Optical Network Control – Concepts, Standardization & Interoperability

Greg Bernstein
Grotto Networking, Fremont CA, USA, gregb@grotto-networking.com

Abstract *This paper provides a short introduction to the key components of the optical control plane: discovery, automated connection management and topology/resource status dissemination. The state of standardization and interoperability of these functions are also reviewed.*

Introduction

Modern optical networks based on SDH and OTN concepts provide a host of important fault management and performance monitoring functions [1]. These typically include: continuity supervision (detection of LOS), connectivity supervision (via trace messages), fault information (AIS, RDI), quality monitoring (forward and backward), and protection control (e.g., via K byte signalling). As optical networks scale in terms of the number of lambdas, fibres, nodes and connections, new functionality is needed to assist in the provisioning process. Hence the "new" optical control plane technology does not replace management plane technology but rather supplements it to provide scaling, better interoperability, new restoration options and new services.

In this paper we will describe three fundamental areas of the optical network control plane [1]: end point discovery, automated connection control and topology/resource status dissemination. We explain the fundamental concepts, the similarities and differences with other technologies such as IP routing and telephony signalling, and the state of standards and interoperability.

Automated Discovery

One of the most fundamental questions one end of a communications link can ask concerning the other end is: who are you? In the case of a single fibre this question seems trivial however in the case of many fibres figuring this out "manually" with the aid of trace overhead or the like can be tedious and error prone. In IP intra-domain routing protocols such as OSPF [2] there is a "hello" sub-protocol that specifically answers this question automatically.

How this works in the optical world is shown in Figure 1, where each link end uses some predefined overhead to announce its identity (e.g., like an IP address) and port number. The first major difference between our optical situation and the IP situation is that in the IP case both data plane and control communications occur over the exact same "channel". In the optical case this is not necessarily true, and hence we need to make sure that whatever mechanism used for optical discovery is applicable to the data plane.

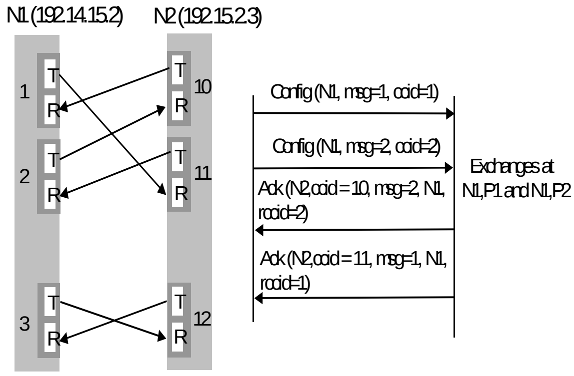

Figure 1. *Automated Discovery via OIF UNI 1.0.*

The second major complication for endpoint discovery in optical networks is that, unlike IP networks, optical networks are layered and not flat. Hence we have the additional question of "who's my neighbour at layer X" as show in Figure 2. For example in an SDH layered network we can speak of neighbour endpoints at the regenerator section layer, the multiplex section layer, the higher order VC layer, and the lower order VC layer.

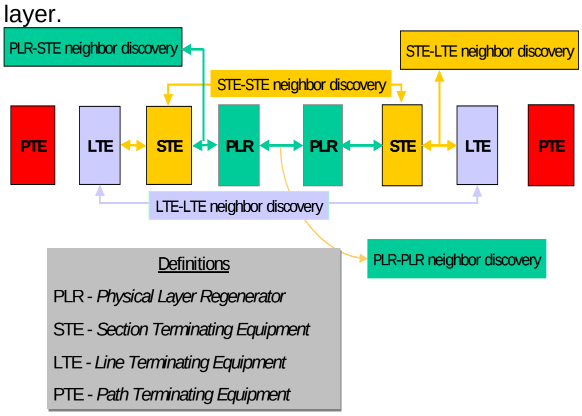

Figure 2. *Discovery of neighbors at different layers.*

The concepts of multi-layer optical discovery are given and properly formalized in ITU-T G.7714 [3]. Specific techniques applicable to SDH are given in G.7714.1. Since G.7714.1 is relatively new vendor interoperability is relatively non-existent. The OIF's UNI 1.0 included a discovery mechanism based on modified LMP [4] messages running over SDH RS or MS DCC however this hasn't seen widespread deployment. It should also be noted that although the OIF based their UNI discovery on modified LMP messages LMP doesn't provide for discovery instead it provides SDH like management functions for non-SDH/OTN links.

