        Impairments and Other Constraints on Optical Layer Routing

Status of This Memo

Copyright Notice

Abstract

   Optical networking poses a number challenges for Generalized Multi-
   Protocol Label Switching (GMPLS). Fundamentally, optical technology
   is an analog rather than digital technology whereby the optical layer
   is lowest in the transport hierarchy and hence has an intimate
   relationship with the physical geography of the network. This
   contribution surveys some of the aspects of optical networks that
   impact routing and identifies possible GMPLS responses for each: (1)
   Constraints arising from the design of new software controllable
   network elements, (2) Constraints in a single all-optical domain
   without wavelength conversion, (3) Complications arising in more
   complex networks incorporating both all-optical and opaque
   architectures, and (4) Impacts of diversity constraints.

Table of Contents

1. Introduction

   Generalized Multi-Protocol Label Switching (GMPLS) [Mannie04] aims to
   extend MPLS to encompass a number of transport architectures,
   including optical networks that incorporate a number of all-optical
   and opto-electronic elements, such as optical cross-connects with
   both optical and electrical fabrics, transponders, and optical add-
   drop multiplexers. Optical networking poses a number of challenges
   for GMPLS. Fundamentally, optical technology is an analog rather
   than digital technology whereby the optical layer is lowest in the
   transport hierarchy and hence has an intimate relationship with the
   physical geography of the network.

   GMPLS already has incorporated extensions to deal with some of the
   unique aspects of the optical layer. This contribution surveys some
   of the aspects of optical networks that impact routing and identifies
   possible GMPLS responses for each. Routing constraints and/or
   complications arising from the design of network elements, the
   accumulation of signal impairments, and the need to guarantee the
   physical diversity of some circuits are discussed.

   Since the purpose of this document is to further the specification of
   GMPLS, alternative approaches to controlling an optical network are
   not discussed. For discussions of some broader issues, see
   [Gerstel2000] and [Strand02].

   The organization of the contribution is as follows:

   - Section 2 is a section requested by the sub-IP Area management for
      all new documents. It explains how this document fits into the
      Area and into the IPO WG, and why it is appropriate for these
      groups.

- Section 3 describes constraints arising from the design of new
  software controllable network elements.

- Section 4 addresses the constraints in a single all-optical domain
  without wavelength conversion.

- Section 5 extends the discussion to more complex networks and
  incorporates both all-optical and opaque architectures.

- Section 6 discusses the impacts of diversity constraints.

- Section 7 deals with security requirements.

- Section 8 contains acknowledgments.

- Section 9 contains references.

- Section 10 contains contributing authors' addresses.

2. Sub-IP Area Summary and Justification of Work

This document merges and extends two previous expired Internet-Drafts
that were made IPO working group documents to form a basis for a
design team at the Minneapolis IETF meeting, where it was also
requested that they be merged to create a requirements document for
the WG.

In the larger sub-IP Area structure, this merged document describes
specific characteristics of optical technology and the requirements
they place on routing and path selection. It is appropriate for the
IPO working group because the material is specific to optical
networks. It identifies and documents the characteristics of the
optical transport network that are important for selecting paths for
optical channels, which is a work area for the IPO WG. The material
covered is directly aimed at establishing a framework and
requirements for routing in an optical network.

3. Reconfigurable Network Elements

3.1. Technology Background

Control plane architectural discussions (e.g., [Awduche99]) usually
assume that the only software reconfigurable network element is an
optical layer cross-connect (OLXC). There are however other software
reconfigurable elements on the horizon, specifically tunable lasers
and receivers and reconfigurable optical add-drop multiplexers

(OADM). These elements are illustrated in the following simple
example, which is modeled on announced Optical Transport System (OTS)
products:

```
                    +                               +
---+---+      |\                            /|    +---+---
---| A |----|D|        X             Y     |D|----| A |---
---+---+      |W|    +--------+    +--------+    |W|    +---+---
    :         |D|-----|  OADM  |-----|  OADM  |-----|D|      :
---+---+      |M|    +--------+    +--------+    |M|    +---+---
---| A |----| |     |       |     |       |     | |----| A |---
---+---+      |/      |       |     |       |     \|    +---+---
                +      +---+ +---+   +---+ +---+    +
             D     | A | | A |   | A | | A |    E
                     +---+ +---+   +---+ +---+
                      | |   | |     | |   | |
```
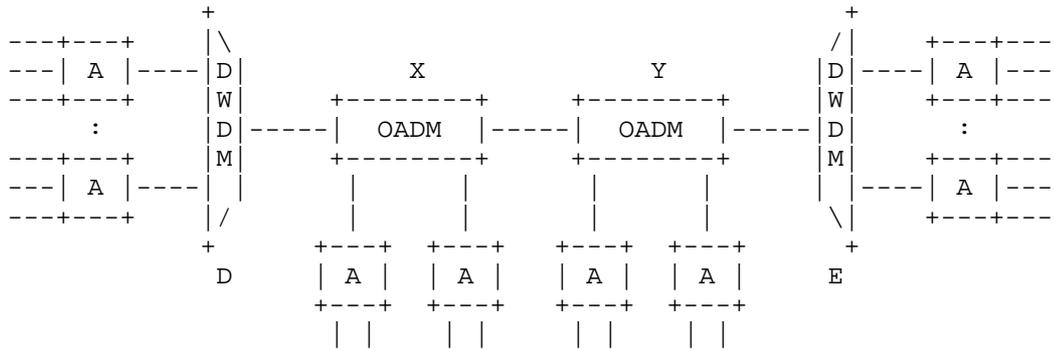
        Figure 3-1: An OTS With OADMs - Functional Architecture

In Fig. 3-1, the part that is on the inner side of all boxes labeled
"A" defines an all-optical subnetwork. From a routing perspective
two aspects are critical:

- Adaptation: These are the functions done at the edges of the
   subnetwork that transform the incoming optical channel into the
   physical wavelength to be transported through the subnetwork.

- Connectivity: This defines which pairs of edge Adaptation
   functions can be interconnected through the subnetwork.

In Fig. 3-1, D and E are DWDMs and X and Y are OADMs. The boxes
labeled "A" are adaptation functions. They map one or more input
optical channels assumed to be standard short reach signals into a
long reach (LR) wavelength or wavelength group that will pass
transparently to a distant adaptation function. Adaptation
functionality that affects routing includes:

- Multiplexing: Either electrical or optical TDM may be used to
   combine the input channels into a single wavelength. This is done
   to increase effective capacity: A typical DWDM might be able to
   handle 100 2.5 Gb/sec signals (250 Gb/sec total) or 50 10 Gb/sec
   (500 Gb/sec total); combining the 2.5 Gb/sec signals together thus
   effectively doubles capacity. After multiplexing the combined
   signal must be routed as a group to the distant adaptation
   function.

