# Application of Control Plane Technology to Dynamic Configuration Management

***Alan McGuire, Shehzad Mirza, and Darren Freeland***

***BTexact Technologies***

## ABSTRACT

The distinction between switched-services-based and leased line services is beginning to disappear. Many network operators and suppliers are developing control plane technology for application in transport networks. This will allow faster service provisioning, particularly between network operators, and the creation of new network services. However, such systems will still require comprehensive management systems, and successful operators and vendors of the future will be those that are capable of developing operational support systems that complement the control plane with service management capabilities, automated plan and build processes, inventory management, and capacity planning. This article examines the distribution of functionality between the management and control planes for support of soft permanent and switched connections.

## INTRODUCTION

The distinction between switched and leased line services continues to blur. Historically, switched services have been considered connections that are set up and torn down using a control plane, while the setup and teardown of leased line services is by means of network management protocols. In technological terms, such a distinction is rather artificial; in many ways the only reason for it remaining is an inertia that divides this function into transmission and switching. This distinction is now being challenged in the marketplace. Many network operators and vendors now wish to apply control plane technology to the emerging optical network [1]. For some, optical networking is seen as an opportunity to carry IP directly on top of a high-capacity transport network, thereby simplifying the network structure. For others there is the commercial reality that, first, optical networking should be able to support any client layer technology, and second, multiple layers are needed to create large-scale networks with quality of service guarantees. In reality, both types of network will be deployed, but the application of control plane technology is not optical-network-specific; it can be applied generically to any layer of the transport network, from synchronous digital hierarchy (SDH) VC-12 connections to 10 Gb/s wavelengths. This article begins by describing the relationship between the management and control planes, and how functionality should be distributed between them. This will depend on the extent to which end users have access to bandwidth on demand or "dialup bandwidth." We suggest that this might not occur in the near future, and that control planes will be used predominantly for connection control by carriers rather than end users. Nevertheless, successful large-scale networks require a marriage between control and management.

## NETWORK AND CONFIGURATION MANAGEMENT

The concept of a switched transport network is not new; it was previously proposed during development of the synchronous digital hierarchy (SDH) standards around 1989 [2]. Indeed, one can look at many of the requirements originally included for the management of the SDH network [3] and conclude that they are equally applicable to a switched network (the intent was that SDH could be switched, but this was never pursued). They include:

- An ability to set up SDH VC paths between client access points automatically on request and across operator boundaries. The client will generally be another network layer, but may, in the case of leased line services, be an end user.
- An ability to maintain these paths to a very high availability, restoring failed paths automatically if appropriate to the quality of service.

0163-6804/01/$10.00 © 2001 IEEE

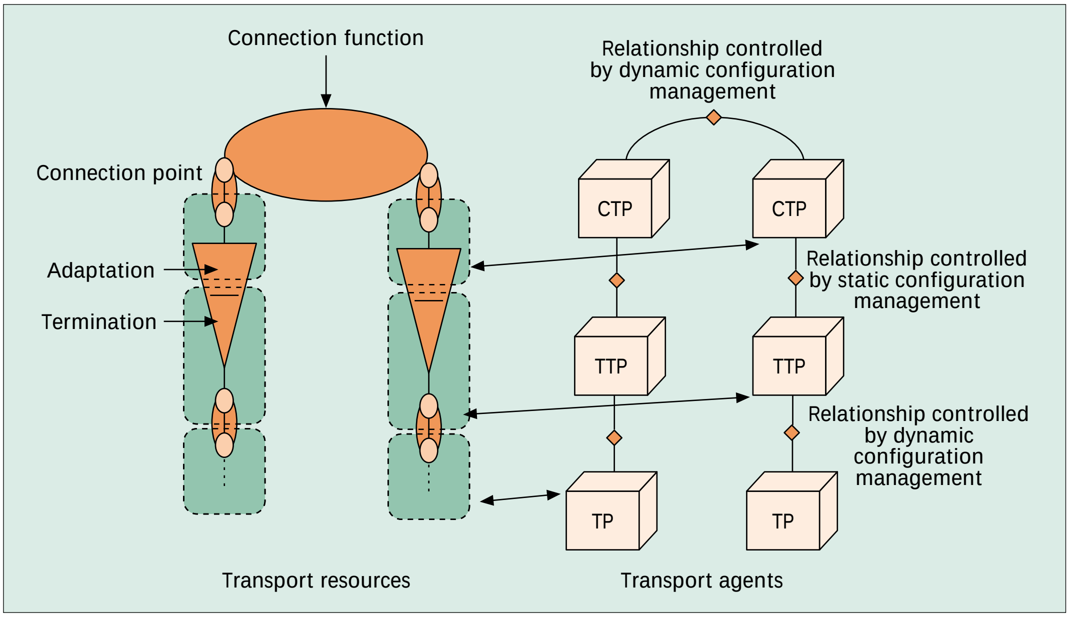

**Figure 1.** *Relationship between transport resources and transport agents and their role in configuration management. CTP: connection termination point; TTP: trail termination point; TP: termination point.*