It should be pointed out that most vendors with relatively modern equipment offer some type of automated discovery at, at least, one layer of the optical network. This is typically done by running a modified intra-domain IP routing protocol over either the RS or MS DCC channels. There are a couple problems with such an approach as a general solution for discovery. First, this method of discovery only works at one layer. Second, there may be endpoints that we wish to discover but not establish a "routing adjacency". For example end systems typically do not participate in routing protocols, nor would we necessarily want to use this type of routing protocol between administrative domains.

Automated Connection Control

Given a request to establish a connection between two points, the A-end and the Z-end of the connection, a path through the network is computed for the connection taking into account bandwidth, reliability, survivability and possibly even diversity constraints. Given this path, cross connects on the individual switches are set up to establish the connection.

There are two ways that are currently in wide spread use in telecommunications systems to set up the cross connects that form an end to end connection. The first has a centralized management system talking to each of the switches involved in the connection individually and commanding them to create the necessary cross connects as shown in Figure 3. The second has a management system or end system talking to the switch at the A-end of the connection and telling that switch the details concerning the connection desired. This switch then talks to the next switch along the path, until the connection request has propagated all the way to the Z-end of the connection as shown in Figure 4.

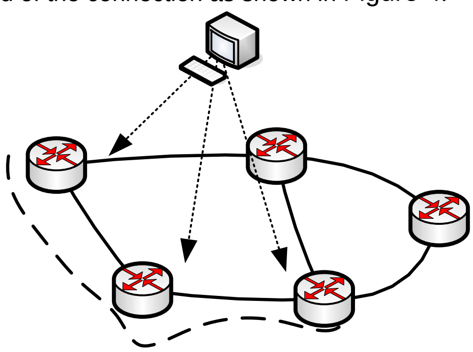

Figure 3. *Management based connection control.*

The case where an end system makes the request to the first switch is generally known as the User Network Interface (UNI) model. The case where a management system makes the request to the first switch is generally know as the Soft Permanent Connection (SPC) model.

The advantage of the management system approach is its simplicity particularly for networks with a small number of nodes and connections. As the number of nodes or number of connections needing to be established or torn down increases a centralized management system becomes a performance bottleneck. This is why the worldwide telephone system is based on signalling.

The disadvantage of the signalled approach comes with added complexity. Each switch must implement signalling protocol and call processing functionality. The signalling protocol is the "standardized" language use to communicate the connection request from node to node. Call processing, not subject to standardization, is the functionality on the node that acts on the signalling requests.

Although signalling/call processing has been used for ages in the telephone system the optical network adds a few special requirements of its own. First and foremost is an emphasis on reliability. In optical networks a single connection can be carrying significantly more traffic than its 64kbps telephony brethren. Hence the usual, though rarely used, technique of "call clearing" in the case of software faults or upgrades is completely unacceptable in the optical case. Second, optical bandwidth is relatively expensive and hence stranding bandwidth by not cleaning up all the cross connects on inactive connections in a timely manner is also unacceptable.

Although the optical signalling standards have been around a bit longer than the discovery standards, e.g., OIF UNI 1.0, ITU-T G.7713(.1-.3) [5] and IETF GMPLS RSVP-TE and CR-LDP. Interoperability has been hampered by having three different signalling protocols to choose from and a number of subtle but important issues still remaining to be worked out between the standard organizations defining these protocols.