- Adaptation Grouping: In this technique, groups of k (e.g., 4)
  wavelengths are managed as a group within the system and must be
  added/dropped as a group. We will call such a group an
  "adaptation grouping". Examples include so called "wave group"
  and "waveband" [Passmore01]. Groupings on the same system may
  differ in basics such as wavelength spacing, which constrain the
  type of channels that can be accommodated.

- Laser Tunability: The lasers producing the LR wavelengths may have
  a fixed frequency, may be tunable over a limited range, or may be
  tunable over the entire range of wavelengths supported by the
  DWDM. Tunability speeds may also vary.

Connectivity between adaptation functions may also be limited:

- As pointed out above, TDM multiplexing and/or adaptation grouping
  by the adaptation function forces groups of input channels to be
  delivered together to the same distant adaptation function.

- Only adaptation functions whose lasers/receivers are tunable to
  compatible frequencies can be connected.

- The switching capability of the OADMs may also be constrained.

For example:

o There may be some wavelengths that can not be dropped at all.

o There may be a fixed relationship between the frequency dropped
   and the physical port on the OADM to which it is dropped.

o OADM physical design may put an upper bound on the number of
   adaptation groupings dropped at any single OADM.

For a fixed configuration of the OADMs and adaptation functions
connectivity will be fixed: Each input port will essentially be
hard-wired to some specific distant port. However this connectivity
can be changed by changing the configurations of the OADMs and
adaptation functions. For example, an additional adaptation grouping
might be dropped at an OADM or a tunable laser retuned. In each case
the port-to-port connectivity is changed.

These capabilities can be expected to be under software control.
Today the control would rest in the vendor-supplied Element
Management system (EMS), which in turn would be controlled by the
operator's OSes. However in principle the EMS could participate in
the GMPLS routing process.

3.2. Implications for Routing

   An OTS of the sort discussed in Sec. 3.1 is essentially a
   geographically distributed but blocking cross-connect system. The
   specific port connectivity is dependent on the vendor design and also
   on exactly what line cards have been deployed.

   One way for GMPLS to deal with this architecture would be to view the
   port connectivity as externally determined. In this case the links
   known to GMPLS would be groups of identically routed wavebands. If
   these were reconfigured by the external EMS the resulting
   connectivity changes would need to be detected and advertised within
   GMPLS. If the topology shown in Fig. 3-1 became a tree or a mesh
   instead of the linear topology shown, the connectivity changes could
   result in Shared Risk Link Group (SRLG - see Section 6.2) changes.

   Alternatively, GMPLS could attempt to directly control this port
   connectivity. The state information needed to do this is likely to
   be voluminous and vendor specific.

4. Wavelength Routed All-Optical Networks

   The optical networks deployed until recently may be called "opaque"
   ([Tkach98]): each link is optically isolated by transponders doing
   O/E/O conversions. They provide regeneration with retiming and
   reshaping, also called 3R, which eliminates transparency to bit rates
   and frame format. These transponders are quite expensive and their
   lack of transparency also constrains the rapid introduction of new
   services. Thus there are strong motivators to introduce "domains of
   transparency" - all-optical subnetworks - larger than an OTS.

   The routing of lightpaths through an all-optical network has received
   extensive attention. (See [Yates99] or [Ramaswami98]). When
   discussing routing in an all-optical network it is usually assumed
   that all routes have adequate signal quality. This may be ensured by
   limiting all-optical networks to subnetworks of limited geographic
   size that are optically isolated from other parts of the optical
   layer by transponders. This approach is very practical and has been
   applied to date, e.g., when determining the maximum length of an
   Optical Transport System (OTS). Furthermore operational
   considerations like fault isolation also make limiting the size of
   domains of transparency attractive.

   There are however reasons to consider contained domains of
   transparency in which not all routes have adequate signal quality.
   From a demand perspective, maximum bit rates have rapidly increased
   from DS3 to OC-192 and soon OC-768 (40 Gb/sec). As bit rates
   increase it is necessary to increase power. This makes impairments

and nonlinearities more troublesome. From a supply perspective,
optical technology is advancing very rapidly, making ever-larger
domains possible. In this section, we assume that these
considerations will lead to the deployment of a domain of
transparency that is too large to ensure that all potential routes
have adequate signal quality for all circuits. Our goal is to
understand the impacts of the various types of impairments in this
environment.

Note that, as we describe later in the section, there are many types
of physical impairments. Which of these needs to be dealt with
explicitly when performing on-line distributed routing will vary
considerably and will depend on many variables, including:

- Equipment vendor design choices,
- Fiber characteristics,
- Service characteristics (e.g., circuit speeds),
- Network size,
- Network operator engineering and deployment strategies.

For example, a metropolitan network that does not intend to support
bit rates above 2.5 Gb/sec may not be constrained by any of these
impairments, while a continental or international network that wished
to minimize O/E/O regeneration investment and support 40 Gb/sec
connections might have to explicitly consider many of them. Also, a
network operator may reduce or even eliminate their constraint set by
building a relatively small domain of transparency to ensure that all
the paths are feasible, or by using some proprietary tools based on
rules from the OTS vendor to pre-qualify paths between node pairs and
put them in a table that can be accessed each time a routing decision
has to be made through that domain.

4.1. Problem Formulation

   We consider a single domain of transparency without wavelength
   translation. Additionally, due to the proprietary nature of DWDM
   transmission technology, we assume that the domain is either single
   vendor or architected using a single coherent design, particularly
   with regard to the management of impairments.

   We wish to route a unidirectional circuit from ingress client node X
   to egress client node Y. At both X and Y, the circuit goes through
   an O/E/O conversion that optically isolates the portion within our
   domain. We assume that we know the bit rate of the circuit. Also,
   we assume that the adaptation function at X may apply some Forward
   Error Correction (FEC) method to the circuit. We also assume we know
   the launch power of the laser at X.

   Impairments can be classified into two categories, linear and
   nonlinear. (See [Tkach98] or [Kaminow02] for more on impairment
   constraints.) Linear effects are independent of signal power and
   affect wavelengths individually. Amplifier spontaneous emission
   (ASE), polarization mode dispersion (PMD), and chromatic dispersion
   are examples. Nonlinearities are significantly more complex: they
   generate not only impairments on each channel, but also crosstalk
   between channels.

   In the remainder of this section we first outline how two key linear
   impairments (PMD and ASE) might be handled by a set of analytical
   formulae as additional constraints on routing. We next discuss how
   the remaining constraints might be approached. Finally we take a
   broader perspective and discuss the implications of such constraints
   on control plane architecture and also on broader constrained domain
   of transparency architecture issues.