*The functional architecture of the transport network describes the basic transport functions in a manner that makes no reference to the control and management of those functions.*

- An ability to continuously monitor performance of allocated paths, while in service, and validate compliance with service commitments.
- The capability to generate resource utilization information to support routing and billing between operators, and planning and cost accounting within a domain.

Of these requirements the first can be achieved where each layer network access point has an identity that:

- Is globally unambiguous
- Is available to other network operators for purposes of cross-domain path setup
- Identifies the country and network operator who is responsible for routing to and from the access point

To meet these requirements the SDH path and section layer networks should each have independent access point identifier schemes.

SDH connections today are predominantly created by management systems rather than control planes. Despite claims that management-system-based solutions are slow, large-scale automated operational support systems developed by some network operators are capable of creating hundreds of circuits a day with connection setup taking minutes per connection. In many cases problems in the provisioning cycle lie elsewhere.

If control plane technology were used in the SDH network (and future optical networks), what would its role in configuration management be?

Configuration management can be subdivided into the following activities: configuration resource management (which can be considered static provision/configuration of resources) and configuration connection management (which can be considered dynamic configuration management).

Configuration resource management is concerned with:

- Provisioning of access points
- Provisioning of access groups
- Configuration of access groups
- Provisioning of connection points
- Configuration of connection points
- Provisioning of subnetworks
- Provisioning of links
- Configuration of links
- Provisioning of link connections
- Provisioning of a layer network

Configuration connection management is concerned with:

- Subnetwork connection setup
- Release of subnetwork connections
- Setting up network connections
- Release of network connections

Configuration resource management is concerned with attachment and detachment of resources to a layer network, while configuration connection management is concerned with establishment and disconnection of paths and trails within a layer network. This separation can be used for constructing the architecture of the control plane and describing interaction with the management plane [4, 5].

In terms of the transport network architecture of ITU-T Recommendation G.805 [5], static configuration management configures the adaptation function, while dynamic configuration management configures the connection functions, as shown in Fig. 1.

The functional architecture of the transport network describes the basic transport functions in a manner that makes no reference to the control and management of those functions. For the purposes of control and management, each transport function has a closely coupled agent that represents the role it has to play. The transport function control agents interact with other functions that are participating in the control and management enterprise through interfaces, and present information or execute operations as required. The relationships between transport functions are manipulated as follows:

*Multiple client layer addressing schemes must be supported, so the server layer optical network addressing should be disjoint from any client layer addressing to ensure a true multiclient server network.*

- Client layer connection termination points (CTPs) and server layer trail termination points (TTPs) are controlled by static configuration management.
- Two CTPs within the same layer network and the same subnetwork are controlled by dynamic configuration management.
- A TTP and a CTP within the same layer network are controlled by dynamic configuration management.

Currently, connection control is achieved using network management protocols within a single domain (although such a domain may be large, e.g., 25,000 network elements). Network management protocols are excellent for vertical (centralized) relationships, such as network element to element manager to network manager, but are less suited to horizontal relationships between peers, as required in connection control. Hierarchical network management of connections spanning several domains requires either an overseeing organization with global visibility or mechanisms that provide a high degree of security for transferring and sharing information. By contrast, control planes confer no special status to any single operator involved in the process. Protocols designed specifically for the control plane with optimized message sequences should improve provisioning times. As such, the control plane is potentially a more promising candidate for global connectivity. Network management protocols will, however, continue to be used for other applications such as fault and performance management (and connection management for permanent circuits). In the following sections we consider the functions of the control plane.

## THE CONTROL PLANE

The main functions of an optical control are targeted toward solving the problem of "find, route, and connect." To achieve this, the following are required:

- A well defined naming and addressing scheme (find)
- A routing process to handle topology/resource usage and route calculation (route)
- A signaling network that provides communications between entities requesting service and those that provision those services
- A signaling protocol for the setup, maintenance, and teardown of trails (connect)

We consider these aspects below.

