Optimal Design of Optical Ring Networks with Differentiated Reliability (DiR)*

Andrea Fumagalli, Marco Tacca

Center for Advanced Telecommunications Systems and Services (CATSS)
University of Texas at Dallas
Erik Jonsson School of Engineering and Computer Science
E-mail: {andreaf,mtacca}@utdallas.edu

Abstract. Current optical networks typically offer two degrees of service reliability: full protection in presence of a single fault in the network, and no protection at all. This situation reflects the historical duality that has its roots in the once divided telephone and data environment. The circuit oriented service required *protection,* i.e., provisioning of readily available spare resources to replace working resources in case of a fault. The datagram oriented service relied upon *restoration,* i.e., dynamic search for and reallocation of affected resources via such actions as routing table updates.
The current development trend, however, is gradually driving the design of networks towards a unified solution that will jointly support traditional voice and data services as well as a variety of novel multimedia applications. The growing importance of concepts, such Quality of Service (QoS) and Differentiated Services that provide varying levels of service performance in the same network evidences this trend.
Consistently with this pattern, the novel concept of *Differentiated Reliability (DiR)* is formally introduced in the paper and applied to provide multiple reliability degrees (classes) in the same network layer using a common protection mechanism, i.e., path switching. According to the DiR concept, each connection in the layer under consideration is guaranteed a minimum reliability degree, defined as the Maximum Failure Probability allowed for that connection. The reliability degree chosen for a given connection is thus determined by the application requirements, and not by the actual network topology, design constraints, robustness of the network components, and span of the connection.
An efficient algorithm is proposed to design the Wavelength Division Multiplexing (WDM) layer of a DiR ring.

1 Introduction

Current statistics reveal that businesses increasingly rely on communications networks. E-commerce and e-business are heavily dependent upon the availabil-

* The work was partially supported by the Texas Higher Education Coordinating Board under contract # 009741-0139-1999, TxTEC, and NSF under contract # ANI-0082085.

ity of communication resources. Thus, reliable communication networks are a must from the end user's point of view.

In a multi-layer network architecture, protection switching techniques may be implemented in more than a single layer. For instance, in the IP (Internet Protocol) over WDM (Wavelength Division Multiplexing) architecture it is possible to implement protection[1] mechanisms at the WDM layer [1], or a set of distinct protection mechanisms via Multi Protocol Label Switching (MPLS) at the IP layer [2], or restoration[2] mechanisms in the conventional IP manner [3]. With these options available, a competitive and cost effective design of the reliable network is a critical goal that may determine the success or failure of the proposed network architecture. In other words, which layer should provide protection and what degree of protection does each layer need? It is well known that distinct protection mechanisms will coexist in the same network, with each mechanism being implemented at a distinct layer [4]. One rule of the thumb is that using the protection mechanism at the lowest possible layer minimizes the response time [5, 6].

Little is known, however, about the degree of reliability required at each layer, that is, the level of reliability that each layer should provide to higher layers. Clearly, there is a tradeoff between the degree of reliability that is offered by each layer and the cost of the resources required at that layer. A higher degree of reliability comes at higher cost. Another factor to consider is that the same network (or layer) is designed to support multiple services, and each service might require a distinct degree of reliability. For example, the WDM layer may support both SONET and IP layers, and may provide to each higher layer a different degree of reliability.

Currently, only two degrees of service reliability are commonly considered: 100% protection in presence of a single fault in the network and none at all. These two approaches have their historical origin in the once divided telephone and data environments. The former service, being circuit oriented, requires protection for immediate network restoration upon a fault. The latter service, being datagram oriented, relies upon restoration, i.e., dynamic updates of the network routing tables. The current trend is gradually driving the design of networks towards a unified solution that will support, not only traditional voice and data services, but also a variety of novel multimedia applications. Evidence of this trend over the last decade is the introduction of concepts like Quality of Service (QoS) [7, 8] and Differentiated Services [9, 10] to provide varying levels of service performance in the same network. Surprisingly, little has been discussed or proposed for network reliability schemes to accommodate this change in the way networks are being designed.

[1] *Protection* indicates a technique based on the provisioning of readily available resources to replace specific working connections in case of a fault.

[2] *Restoration* indicates a technique based on the dynamic search for alternative resources that will allow to continue the connections affected by the fault.

Today's competitive networks can no longer provide just pure circuit switching and datagram services, nor they can limit the options of reliability to only two degrees, either fully protected or unprotected.

In this paper the problem of designing cost effective multi-layer networks that are capable of providing various reliability degrees — as opposed to 0% and 100% only — is addressed. The concept of *Differentiated Reliability (DiR)* is for the first time formally introduced and applied to provide multiple reliability degrees (or classes) at the same layer using a common protection mechanism, i.e., path switching [11, 12].