4.2. Polarization Mode Dispersion (PMD)

   For a transparent fiber segment, the general PMD requirement is that
   the time-average differential group delay (DGD) between two
   orthogonal state of polarizations should be less than some fraction a
   of the bit duration, T=1/B, where B is the bit rate. The value of
   the parameter a depends on three major factors: 1) margin allocated
   to PMD, e.g., 1dB; 2) targeted outage probability, e.g., 4×10^{-5}, and
   3) sensitivity of the receiver to DGD. A typical value for a is 10%
   [ITU]. More aggressive designs to compensate for PMD may allow
   values higher than 10%. (This would be a system parameter dependent
   on the system design. It would need to be known to the routing
   process.)

   The PMD parameter (Dpmd) is measured in pico-seconds (ps) per
   sqrt(km). The square of the PMD in a fiber span, denoted as span-
   PMD-square is then given by the product of Dpmd**2 and the span
   length. (A fiber span in a transparent network refers to a segment
   between two optical amplifiers.) If Dpmd is constant, this results
   in a upper bound on the maximum length of an M-fiber-span transparent
   segment, which is inversely proportional to the square of the product
   of bit rate and Dpmd (the detailed equation is omitted due to the
   format constraint - see [Strand01] for details).

   For older fibers with a typical PMD parameter of 0.5 picoseconds per
   square root of km, based on the constraint, the maximum length of the
   transparent segment should not exceed 400km and 25km for bit rates of
   10Gb/s and 40Gb/s, respectively. Due to recent advances in fiber
   technology, the PMD-limited distance has increased dramatically. For
   newer fibers with a PMD parameter of 0.1 picosecond per square root
   of km, the maximum length of the transparent segment (without PMD

compensation) is limited to 10000km and 625km for bit rates of 10Gb/s
and 40Gb/, respectively. Still lower values of PMD are attainable in
commercially available fiber today, and the PMD limit can be further
extended if a larger value of the parameter a (ratio of DGD to the
bit period) can be tolerated. In general, the PMD requirement is not
an issue for most types of fibers at 10Gb/s or lower bit rate. But
it will become an issue at bit rates of 40Gb/s and higher.

If the PMD parameter varies between spans, a slightly more
complicated equation results (see [Strand01]), but in any event the
only link dependent information needed by the routing algorithm is
the square of the link PMD, denoted as link-PMD-square. It is the
sum of the span-PMD-square of all spans on the link.

Note that when one has some viable PMD compensation devices and
deploy them ubiquitously on all routes with potential PMD issues in
the network, then the PMD constraint disappears from the routing
perspective.

4.3. Amplifier Spontaneous Emission

ASE degrades the optical signal to noise ratio (OSNR). An acceptable
optical SNR level (SNRmin), which depends on the bit rate,
transmitter-receiver technology (e.g., FEC), and margins allocated
for the impairments, needs to be maintained at the receiver. In
order to satisfy this requirement, vendors often provide some general
engineering rule in terms of maximum length of the transparent
segment and number of spans. For example, current transmission
systems are often limited to up to 6 spans each 80km long. For
larger transparent domains, more detailed OSNR computations will be
needed to determine whether the OSNR level through a domain of
transparency is acceptable. This would provide flexibility in
provisioning or restoring a lightpath through a transparent
subnetwork.

Assume that the average optical power launched at the transmitter is
P. The lightpath from the transmitter to the receiver goes through M
optical amplifiers, with each introducing some noise power. Unity
gain can be used at all amplifier sites to maintain constant signal
power at the input of each span to minimize noise power and
nonlinearity. A constraint on the maximum number of spans can be
obtained [Kaminow97] which is proportional to P and inversely
proportional to SNRmin, optical bandwidth B, amplifier gain G-1 and
spontaneous emission factor n of the optical amplifier, assuming all
spans have identical gain and noise figure. (Again, the detailed
equation is omitted due to the format constraint - see [Strand01] for
details.) Let's take a typical example. Assuming P=4dBm,
SNRmin=20dB with FEC, B=12.5GHz, n=2.5, G=25dB, based on the

constraint, the maximum number of spans is at most 10. However, if
FEC is not used and the requirement on SNRmin becomes 25dB, the
maximum number of spans drops down to 3.

For ASE the only link-dependent information needed by the routing
algorithm is the noise of the link, denoted as link-noise, which is
the sum of the noise of all spans on the link. Hence the constraint
on ASE becomes that the aggregate noise of the transparent segment
which is the sum of the link-noise of all links can not exceed
P/SNRmin.

4.4. Approximating the Effects of Some Other Impairment Constraints

There are a number of other impairment constraints that we believe
could be approximated with a domain-wide margin on the OSNR, plus in
some cases a constraint on the total number of networking elements
(OXC or OADM) along the path. Most impairments generated at OXCs or
OADMs, including polarization dependent loss, coherent crosstalk, and
effective passband width, could be dealt with using this approach.
In principle, impairments generated at the nodes can be bounded by
system engineering rules because the node elements can be designed
and specified in a uniform manner. This approach is not feasible
with PMD and noise because neither can be uniformly specified.
Instead, they depend on node spacing and the characteristics of the
installed fiber plant, neither of which are likely to be under the
system designer's control.

Examples of the constraints we propose to approximate with a domain-
wide margin are given in the remaining paragraphs in this section.
It should be kept in mind that as optical transport technology
evolves it may become necessary to include some of these impairments
explicitly in the routing process. Other impairments not mentioned
here at all may also become sufficiently important to require
incorporation either explicitly or via a domain-wide margin.

Other Polarization Dependent Impairments
   Other polarization-dependent effects besides PMD influence system
   performance. For example, many components have polarization-
   dependent loss (PDL) [Ramaswami98], which accumulates in a system
   with many components on the transmission path. The state of
   polarization fluctuates with time and its distribution is very
   important also. It is generally required that the total PDL on
   the path be maintained within some acceptable limit, potentially
   by using some compensation technology for relatively long
   transmission systems, plus a small built-in margin in OSNR. Since
   the total PDL increases with the number of components in the data
   path, it must be taken into account by the system vendor when
   determining the maximum allowable number of spans.