## NAMING AND ADDRESSING

Transport networks provide a service to client layer networks whereby a client layer link is supported by a server layer trail. The number of trails per second that are set up in the future network may in the long term be equivalent to connection volumes in existing large-scale SDH/synchronous optical network (SONET) networks (~ 500 a day). Control planes should be designed to receive requests to set up hundreds of trails per day per domain, with the potential to scale to many more in the future. To allow for the possibility of global connectivity, the transport network must have a scalable naming and addressing scheme, capable of meeting expected demand for new names and addresses over decades to come. In the transport network the access points that delimit a layer network are bound together to form trails, and as such are entities with addresses.

A proportion of carriers will require transport networks that support multiclient environments — not just IP. In reality routing, addressing (where you are) and naming (who you are) should not be considered in isolation. Multiple client layer addressing schemes must be supported, so the server layer optical network addressing should be disjoint from any client layer addressing to ensure a true multiclient server network.

This does not imply that the same type of addressing cannot be used in both the client and optical layer control planes; it simply means address spaces for each layer must be disjoint from one another.

Addresses in the transport network can be divided between those that are public and refer to endpoints, and those used by providers for internal addressing. The requirement for the former is that the address space must provide globally unique addresses (within a layer network), support flexible summarization or aggregation, and be scalable to support very large numbers of end points. Provider internal addresses represent the internal resources of the network and are not visible to users of the network. In principle the internal addressing scheme need not match that of the endpoints, but must be scalable to support any network operator's slice of the endpoint address cake.

## SIGNALING

In connection-oriented networks, connections are established prior to information transfer; this and call admission control distinguish them from connectionless networks. The signaling system can be separated into two parts:

- Communications between end systems and ingress/egress switches. This form of signaling occurs at a user–network interface (UNI).
- Communications between switches in the network that set up and tear down the end-to-end path through the switches. This form of signaling occurs at a network–network interface (NNI).

Both UNIs and NNIs should be capable of supporting dynamic connection requests. The requesting party must specify the connection parameters required (e.g., unidirectional or bidirectional, bandwidth, signal type). The control interface should support some form of call admission control (CAC) to provide authentication and permissibility of the user and verification of the service-level parameters, and inform the requesting party of whether or not the connection request was successful. This may or may not require resolution of user names into routable network addresses. In the event of an unsuccessful attempt, the requesting party should be informed of the particular reasons the request was rejected (network busy, authentication failure, etc).

Signaling within a network can be carried in either a channel associated or common channel form. The former is a simple method used in early transmission systems, where limited numbers of signaling bits were added to the frame structure (which also contained the associated traffic channel; i.e., the user-plane and control-plane share a common routing). The disadvantage is that the channel cannot readily be extended to include advanced service features and is also not generally compatible with transparent optical networking. It is desirable for common channel signaling to transfer messages in channels that are external to the user plane, which reduces the number of signaling interfaces since signaling information for multiple connections can now be statistically multiplexed into a single channel. This can be achieved using an optical supervisory channel, which is a wavelength that is terminated or digitally processed at every switch node (out-of-fiber mechanisms can also be used where appropriate). The signaling network can pass through the physical links between switches in the optical network using the supervisory channel. However, if the signaling messages are always carried along the same physical links (associated signaling) as the traffic, a failure in a physical link will result in loss of both the control and user planes. Furthermore, within a link an optical supervisory channel can fail independent of the traffic channels, since it has its own associated transmitter and receiver. Failure of the optical supervisory channel must not have any traffic-affecting consequences in the user plane, such as generating protection switching events. At the same time it must still be possible to communicate information regarding the status of individual connections within the affected link. This can be achieved using rerouting of the signaling messages onto other links. In general, signaling messages do not need to follow the traffic route, and the design of the signaling network should be such that it is more reliable than that of the user plane. Indeed, at the lowest layers of the network this is a critical consideration in the construction of the signaling network.

While the above applies to switches within the network, it may well be the case that network elements at the ingress and egress points of the network are opaque to the user plane (i.e., perform optical-electronic-optical conversions), and can therefore easily access channel associated signaling. Signaling for a UNI may therefore be channel associated or common channel (in or out of fiber).

