.2 interfaces are able to concentrate bearer traffic so that the bearers at the user ports are supported by a smaller number of bearer channels on the interface. The allocation of bearer channels to user ports on a V5.2 interface may change with each call, unlike the allocation for a V5.1 interface. This dynamic allocation of bearer channels also makes the V5.2 interfaces more secure. It does not determine the amount of concentration at the interface, and whether or not concentration is used at the V5.2 interface is not determined by whether concentration is or is not used within an access network.

V5.2 interfaces also allow their individual links to be identified so that the integrity of the interface can be checked. Individual links can also be removed from service, either temporarily for maintenance or permanently to make the interface smaller. Likewise, new links for a V5.2 interface can be brought into service.

The V5.2 interface provides for protection of the time slots reserved for signaling. The time slot used for the most sensitive protocols is always protected—the active and standby time slots forming protection group 1. Other

signaling time slots may also be protected, and their active and standby time slots form protection group 2.

## Selected Bibliography

*V5.1 Interface Specification for the Support of Access Networks*, ITU-T Recommendation G.964.

*V5.1 Interface Specification for the Support of Access Networks*, ETSI Specification ETS 300 324-1.

*V5.2 Interface Specification for the Support of Access Networks*, ITU-T Recommendation G.965.

*V5.2 Interface Specification for the Support of Access Networks*, ETSI Specification ETS 300 347-1.