Chromatic Dispersion
   In general this impairment can be adequately (but not optimally)
   compensated for on a per-link basis, and/or at system initial
   setup time. Today most deployed compensation devices are based on
   Dispersion Compensation Fiber (DCF). DCF provides per fiber
   compensation by means of a spool of fiber with a CD coefficient
   opposite to the fiber. Due to the imperfect matching between the
   CD slope of the fiber and the DCF some lambdas can be over
   compensated while others can be under compensated. Moreover DCF
   modules may only be available in fixed lengths of compensating
   fiber; this means that sometimes it is impossible to find a DCF
   module that exactly compensates the CD introduced by the fiber.
   These effects introduce what is known as residual CD. Residual CD
   varies with the frequency of the wavelength. Knowing the
   characteristics of both of the fiber and the DCF modules along the
   path, this can be calculated with a sufficient degree of
   precision. However this is a very challenging task. In fact the
   per-wavelength residual dispersion needs to be combined with other
   information in the system (e.g., types fibers to figure out the
   amount of nonlinearities) to obtain the net effect of CD either by
   simulation or by some analytical approximation. It appears that
   the routing/control plane should not be burdened by such a large
   set of information while it can be handled at the system design
   level. Therefore it will be assumed until proven otherwise that
   residual dispersion should not be reported. For high bit rates,
   dynamic dispersion compensation may be required at the receiver to
   clean up any residual dispersion.

Crosstalk
   Optical crosstalk refers to the effect of other signals on the
   desired signal. It includes both coherent (i.e., intrachannel)
   crosstalk and incoherent (i.e., interchannel) crosstalk. Main
   contributors of crosstalk are the OADM and OXC sites that use a
   DWDM multiplexer/demultiplexer (MUX/DEMUX) pair. For a relatively
   sparse network where the number of OADM/OXC nodes on a path is
   low, crosstalk can be treated with a low margin in OSNR without
   being a binding constraint. But for some relatively dense
   networks where crosstalk might become a binding constraint, one
   needs to propagate the per-link crosstalk information to make sure
   that the end-to-end path crosstalk which is the sum of the
   crosstalks on all the corresponding links to be within some limit,
   e.g., -25dB threshold with 1dB penalty ([Goldstein94]). Another
   way to treat it without having to propagate per-link crosstalk
   information is to have the system evaluate what the maximum number
   of OADM/OXC nodes that has a MUX/DEMUX pair for the worst route in
   the transparent domain for a low built-in margin. The latter one
   should work well where all the OXC/OADM nodes have similar level
   of crosstalk.

Effective Passband
    As more and more DWDM components are cascaded, the effective
    passband narrows. The number of filters along the link, their
    passband width and their shape will determine the end-to-end
    effective passband. In general, this is a system design issue,
    i.e., the system is designed with certain maximum bit rate using
    the proper modulation format and filter spacing. For linear
    systems, the filter effect can be turned into a constraint on the
    maximum number of narrow filters with the condition that filters
    in the systems are at least as wide as the one in the receiver.
    Because traffic at lower bit rates can tolerate a narrower
    passband, the maximum allowable number of narrow filters will
    increase as the bit rate decreases.

Nonlinear Impairments
    It seems unlikely that these can be dealt with explicitly in a
    routing algorithm because they lead to constraints that can couple
    routes together and lead to complex dependencies, e.g., on the
    order in which specific fiber types are traversed [Kaminow97].
    Note that different fiber types (standard single mode fiber,
    dispersion shifted fiber, dispersion compensated fiber, etc.) have
    very different effects from nonlinear impairments. A full
    treatment of the nonlinear constraints would likely require very
    detailed knowledge of the physical infrastructure, including
    measured dispersion values for each span, fiber core area and
    composition, as well as knowledge of subsystem details such as
    dispersion compensation technology. This information would need
    to be combined with knowledge of the current loading of optical
    signals on the links of interest to determine the level of
    nonlinear impairment. Alternatively, one could assume that
    nonlinear impairments are bounded and result in X dB margin in the
    required OSNR level for a given bit rate, where X for performance
    reasons would be limited to 1 or 2 dB, consequently setting a
    limit on the maximum number of spans. For the approach described
    here to be useful, it is desirable for this span length limit to
    be longer than that imposed by the constraints which can be
    treated explicitly. When designing a DWDM transport system, there
    are tradeoffs between signal power launched at the transmitter,
    span length, and nonlinear effects on BER that need to be
    considered jointly. Here, we assume that an X dB margin is
    obtained after the transport system has been designed with a fixed
    signal power and maximum span length for a given bit rate. Note
    that OTSs can be designed in very different ways, in linear,
    pseudo-linear, or nonlinear environments. The X-dB margin
    approach may be valid for some but not for others. However, it is
    likely that there is an advantage in designing systems that are

less aggressive with respect to nonlinearities, and therefore
somewhat sub-optimal, in exchange for improved scalability,
simplicity and flexibility in routing and control plane design.

4.5. Other Impairment Considerations

There are many other types of impairments that can degrade
performance. In this section, we briefly mention one other type of
impairment, which we propose be dealt with by either the system
designer or by the transmission engineers at the time the system is
installed. If dealt with successfully in this manner they should not
need to be considered in the dynamic routing process.

Gain Nonuniformity and Gain Transients For simple noise estimates to
be of use, the amplifiers must be gain-flattened and must have
automatic gain control (AGC). Furthermore, each link should have
dynamic gain equalization (DGE) to optimize power levels each time
wavelengths are added or dropped. Variable optical attenuators on
the output ports of an OXC or OADM can be used for this purpose, and
in-line devices are starting to become commercially available.
Optical channel monitors are also required to provide feedback to the
DGEs. AGC must be done rapidly if signal degradation after a
protection switch or link failure is to be avoided.

Note that the impairments considered here are treated more or less
independently. By considering them jointly and varying the tradeoffs
between the effects from different components may allow more routes
to be feasible. If that is desirable or the system is designed such
that certain impairments (e.g., nonlinearities) need to be considered
by a centralized process, then distributed routing is not the one to
use.

4.6. An Alternative Approach - Using Maximum Distance as the Only
      Constraint

Today, carriers often use maximum distance to engineer point-to-point
OTS systems given a fixed per-span length based on the OSNR
constraint for a given bit rate. They may desire to keep the same
engineering rule when they move to all-optical networks. Here, we
discuss the assumptions that need to be satisfied to keep this
approach viable and how to treat the network elements between two
adjacent links.

In order to use the maximum distance for a given bit rate to meet an
OSNR constraint as the only binding constraint, the operators need to
satisfy the following constraints in their all-optical networks:

   - All the other non-OSNR constraints described in the previous
      subsections are not binding factors as long as the maximum
      distance constraint is met.

   - Specifically for PMD, this means that the whole all-optical
      network is built on top of sufficiently low-PMD fiber such that
      the upper bound on the mean aggregate path DGD is always satisfied
      for any path that does not exceed the maximum distance, or PMD
      compensation devices might be used for routes with high-PMD
      fibers.