## ROUTING

Routing can generally be divided into three basic schemes: hierarchical, source-based, and step-by-step. Taking into account speed, flexibility of routing algorithms, control of routes, and the ability to support disjoint connections and variations in transport topology, source routing is potentially the most attractive solution.

Dynamic routing protocols can perform resource discovery in addition to their standard route computation function. This allows the control plane to react dynamically to changes in network state/loading, maintain up-to-date routing tables, and (in certain scenarios) give a particular node visibility of network topology and available resource within its domain.

In a network of electronic nodes (or all-optical nodes surrounded by a wall of transponders connecting them to optical line systems) there is no need to distribute on a networkwide basis information regarding the analog properties of a link since they are compensated for on a link-by-link basis. In the all-optical case transmission impairments accumulate end to end and need to be considered as constraints when calculating the route. This is not dissimilar in concept to the original transmission plans used between telephone exchanges. Two sets of parameters need to be considered: first, those that are more or less static such as fiber type, length, dispersion, and so on. These can be stored in some form of a database. Second are those parameters that are dynamic, such as signal-to-noise ratio and signal level; unfortunately these only give information regarding the state of active channels, and the impact of a new channel on such a dynamic system can only be inferred. Many operators may choose to test the newly provisioned channel prior to releasing it to a customer.

It is worth noting that unless optical mid-span meet is achieved between vendors in all-optical networks, it is questionable if it is necessary to standardize the messages related to analog engineering within such an optical island.

In the context of the transport network, resource discovery is only supported at NNIs within a network operator's domain. Operators are unlikely to allow another operator or a private domain visibility of either topology or resources (other than reachability). Therefore, topology/resource discovery and analog engineering information is unlikely to be supported between different administrative domains.

*In the context of the transport network, resource discovery is only supported at NNIs within a network operator's domain. Operators are unlikely to allow another operator or a private domain visibility of either topology or resources.*

## SOFT PERMANENT CONNECTIONS

One method by which control plane technology can be introduced into the network is through support of soft permanent connections. This connection type is similar to a switched connection in that routing, setup, and teardown are provided by the control plane, but the endpoints of the switched connection are contained within the network and do not cross a UNI. Preprovisioned tails are provided at the edge of the network, as illustrated in Fig. 2. Furthermore, the management system initiates the connection request rather than the user. A soft permanent circuit should appear no different from the user's perspective than a management-controlled permanent connection. The introduction of soft permanent connections requires the development of an NNI but avoids the need to introduce a UNI, and thus avoids the need to develop a new commercial interface between the network and the user. This means that no new billing systems or charging mechanisms need be introduced, public name/address resolution mechanisms are avoided, and no new security/authentication mechanisms between the network and the end user are required.

Soft permanent connections may also be used with network protection and restoration mechanisms. In the case of protection, dedicated capacity is assigned to both a working and a protection path within the network. Network protection

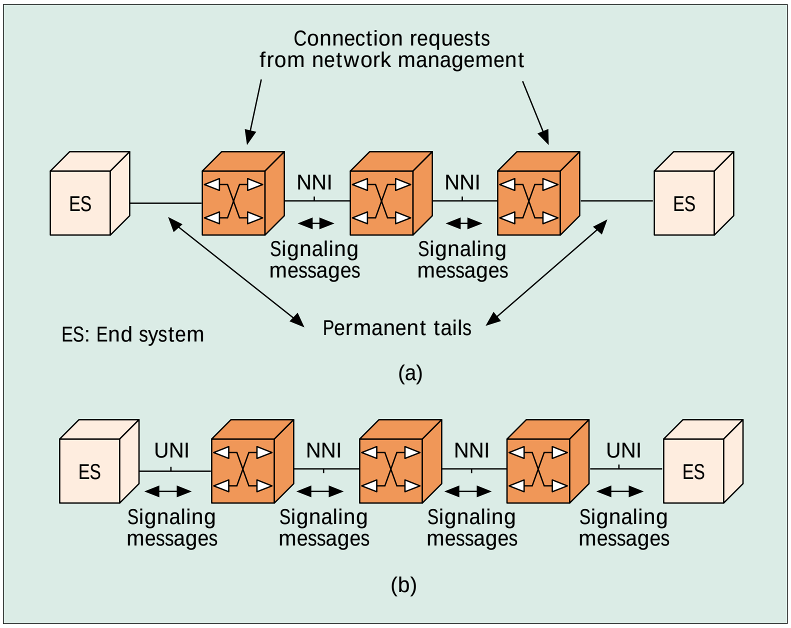

**Figure 2.** *a) Soft permanent connections; b) switched connections.*

mechanisms are not dependent on connection control in the control or management plane (in case of response to a failure). This is achieved by means of specialized protection agents and protocols that communicate between network elements over an open interface. These protocols are generally bit-level and implemented in firmware (allowing fast processing, but not very amenable to extensions) rather than message-based.