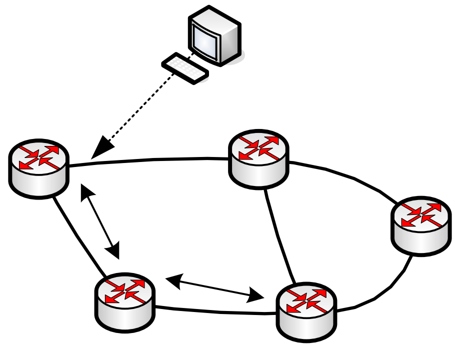

Figure 4. *Signaling based connection control.*

Topology and Resource Status Dissemination

For a switch to be able to respond to a connection request it first must be able to find a path between the A-end and Z-end of the connection satisfying the connections bandwidth and other constraints. Hence it needs a map of the network, like that shown in Figure 4, that gives the topology of the network as well as the status of link resources. Since any switch could be the A-end of a connection all switches need such a map if they are to act as the first switch in a signalled connection.

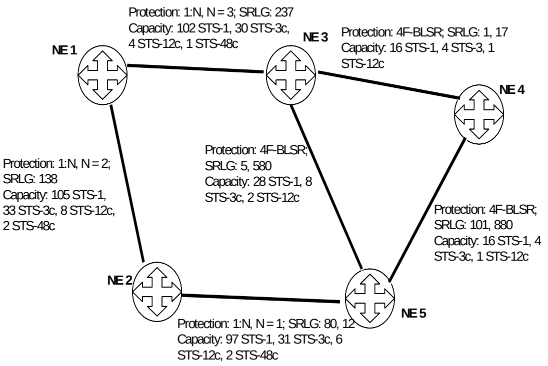

Figure 5. *Example network map.*

It turns out that, as a by product of their operation, link state intra-domain IP route protocols such as IS-IS and OSPF [2] disseminate network topology to all the nodes participating in the protocol.

Now for an IP routing protocol the end result is an IP forwarding table, i.e., a table with aggregated address information that tells the IP router which link to forward an IP packet with a given destination address. An intra-domain IP link state routing protocol has the following components: discovery, topology dissemination, path computation, and creation of the data plane forwarding table. The similarities and differences between IP routing and optical switching as follows:

Discovery:

- IP link state protocol include “hello” protocol to determine neighbours (flat network, control and data planes communication channels the same).
- Optical: separate discovery procedure for each layer is needed and care needs to be taken due to differences between data plane and control plane communications channels.

Topology Dissemination:

- Both can use the same efficient and reliable mechanisms for information dissemination.

Path Computation:

- IP: path computation is standardized and must be computed the same way on all IP routers.
- Optical: a path is computed for each connection at a time. Doesn’t have to be computed the same way on all switches (note that only the lead switch performs this calculation), each connection can be computed according to different criteria. No standardization is needed.

Data plane forwarding:

- IP: Routing protocol sets up forwarding table
- Optical: Signalling (or management) protocol sets up cross connect components of the connection.

From the previous items we see that an optical “routing” protocol really is concerned with disseminating topology information while an IP routing protocol has many other duties including a direct impact on data plane forwarding. Hence in the IP case we see that the routing protocol is “service impacting” while in the optical case the signalling protocol carries this burden.

Another consequence of these differences comes in the amount of information about a link that is useful in setting up the data plane forwarding. In the IP case due to the standardized path computation and the fact that we do not compute paths for individual connection the only additional information we have about a link is the “link weight”, a fixed quantity assigned to a link that influences the path chosen to a destination by a routing protocol. In the optical case since we compute the path for each connection individually we can make use of a good deal more information concerning each link to meet the requirements of a particular connection. This includes information concerning: capacity (bandwidth), underlying protection, and which links may share a similar fate with respect to failures.

The bandwidth accounting needed in optical circuit-switched networks is different than in packet networks. In packet networks using either ATM QoS or MPLS-TE, statistical measures are used to characterize the load on a link, often with varying degrees of success. The inexactness of such measures and the “compressibility” of statistically multiplexed traffic imply that a small percentage change in link utilization can usually be absorbed by the network. By contrast, if an OC-192 link has just one STS-1 path occupied (less than 1% of the link bandwidth) it cannot accommodate an STS-192c path. Due to the relatively simple finite multiplex structures currently used in optical networks tracking bandwidth resources is much easier than packet switched networks, however much stricter bandwidth accounting is required on circuit switched links. In particular, for an individual optical circuit switched link it is expected that this link can be fully utilized. While due to queuing effects a packet switched link on average can never be run at full capacity and is typically run at less than 70% of capacity. Understanding which links may share a similar fate with respect to failure scenarios, such as cable cuts, is important when diversely routing connections.