   - In terms of the ASE/OSNR constraint, in order to convert the ASE
      constraint into a distance constraint directly, the network needs
      to have a fixed fiber distance D for each span (so that ASE can be
      directly mapped by the gain of the amplifier which equals to the
      loss of the previous fiber span), e.g., 80km spacing which is
      commonly chosen by carriers. However, when spans have variable
      lengths, certain adjustment and compromise need to be made in
      order to avoid treating ASE explicitly as in section 4.3. These
      include: 1) Unless a certain mechanism is built in the OTS to take
      advantage of shorter spans, spans shorter than a typical span
      length D need to be treated as a span of length D instead of with
      its real length. 2) Spans that are longer than D would have a
      higher average span loss. In general, the maximum system reach
      decreases when the average span loss increases. Thus, in order to
      accommodate longer spans in the network, the maximum distance
      upper bound has to be set with respect to the average span loss of
      the worst path in the network. This sub-optimality may be
      acceptable for some networks if the variance is not too large, but
      may be too conservative for others.

   If these assumptions are satisfied, the second issue we need to
   address is how to treat a transparent network element (e.g., MEMS-
   based switch) between two adjacent links in terms of a distance
   constraint since it also introduces an insertion loss. If the
   network element cannot somehow compensate for this OSNR degradation,
   one approach is to convert each network element into an equivalent
   length of fiber based on its loss/ASE contribution. Hence, in
   general, introducing a set of transparent network elements would
   effectively result in reducing the overall actual transmission
   distance between the OEO edges.

   With this approach, the link-specific state information is link-
   distance, the length of a link. It equals the distance sum of all
   fiber spans on the link and the equivalent length of fiber for the
   network element(s) on the link. The constraint is that the sum of

all the link-distance over all links of a path should be less than
the maximum-path-distance, the upper bound of all paths.

4.7. Other Considerations

   Routing in an all-optical network without wavelength conversion
   raises several additional issues:

   - Since the route selected must have the chosen wavelength available
      on all links, this information needs to be considered in the
      routing process. One approach is to propagate information
      throughout the network about the state of every wavelength on
      every link in the network. However, the state required and the
      overhead involved in processing and maintaining this information
      is proportional to the total number of links (thus, number of
      nodes squared), maximum number of wavelengths (which keeps
      doubling every couple of years), and the frequency of wavelength
      availability changes, which can be very high. Instead
      [Hjalmtysson00], proposes an alternative method which probes along
      a chosen path to determine which wavelengths (if any) are
      available. This would require a significant addition to the
      routing logic normally used in OSPF. Others have proposed
      simultaneously probing along multiple paths.

   - Choosing a path first and then a wavelength along the path is
      known to give adequate results in simple topologies such as rings
      and trees ([Yates99]). This does not appear to be true in large
      mesh networks under realistic provisioning scenarios, however.
      Instead significantly better results are achieved if wavelength
      and route are chosen simultaneously ([Strand01b]). This approach
      would however also have a significant effect on OSPF.

4.8. Implications For Routing and Control Plane Design

   If distributed routing is desired, additional state information will
   be required by the routing to deal with the impairments described in
   Sections 4.2 - 4.4:

   - As mentioned earlier, an operator who wants to avoid having to
      provide impairment-related parameters to the control plane may
      elect not to deal with them at the routing level, instead treating
      them at the system design and planning level if that is a viable
      approach for their network. In this approach the operator can
      pre-qualify all or a set of feasible end-to-end optical paths
      through the domain of transparency for each bit rate. This
      approach may work well with relatively small and sparse networks,
      but it may not be scalable for large and dense networks where the
      number of feasible paths can be very large.

- If the optical paths are not pre-qualified, additional link-
  specific state information will be required by the routing
  algorithm for each type of impairment that has the potential of
  being limiting for some routes. Note that for one operator, PMD
  might be the only limiting constraint while for another, ASE might
  be the only one, or it could be both plus some other constraints
  considered in this document. Some networks might not be limited
  by any of these constraints.

- For an operator needing to deal explicitly with these constraints,
  the link-dependent information identified above for PMD is link-
  PMD-square which is the square of the total PMD on a link. For
  ASE the link-dependent information identified is link-noise which
  is the total noise on a link. Other link-dependent information
  includes link-span-length which is the total number of spans on a
  link, link-crosstalk or OADM-OXC-number which is the total
  crosstalk or the number of OADM/OXC nodes on a link, respectively,
  and filter-number which is the number of narrow filters on a link.
  When the alternative distance-only approach is chosen, the link-
  specific information is link-distance.

- In addition to the link-specific information, bounds on each of
  the impairments need to be quantified. Since these bounds are
  determined by the system designer's impairment allocations, these
  will be system dependent. For PMD, the constraint is that the sum
  of the link-PMD-square of all links on the transparent segment is
  less than the square of (a/B) where B is the bit rate. Hence, the
  required information is the parameter "a". For ASE, the
  constraint is that the sum of the link-noise of all links is no
  larger than P/SNRmin. Thus, the information needed include the
  launch power P and OSNR requirement SNRmin. The minimum
  acceptable OSNR, in turn, depends on the strength of the FEC being
  used and the margins reserved for other types of impairments.
  Other bounds include the maximum span length of the transmission
  system, the maximum path crosstalk or the maximum number of
  OADM/OXC nodes, and the maximum number of narrow filters, all are
  bit rate dependent. With the alternative distance-only approach,
  the upper bound is the maximum-path-distance. In single-vendor
  "islands" some of these parameters may be available in a local or
  EMS database and would not need to be advertised

- It is likely that the physical layer parameters do not change
  value rapidly and could be stored in some database; however these
  are physical layer parameters that today are frequently not known
  at the granularity required. If the ingress node of a lightpath
  does path selection these parameters would need to be available at
  this node.

-   The specific constraints required in a given situation will depend
    on the design and engineering of the domain of transparency; for
    example it will be essential to know whether chromatic dispersion
    has been dealt with on a per-link basis, and whether the domain is
    operating in a linear or nonlinear regime.

-   As optical transport technology evolves, the set of constraints
    that will need to be considered either explicitly or via a
    domain-wide margin may change. The routing and control plane
    design should therefore be as open as possible, allowing
    parameters to be included as necessary.

-   In the absence of wavelength conversion, the necessity of finding
    a single wavelength that is available on all links introduces the
    need to either advertise detailed information on wavelength
    availability, which probably doesn't scale, or have some mechanism
    for probing potential routes with or without crankback to
    determine wavelength availability. Choosing the route first, and
    then the wavelength, may not yield acceptable utilization levels
    in mesh-type networks.

5. More Complex Networks

   Mixing optical equipment in a single domain of transparency that has
   not been explicitly designed to interwork is beyond the scope of this
   document. This includes most multi-vendor all-optical networks.