Full diversity for protected connections requires knowledge of all the underlying layer networks used to support the connections. This is because the working and standby connections that appear disjoint in one layer network may be commonly routed at a lower layer, and therefore subject to failures affecting both working and standby circuits. This is shown in Fig. 3. Indeed, server layers may contain some features such as passive splitters that cannot be detected by topology discovery protocols. In order to ensure that the control plane calculates diverse routes, it is necessary to have access to all the information regarding the server layers all the way down to the fiber, cable, and duct. It is also necessary to know the relationship of these layers to building locations. At the infrastructure level this topology has to be manually entered, and is therefore prone to error and particularly difficult to obtain where operators lease capacity from third parties.

As an alternative to holding this information in every switch, many existing operational support systems maintain large network databases of these lower layers. These can be (and already are) used to precalculate the working and standby paths in the management domain before presenting them to the control plane. This approach minimizes the impact on existing operational processes and systems used, for example, in network planning and fault correlation.

The network management system is notified in response to a change in protection status, such as a failure on the standby connection. It is also to be noted that in many cases the majority of protection switching events are initiated not as a result of network failures, but rather by the network management systems in response to planned engineering work. Restoration mechanisms can be used for unprotected connections. In contrast to network protection, which requires dedicated and preallocated resources, restoration seeks alternative resources from the spare capacity in the network. This requires network-level processes that have visibility of the network topology and are responsive to changes in topology. Restoration is therefore ideally suited to the control plane since it is much faster and more scalable than centralized management approaches that require time to collect alarms, correlate with network connectivity data to determine affected services, and search for new routes. For network restoration it is necessary to have visibility of connectivity within the layer network associated with the connection, but not always to have visibility of the fiber, cable, and duct.

From the discussion above it is clear that soft permanent connections allow much of the functionality of existing operational support systems to be reused while transferring much of the connection management to a more suitable environment. They are simpler to introduce into large networks than switched connections (as discussed later), and as such may see earlier large-scale deployment.

## SWITCHED NETWORK CONNECTIONS

Soft permanent connections offer the potential to streamline operational processes and do not impact significantly on business models and services, in contrast to switched services. Traditionally transport networks have been designed to efficiently use bandwidth. Switched networks, however, require a pool of spare capacity and preprovisioning (with the cost this incurs) to ensure that an acceptable level of blocking can be maintained during the busy hour. To ensure that switches and links between them are correctly dimensioned, it is necessary to understand the traffic patterns of the network. However, the “calling patterns” and holding times of a 2 Mb/s or 10 Gb/s “phone” network are, as of now, unknown, making it difficult to predict the size of network required to meet a certain grade of service. This is made all the more challenging because, in contrast to the telephone network, the number of circuits in the transport network is much lower, and there is little or no statistical gain in the access network, so the benefits of scale are less likely to be realized in the optical network. The issue of dimensioning is likely to become an area of active research.

Switched networks are, for the above reasons, potentially more attractive in practice for layer networks with high volumes of connections and churn, which therefore put more strain on management systems. These layers will tend to be at the higher layers of the network such as VC-12 and VC-4. As we descend through the layers of the transport network, the time constants of the “connections” increase dramatically and churn rates decrease (the duct being the slowest).

In the previous section it was noted that both protection and restoration could be employed. While there has been considerable discussion on the need for on-demand circuits that are protected end-to-end, the authors question whether this is really necessary. Mission-critical applications such as air traffic control will normally be provisioned between dedicated sites with dedicated protection, rather than switched. Rather than complicate the function of the control plane, the following strategies can be employed:

- Soft permanent connections are used to provide protected circuits.
- Unprotected soft permanent connections receive first call on restoration capacity.
- Switched services use whatever restoration capacity remains.

Any switched calls that are dropped (which may occur if restoration takes too long) can attempt to reconnect. Such a strategy is consistent with avoiding the need to include fiber, cable, and duct information in every switch, and thus may allow for a more scalable control plane.