One thing not mentioned in the section on discovery is to what use can be made of information concerning lower layer network neighbours. For example consider a large SDH cross connect connected to several different WDM multiplexers. From discovering its optical neighbours the cross connect can establish which of its SDH link will be carried over the same fibre and hence subject to similar risks. This information is useful link information and is codified in the shared risk link group (SRLG) attribute of a link [1]. Note that information such as conduit and right of way dependencies can also be represented in the SRLG information (though this information cannot be discovered).

As was done with ATM's PNNI routing and MPLS TE routing, the extension of IP routing protocols to the optical situation involves adding the previously mentioned extra information concerning the links to the protocols. This work has been moving forward at both the IETF [6] and ITU-T [7]. Currently, however, only pre-standard non-interoperable implementations exist. However due to the usefulness of being able to get a view of the entire network by communicating with a single switching node, the efficiency of the link state protocols in terms overhead bandwidth consumed and the non-service impacting nature of these protocols in the optical situation, even non-standard protocols have proved popular with the carriers that have tried them.

New Service and Restoration Options

The ability to rapidly and efficiently establish and teardown optical connections, leads to the viability of short duration services that can be economically feasible for both carrier and customer alike. Another type of service enabled by the rapid provisioning enabled by the control plan is time of day based bandwidth allocation. This may be particular valuable when used with data services taking advantage of SDH's new virtual concatenation (VCAT) and link capacity adjustment scheme (LCAS) features. These features allow the "hitless" addition and removal of bandwidth, in say 50Mbps or 150Mbps chunks, from data links [8].

There are three fundamental qualities that we can use to evaluate any restoration scheme: robustness, bandwidth efficiency and recovery time. Let's evaluate these criteria for the simplest type of restoration that we can implement given both automated connection control and topology dissemination/resource status control plane features. We'll call this source re-route protection and in the case of a failure in the network the affected connections are released, a new path is computed around the failure from the up to date topology and resource information, and a replacement connection is established via signalling.

From the point of robustness source re-route can deal with more types of network failures than either standard linear or ring protection. From the point of bandwidth efficiency source re-route works at the finest granularity possible, the connection level. For an SDH network this compares with MS section based linear or ring protection which works at the MS level rather than the VC level. In sufficiently "meshy" networks additional bandwidth savings are possible [1]. The compromise for these increases in robustness and bandwidth efficiency is longer times needed for restoration. However, not all signals need to be restored within the classic SDH 50 mili seconds, but it can be quite important for them be restored within 1-2 seconds. Hence even without optimizations for speed improvement the optical control plane can offer a very useful restoration option.

Conclusion

The "new" optical control plane technologies are a necessary supplement to existing management technologies as optical networks scale up in terms of the number of fibres, lambdas, switching nodes or connections supported. In addition to providing for more efficient operation of larger optical networks, optical control plane technology can also allow for new services and protection options. Although interoperability across vendors is not widespread yet the necessary standards have essentially been developed and continue to be refined at the IETF, ITU-T and the OIF. For SDH and OTN optical technologies the optical control plane can be understood in terms of the processes of discovery, automated connection control and topology and resource status dissemination as reviewed here.

References

1. G. Bernstein, et al, *Optical Network Control,* Addison-Wesley, 2004.
2. J. Moy, *OSPF Version 2*, IETF RFC 2328, 1998.
3. ITU-T recommendation G.7714, *Generalized automatic discovery techniques*, 2001.
4. J. Lange, ed., Link Management Protocol (LMP), IETF work in progress, 2003.
5. ITU-T recommendation G7713, *Distributed Call and Connection Management,* 2001.
6. K. Kompella, Ed., *Routing Extensions in Support of Generalized Multi-Protocol Label Switching*, IETF work in progress, 2003.
7. ITU-T Recommendation G.7715, *Architecture and Requirements for Routing in the Automatic Switched Optical Networks*, June 2002.
8. Paul Bonenfant, et al, "Generic Framing Procedure (GFP): The catalyst for efficient data over transport", IEEE Communications Magazine, May 2002.