   An optical network composed of multiple domains of transparency
   optically isolated from each other by O/E/O devices (transponders) is
   more plausible. A network composed of both "opaque" (optically
   isolated) OLXCs and one or more all-optical "islands" isolated by
   transponders is of particular interest because this is most likely
   how all-optical technologies (such as that described in Sec. 2) are
   going to be introduced. (We use the term "island" in this discussion
   rather than a term like "domain" or "area" because these terms are
   associated with specific approaches like BGP or OSPF.)

   We consider the complexities raised by these alternatives now.

   The first requirement for routing in a multi-island network is that
   the routing process needs to know the extent of each island. There
   are several reasons for this:

-   When entering or leaving an all-optical island, the regeneration
    process cleans up the optical impairments discussed in Sec. 3.

-   Each all-optical island may have its own bounds on each
    impairment.

- The routing process needs to be sensitive to the costs associated
  with "island-hopping".

This last point needs elaboration. It is extremely important to
realize that, at least in the short to intermediate term, the
resources committed by a single routing decision can be very
significant: The equipment tied up by a single coast-to-coast OC-192
can easily have a first cost of $10**6, and the holding times on a
circuit once established is likely to be measured in months.
Carriers will expect the routing algorithms used to be sensitive to
these costs. Simplistic measures of cost such as the number of
"hops" are not likely to be acceptable.

Taking the case of an all-optical island consisting of an "ultra
long-haul" system like that in Fig. 3-1 embedded in an OEO network of
electrical fabric OLXCs as an example: It is likely that the ULH
system will be relatively expensive for short hops but relatively
economical for longer distances. It is therefore likely to be
deployed as a sort of "express backbone". In this scenario a carrier
is likely to expect the routing algorithm to balance OEO costs
against the additional costs associated with ULH technology and route
circuitously to make maximum use of the backbone where appropriate.
Note that the metrics used to do this must be consistent throughout
the routing domain if this expectation is to be met.

The first-order implications for GMPLS seem to be:

- Information about island boundaries needs to be advertised.

- The routing algorithm needs to be sensitive to island transitions
  and to the connectivity limitations and impairment constraints
  particular to each island.

- The cost function used in routing must allow the balancing of
  transponder costs, OXC and OADM costs, and line haul costs across
  the entire routing domain.

Several distributed approaches to multi-island routing seem worth
investigating:

- Advertise the internal topology and constraints of each island
  globally; let the ingress node compute an end-to-end strict
  explicit route sensitive to all constraints and wavelength
  availabilities. In this approach the routing algorithm used by
  the ingress node must be able to deal with the details of routing
  within each island.

- Have the EMS or control plane of each island determine and
  advertise the connectivity between its boundary nodes together
  with additional information such as costs and the bit rates and
  formats supported. As the spare capacity situation changes,
  updates would be advertised. In this approach impairment
  constraints are handled within each island and impairment-related
  parameters need not be advertised outside of the island. The
  ingress node would then do a loose explicit route and leave the
  routing and wavelength selection within each island to the island.

- Have the ingress node send out probes or queries to nearby gateway
  nodes or to an NMS to get routing guidance.

6. Diversity

6.1. Background on Diversity

   "Diversity" is a relationship between lightpaths. Two lightpaths are
   said to be diverse if they have no single point of failure. In
   traditional telephony the dominant transport failure mode is a
   failure in the interoffice plant, such as a fiber cut inflicted by a
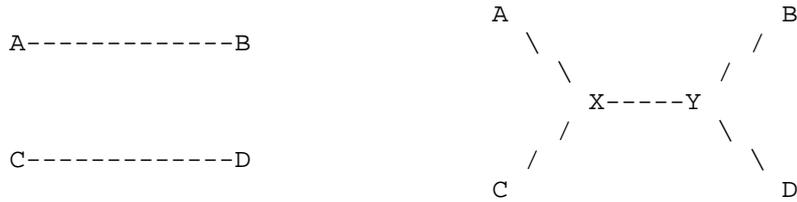   backhoe.

   Why is diversity a unique problem that needs to be considered for
   optical networks? Traditionally, data network operators have relied
   on their private line providers to ensure diversity and so have not
   had to deal directly with the problem. GMPLS makes the complexities
   handled by the private line provisioning process, including
   diversity, part of the common control plane and so visible to all.

   To determine whether two lightpath routings are diverse it is
   necessary to identify single points of failure in the interoffice
   plant. To do so we will use the following terms: A fiber cable is a
   uniform group of fibers contained in a sheath. An Optical Transport
   System will occupy fibers in a sequence of fiber cables. Each fiber
   cable will be placed in a sequence of conduits - buried honeycomb
   structures through which fiber cables may be pulled - or buried in a
   right of way (ROW). A ROW is land in which the network operator has
   the right to install his conduit or fiber cable. It is worth noting
   that for economic reasons, ROWs are frequently obtained from
   railroads, pipeline companies, or thruways. It is frequently the
   case that several carriers may lease ROW from the same source; this
   makes it common to have a number of carriers' fiber cables in close
   proximity to each other. Similarly, in a metropolitan network,
   several carriers might be leasing duct space in the same RBOC
   conduit. There are also "carrier's carriers" - optical networks
   which provide fibers to multiple carriers, all of whom could be
   affected by a single failure in the "carrier's carrier" network. In

a typical intercity facility network there might be on the order of
100 offices that are candidates for OLXCs. To represent the inter-
office fiber network accurately a network with an order of magnitude
more nodes is required. In addition to Optical Amplifier (OA) sites,
these additional nodes include:

- Places where fiber cables enter/leave a conduit or right of way;

- Locations where fiber cables cross; Locations where fiber splices
  are used to interchange fibers between fiber cables.

An example of the first might be:

```
                                   A                   B
   A------------B                   \                 /
                                      \             /
                                        X-----Y
                                      /             \
   C------------D                   /                 \
                                   C                   D

   (a) Fiber Cable Topology       (b) Right-Of-Way/Conduit Topology
```

         Figure 6-1: Fiber Cable vs. ROW Topologies

Here the A-B fiber cable would be physically routed A-X-Y-B and the
C-D cable would be physically routed C-X-Y-D. This topology might
arise because of some physical bottleneck: X-Y might be the Lincoln
Tunnel, for example, or the Bay Bridge.

Fiber route crossing (the second case) is really a special case of
this, where X and Y coincide. In this case the crossing point may
not even be a manhole; the fiber routes might just be buried at
different depths.