## INTELLIGENT NETWORKING

Network operators may introduce some forms of intelligent networking features to differentiate their products, such as supplementary services, including, for example:

- Closed user group
- Call gapping
- Carrier preselect
- Automatic callback
- Call logging
- Attendant

The ability to offer such services will almost completely remove the distinction between transport and switching. The ability to change the topology of the transport network rapidly will bring a new dimension to traffic engineering. It may also allow the network to be used in new ways, such as selling off-peak redundant capacity at discounted rates to encourage usage of the network during periods of low demand. Users may even perhaps initiate real-time negotiations with bandwidth brokers to obtain the lowest cost connections — a real-time spot market for bandwidth.

## CONCLUSIONS

This article examines some of the major features required of transport network control planes to support global connectivity. Central to the capability to route, find, and connect is the development of a naming and addressing scheme that will scale to meet network growth over the next few decades. Control planes will allow faster service provisioning and create opportunities for new innovative network services. However, such systems will still require comprehensive management systems, and the successful operators and vendors of the future will be those capable of developing operational support systems that complement the control plane with service management capabilities, automated plan and build processes, inventory management, and capacity planning.

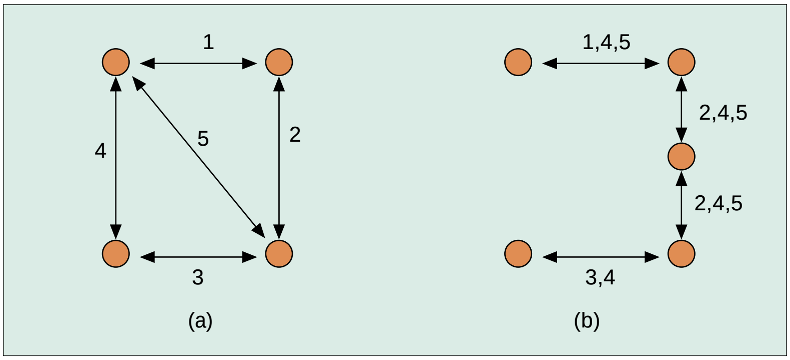

**Figure 3.** *The relationship between client and server layers: a) logical connectivity in the client layer; b) logical connectivity in the server layer. Connections that appear diverse in the client layer share common infrastructure in the server layer.*

## ACKNOWLEDGMENTS

The authors would like to thank colleagues at BT for many valuable discussions, especially Neil Harrison.

## REFERENCES

[1] ITU-T Delayed Contribution COM15-D558 "Dial-Up Connections in the OTN," submitted by BT, Geneva, Switzerland, June 21–July 2, 1999.
[2] M. Sexton and A. Reid, *Broadband Networking: ATM, SDH and SONET*, Artech House, 1997.
[3] ITU-T Rec. G.831, "Management Capabilities of Transport Networks Based on the Synchronous Digital Hierarchy (SDH)," 2000.
[4] ITU-T Rec. G.807, "Requirements for the Automatic Switched Transport Network (ASTN)," Caracas, May 2001.
[5] ITU-T Draft Rec. G.ason, "Architecture for the Automatic Switched Optical Network (ASON)," Turin, Italy, 2001.
[6] ITU-T Rec. G.805, "Generic Functional Architecture of Transport Networks," 2001.

## BIOGRAPHIES

ALAN MCGUIRE (alan.mcguire@bt.com) graduated from the University of St. Andrews with a first class honors degree in physics and an M.Sc. in medical physics from the University of Aberdeen in 1989. He joined BT from university and has been involved in a wide variety of technical areas including optical networking, SDH, ATM, network management, functional architecture, and network design. He is currently the principal engineer for core transport in BTexact Technologies. He is a member of the Institute of Physics and is a Chartered Physicist.

SHEHZAD MIRZA received his B.Sc. honors degree in physics in 1995 from the University of Edinburgh, and joined BT Labs in 1996 after graduating with an M.Sc. in optical electronics from Strathclyde University, Scotland. He has previously worked on optical component design, optical amplifier modeling, data communications network design, WDM deployment, and the launch of new data services. He is currently in the Broadband Architectures & Optical Networks group in BT working on optical network design, network management, and the optical control plane.

DARREN FREELAND graduated from the University of Strathclyde in Glasgow with a first class honors degree in electronic and electrical engineering. He joined BT in the summer of 1998, working on various projects related to core transport networks. His current main interests lie with standardization issues for the optical transport network, specifically network architecture and management/control.