According to the DiR concept, each connection in the layer under consideration is assigned a Maximum Failure Probability (MFP) which is defined as the probability that the connection is unavailable due to the occurrence of a fault in the network. The MFP chosen for a given connection is determined by the application requirements and not by the protection mechanism[3], the network topology, the source-destination distance, or the Mean Time Between Failure (MTBF) of the network components. It is expected that the cost of the connection is inversely proportional to the chosen MFP degree. For example, in a simplified network scenario with only two layers, i.e., IP over WDM, an optical connection between two IP routers shall guarantee the MFP required by the user. Certain connections between routers may be more critical than others, thus, a cost effective solution will rely on some highly reliable optical connections and other less reliable connections whose disruption will not heavily affect the user's IP network.

From a practical point of view, DiR connections can rely upon the protection mechanism preferred by the network designer, or available from technology. Connections are then individually routed and assigned spare resources in such a way that the MFP degree required for each connection is met. Preemption of a less reliable connection to guarantee reliability of a higher reliable connection is allowed when the resulting MFP degree still meets the connection requirement.

In summary, the DiR concept offers the unique combination of the following advantages. The user (or upper layer) determines the desired MFP degree for each connection. Thus, in a multi-layer network, the lower layer may provide the above layers with the desired reliability degree, transparently from the actual network topology, design constraints, device technology and connection span. In addition, DiR allows to dynamically adjust the network configuration to take into account possible improvements of the network component's MTBF that may become available in some portion of the network, without affecting the reliability degree offered to the upper layers. In practical terms, improved MTBF of network components may allow the same network to support additional connections without affecting the guaranteed MFP degree of the already existing connections.

In the rest of the paper the DiR concept is demonstrated using a bidirectional WDM ring in which protection is achieved in the optical layer.

[3] With DiR, the protection mechanism may offer a variety of different MFP degrees

2 Applying the DiR Concept to WDM Rings

First the problem of optimally designing the WDM DiR-based ring is defined, then an efficient algorithm that sub-optimally solves the problem is presented.

2.1 The Problem of DiR Optimal Design

This subsection describes the assumptions made and defines the DiR optimal design problem.

It is assumed that a set of connection demands are given and must be routed across the ring. A demand consists of one or multiple lightpaths[4], that need to be established between two nodes. Each demand is assigned a requested MFP that must be met by the protection mechanism at the optical layer. Only one protection mechanism is available at the optical layer to achieve the demand requested MFP. This requirement appears to be necessary to provide feasible network management as opposed to complex management handling of multiple concurrent recovery actions that take place in the event of a network element failure. The protection mechanism considered in the paper is the $1:1$ path protection applied to lightpaths. With this protection mechanism, a *working lightpath* is assigned a route-disjoint *protection lightpath* that may be alternatively used if the working lightpath fails to work.

First, the set of multiple *reliability classes* is introduced. A reliability class, c, is characterized by a connection maximum failure probability $MFP(c)$, which indicates the maximum acceptable probability that, upon a network element failure, a connection in that class will not survive despite the optical layer protection[5].

Connection demands are assigned to classes according to their required MFP. Traffic demands in higher requirement classes, i.e., lower MFP, are those with the most stringent requirements in terms of recovery speed and need to be restored at the optical layer (e.g., voice traffic), as opposed to traffic demands in lower classes that may also be recovered by means of the (relatively slower) IP restoration mechanisms.

The WDM ring topology is modeled as a graph $G(\mathcal{N}, E)$. Set $\mathcal{N}$ represents the network nodes and link set E represents the WDM ring lines connecting physically adjacent nodes. Each link in E is characterized by two numbers: link cost and link failure probability.

The link cost represents the cost of routing a lightpath on that link. For example, if the link cost is set equal to the link length, it is assumed that the cost of routing a lightpath on that link is proportional to the link length.

The link failure probability is the probability that the considered link is faulted under the condition that a single network line fault has occurred in the ring. (The analysis presented here is based on the assumptions that only single

[4] A lightpath is a path of light between a node pair, whose bandwidth equals the wavelength bandwidth.

[5] The concept of reliability class can be easily generalized to multiple failures.

(line) faults may occur. However, the proposed technique can be extended to handle concurrent faults of varying natures, including node faults.) The link failure probability is estimated on the basis of available failure statistics of the employed optical components. First, the probability of having a single fault in the network is estimated. Once this value is known, every link fault probability is normalized to the probability of having a single fault in the network. We define $P_f(i,j)$ as the failure probability of link (i,j), given the occurrence of one fault in the network. A uniform distribution of faults across the link, would then result in $P_f(i,j) = \frac{1}{|E|} \ \forall (i,j) \in E$.

To demonstrate the DiR concept, the $1:1$ path protection scheme is modified as follows. (Notice that provisioned protection resources are unused in absence of network failure.) A lower class working lightpath may be routed on already provisioned protection wavelengths (Fig. 1), thus improving use efficiency of network resources. Consequently, in case of a link failure, a higher class connection may preempt a lower class connection if the latter is using protection resources dedicated to the former. This mechanism affects the connection failure probability of the lower class connections due to its preemption. From a lower class connection viewpoint, preemption is equivalent to a fault that disrupts its working lightpath without the possibility to resort to a protection lightpath. For each connection, we thus define the connection failure probability as the sum of the failure probabilities of its unprotected links plus the probability of "virtual" failure due to preemption.