Fiber splicing (the third case) often occurs when a major fiber route
passes near to a small office. To avoid the expense and additional
transmission loss only a small number of fibers are spliced out of
the major route into a smaller route going to the small office. This
might well occur in a manhole or hut. An example is shown in Fig.
6-2(a), where A-X-B is the major route, X the manhole, and C the
smaller office. The actual fiber topology would then look like Fig.
6-2(b), where there would typically be many more A-B fibers than A-C
or C-B fibers, and where A-C and C-B might have different numbers of
fibers. (One of the latter might even be missing.)

```
              C                              C
              |                            /   \
              |                          /       \
              |                        /           \
          A------X------B          A--------------B

     (a) Fiber Cable Topology     (b) Fiber Topology
```
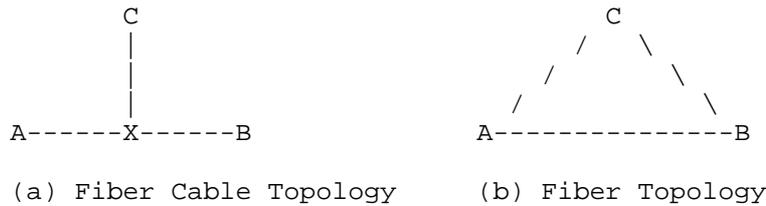
              Figure 6-2. Fiber Cable vs Fiber Topologies

The imminent deployment of ultra-long (>1000 km) Optical Transport
Systems introduces a further complexity: Two OTSes could interact a
number of times. To make up a hypothetical example: A New York -
Atlanta OTS and a Philadelphia - Orlando OTS might ride on the same
right of way for x miles in Maryland and then again for y miles in
Georgia. They might also cross at Raleigh or some other intermediate
node without sharing right of way.

Diversity is often equated to routing two lightpaths between a single
pair of points, or different pairs of points so that no single route
failure will disrupt them both. This is too simplistic, for a number
of reasons:

- A sophisticated client of an optical network will want to derive
   diversity needs from his/her end customers' availability
   requirements. These often lead to more complex diversity
   requirements than simply providing diversity between two
   lightpaths. For example, a common requirement is that no single
   failure should isolate a node or nodes. If a node A has single
   lightpaths to nodes B and C, this requires A-B and A-C to be
   diverse. In real applications, a large data network with N
   lightpaths between its routers might describe their needs in an
   NxN matrix, where (i,j) defines whether lightpaths i and j must be
   diverse.

- Two circuits that might be considered diverse for one application
   might not be considered diverse for in another situation.
   Diversity is usually thought of as a reaction to interoffice route
   failures. High reliability applications may require other types
   of failures to be taken into account. Some examples:

   o Office Outages: Although less frequent than route failures,
      fires, power outages, and floods do occur. Many network
      managers require that diverse routes have no (intermediate)
      nodes in common. In other cases an intermediate node might be
      acceptable as long as there is power diversity within the
      office.

      o Shared Rings: Many applications are willing to allow "diverse"
         circuits to share a SONET ring-protected link; presumably they
         would allow the same for optical layer rings.

      o Disasters: Earthquakes and floods can cause failures over an
         extended area. Defense Department circuits might need to be
         routed with nuclear damage radii taken into account.

   - Conversely, some networks may be willing to take somewhat larger
      risks. Taking route failures as an example: Such a network might
      be willing to consider two fiber cables in heavy duty concrete
      conduit as having a low enough chance of simultaneous failure to
      be considered "diverse". They might also be willing to view two
      fiber cables buried on opposite sides of a railroad track as being
      diverse because there is minimal danger of a single backhoe
      disrupting them both even though a bad train wreck might
      jeopardize them both. A network seeking N mutually diverse paths
      from an office with less than N diverse ROWs will need to live
      with some level of compromise in the immediate vicinity of the
      office.

   These considerations strongly suggest that the routing algorithm
   should be sensitive to the types of threat considered unacceptable by
   the requester. Note that the impairment constraints described in the
   previous section may eliminate some of the long circuitous routes
   sometimes needed to provide diversity. This would make it harder to
   find many diverse paths through an all-optical network than an opaque
   one.

   [Hjalmtysson00] introduced the term "Shared Risk Link Group" (SRLG)
   to describe the relationship between two non-diverse links. The
   above examples and discussion given at the start of this section
   suggests that an SRLG should be characterized by 2 parameters:

   - Type of Compromise: Examples would be shared fiber cable, shared
      conduit, shared ROW, shared optical ring, shared office without
      power sharing, etc.)

   - Extent of Compromise: For compromised outside plant, this would
      be the length of the sharing.

   A CSPF algorithm could then penalize a diversity compromise by an
   amount dependent on these two parameters.

Two links could be related by many SRLGs. (AT&T's experience
indicates that a link may belong to over 100 SRLGs, each
corresponding to a separate fiber group.) Each SRLG might relate a
single link to many other links. For the optical layer, similar
situations can be expected where a link is an ultra-long OTS.

The mapping between links and different types of SRLGs is in general
defined by network operators based on the definition of each SRLG
type. Since SRLG information is not yet ready to be discoverable by
a network element and does not change dynamically, it need not be
advertised with other resource availability information by network
elements. It could be configured in some central database and be
distributed to or retrieved by the nodes, or advertised by network
elements at the topology discovery stage.

6.2. Implications For Routing

Dealing with diversity is an unavoidable requirement for routing in
the optical layer. It requires dealing with constraints in the
routing process, but most importantly requires additional state
information (e.g., the SRLG relationships). The routings of any
existing circuits from which the new circuit must be diverse must
also be available to the routing process.

At present SRLG information cannot be self-discovered. Indeed, in a
large network it is very difficult to maintain accurate SRLG
information. The problem becomes particularly daunting whenever
multiple administrative domains are involved, for instance after the
acquisition of one network by another, because there normally is a
likelihood that there are diversity violations between the domains.
It is very unlikely that diversity relationships between carriers
will be known any time in the near future.

Considerable variation in what different customers will mean by
acceptable diversity should be anticipated. Consequently we suggest
that an SRLG should be defined as follows: (i) It is a relationship
between two or more links, and (ii) it is characterized by two
parameters, the type of compromise (shared conduit, shared ROW,
shared optical ring, etc.) and the extent of the compromise (e.g.,
the number of miles over which the compromise persisted). This will
allow the SRLGs appropriate to a particular routing request to be
easily identified.

7. Security Considerations

We are assuming OEO interfaces to the domain(s) covered by our
discussion (see, e.g., Sec. 4.1 above). If this assumption were to
be relaxed and externally generated optical signals allowed into the

domain, network security issues would arise. Specifically,
unauthorized usage in the form of signals at improper wavelengths or
with power levels or impairments inconsistent with those assumed by
the domain would be possible. With OEO interfaces, these types of
layer one threats should be controllable.

A key layer one security issue is resilience in the face of physical
attack. Diversity, as describe in Sec. 6, is a part of the solution.
However, it is ineffective if there is not sufficient spare capacity
available to make the network whole after an attack. Several major
related issues are:

- Defining the threat: If, for example, an electro-magnetic
   interference (EMI) burst is an in-scope threat, then (in the
   terminology of Sec. 6) all of the links sufficiently close
   together to be disrupted by such a burst must be included in a
   single SRLG. Similarly for other threats: For each in-scope
   threat, SRLGs must be defined so that all links vulnerable to a
   single incident of the threat must be grouped together in a single
   SRLG.

- Allocating responsibility for responding to a layer one failure
   between the various layers (especially the optical and IP layers):
   This must be clearly specified to avoid churning and unnecessary
   service interruptions.

The whole proposed process depends on the integrity of the impairment
characterization information (PMD parameters, etc.) and also the SRLG
definitions. Security of this information, both when stored and when
distributed, is essential.

This document does not address control plane issues, and so control-
plane security is out of scope. IPO control plane security
considerations are discussed in [Rajagopalam04]. Security
considerations for GMPLS, a likely control plane candidate, are
discussed in [Mannie04].

8. Acknowledgments

This document has benefited from discussions with Michael Eiselt,
Jonathan Lang, Mark Shtaif, Jennifer Yates, Dongmei Wang, Guangzhi
Li, Robert Doverspike, Albert Greenberg, Jim Maloney, John Jacob,
Katie Hall, Diego Caviglia, D. Papadimitriou, O. Audouin, J. P.
Faure, L. Noirie, and with our OIF colleagues.

9. References

9.1. Normative References

   [Goldstein94]   Goldstein, E. L., Eskildsen, L., and Elrefaie, A. F.,
                     Performance Implications of Component Crosstalk in
                     Transparent Lightwave Networks", IEEE Photonics
                     Technology Letters, Vol.6, No.5, May 1994.

   [Hjalmtysson00]   Gsli Hjalmtysson, Jennifer Yates, Sid Chaudhuri and
                     Albert Greenberg, "Smart Routers - Simple Optics: An
                     Architecture for the Optical Internet, IEEE/OSA
                     Journal of Lightwave Technology, December 2000, Vo
                     18, Issue 12, Dec. 2000, pp. 1880-1891.

   [ITU]   ITU-T Doc. G.663, Optical Fibers and Amplifiers,
                     Section II.4.1.2.

   [Kaminow97]   Kaminow, I. P. and Koch, T. L., editors, Optical
                     Fiber Telecommunications IIIA, Academic Press, 1997.

   [Mannie04]   Mannie, E., Ed., "Generalized Multi-Protocol Label
                     Switching (GMPLS) Architecture", RFC 3945, October
                     2004.

   [Rajagopalam04]   Rajagopalan, B., Luciani, J., and D. Awduche, "IP
                     over Optical Networks: A Framework", RFC 3717, March
                     2004.

   [Strand01]   Strand, J., Chiu, A., and R. Tkach, "Issues for
                     Routing in the Optical Layer", IEEE Communications
                     Magazine, Feb. 2001, vol. 39 No. 2, pp. 81-88.

   [Strand01b]   Strand, J., Doverspike, R., and G. Li, "Importance of
                     Wavelength Conversion In An Optical Network", Optical
                     Networks Magazine, May/June 2001, pp. 33-44.

   [Yates99]   Yates, J. M., Rumsewicz, M. P., and J. P. R. Lacey,
                     "Wavelength Converters in Dynamically-Reconfigurable
                     WDM Networks", IEEE Communications Surveys, 2Q1999
                     (online at
                     www.comsoc.org/pubs/surveys/2q99issue/yates.html).

9.2. Informative References

   [Awduche99]   Awduche, D. O., Rekhter, Y., Drake, J., R. and
                    Coltun, "Multi-Protocol Lambda Switching: Combining
                    MPLS Traffic Engineering Control With Optical
                    Crossconnects", Work in Progress.

   [Gerstel2000]   Gorstel, O., "Optical Layer Signaling: How Much Is
                    Really Needed?" IEEE Communications Magazine, vol. 38
                    no. 10, Oct. 2000, pp. 154-160

   [Kaminow02]   Ivan P. Kaminow and Tingye Li (editors), "Optical
                    Fiber Communications IV: Systems and Impairments",
                    Elsevier Press, 2002.

   [Passmore01]   Passmore, D., "Managing Fatter Pipes," Business
                    Communications Review, August 2001, pp. 20-21.

   [Ramaswami98]   Ramaswami, R. and K. N. Sivarajan, Optical Networks:
                    A Practical Perspective, Morgan Kaufmann Publishers,
                    1998.

   [Strand02]   John Strand, "Optical Network Architecture
                    Evolution", in [Kaminow02].

   [Tkach98]   Tkach, R., Goldstein, E., Nagel, J., and J. Strand,
                    "Fundamental Limits of Optical Transparency", Optical
                    Fiber Communication Conf., Feb. 1998, pp. 161-162.

10. Contributing Authors

   This document was a collective work of a number of people. The text
   and content of this document was contributed by the editors and the
   co-authors listed below.

   Ayan Banerjee
   Calient Networks
   6620 Via Del Oro
   San Jose, CA 95119
   EMail: abanerjee@calient.net

   Prof. Dan Blumenthal
   Eng. Science Bldg., Room 2221F
   Department of Electrical and Computer Engineering
   University of California
   Santa Barbara, CA 93106-9560
   EMail: danb@ece.ucsb.edu

Dr. John Drake
Boeing
2260 E Imperial Highway
El Segundo, Ca 90245
EMail: John.E.Drake2@boeing.com

Andre Fredette
Hatteras Networks
PO Box 110025
Research Triangle Park, NC 27709
EMail: afredette@hatterasnetworks.com

Change Nan Froberg's reach info to:
Dr. Nan Froberg
Photonic Systems, Inc.
900 Middlesex Turnpike, Bldg #5
Billerica, MA 01821
EMail: nfroberg@photonicsinc.com

Dr. Taha Landolsi
King Fahd University
KFUPM Mail Box 1026
Dhahran 31261, Saudi Arabia
EMail: landolsi@kfupm.edu.sa

James V. Luciani
900 Chelmsford St.
Lowell, MA 01851
EMail: james_luciani@mindspring.com

Dr. Robert Tkach
32 Carriage House Lane
Little Silver, NJ 07739
908 246 5048
EMail: tkach@ieee.org

   Yong Xue
   Dr. Yong Xue
   DoD/DISA
   5600 Columbia Pike
   Falls Church VA 22041
   EMail: yong.xue@disa.mil

Editors' Addresses

   Angela Chiu
   AT&T Labs
   200 Laurel Ave., Rm A5-1F13
   Middletown, NJ 07748

   Phone: (732) 420-9061
   EMail: chiu@research.att.com

   John Strand
   AT&T Labs
   200 Laurel Ave., Rm A5-1D33
   Middletown, NJ 07748

   Phone: (732) 420-9036
   EMail: jls@research.att.com