In Fig. 1, we show an instance of the ring with uniform line failure distribution. The higher class connection (thick line) has $MFP = 0$, i.e., it must be 100% protected against any single line fault. The lower class connection (thin line) reuses protection wavelengths assigned to the higher class connection (lines (D,C) and (C,B)), thus its failure probability is given by the sum $P_f(D,C) + P_f(C,B)$ (fault of the unprotected links of the lower class connection), plus the probability of being preempted by the higher class connection which is $P_f(D,E) + P_f(E,A)$.

The objective of the DiR problem is to determine the routing for and the resources used by each lightpath request in order to minimize the ring total *wavelength-mileage*, or Λ — more generally, the network cost — subject to guaranteeing the requested reliability degree (MFP) of each traffic class c.

The objective function is formally described next under the assumption that wavelength converters are available at all nodes and any (working or protection) lightpath may be assigned a distinct wavelength on each of its links.

$$min\ \Lambda = \sum_{i,j} \left(\sum_{c=1}^{C} \sum_{s,d} (w_{i,j}^{s,d,c} + p_{i,j}^{s,d,c} - r_{i,j}^{s,d,c}) \cdot l_{i,j} \right) \tag{1}$$

where:

- $w_{i,j}^{s,d,c}$ is the number of (working) wavelengths on line (i,j) assigned to the working lightpaths that belong to the connection demand from s to d of class c

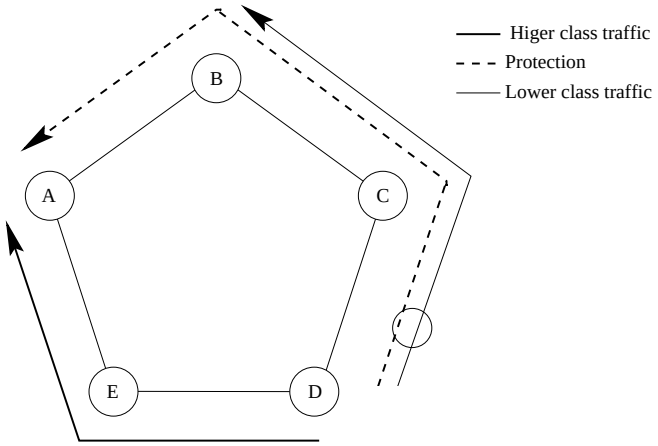

Fig. 1. Protection and preemption mechanism

- $p_{i,j}^{s,d,c}$ is the number of (protection) wavelengths on line (i,j) assigned to the protection lightpaths that belong to the connection demand from s to d of class c
- $r_{i,j}^{s,d,c}$ is the number of protection wavelengths on line (i,j) that are reused by working lightpaths that belong to the connection demand from s to d of class c
- $l_{i,j}$ is the length of line (i,j).

The objective function must be minimized under the constraint that the probability failure of every connection demand in class c does not exceed the requested maximum failure probability $MPF(c)$.

2.2 The Difficult-Reuse-First Algorithm

This section presents an efficient greedy algorithm that may be used to sub-optimally[6] solve the DiR design problem in WDM rings. The name of the algorithm is Difficult-Reuse-First (DRF) to indicate that, prior to being routed, demands are sorted considering how difficult it is for the corresponding working lightpaths to achieve reuse of protection wavelengths. Lightpaths that are more difficult in that regard, are routed first, thus giving them a better chance to achieve a more efficient wavelength reuse.
The DRF algorithm is based on two observations.

1. Due to the topological layout of bidirectional ring, only two disjoint routes exist between any node pair (Fig. 2). Routing is thus a binary problem and each demand that requires protection must employ both routes, one for the working lightpath and the other for the protection lightpath (Fig. 2(b)).

[6] The complexity of the DiR design problem is under investigation.

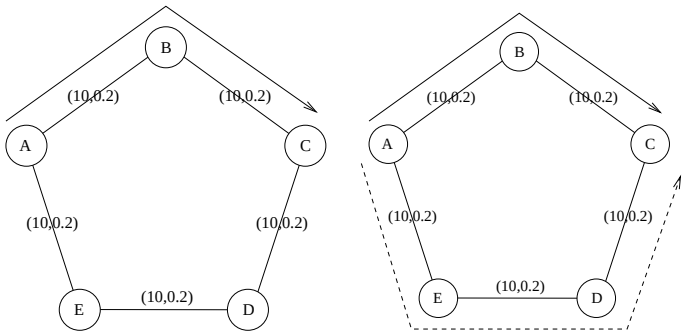

(a) Unprotected traffic (b) Protected traffic

Fig. 2. Ring protection mechanism

2. When the reliability degree requested for the connection cannot be achieved using the already provisioned wavelengths, an additional wavelength must be added to the ring. Under this circumstance, irrespective of the routing (clockwise or counterclockwise) chosen for the working (protection) lightpath, one wavelength must be added to every line of the ring. The cost (working and protection) for setting the connection thus equals the cost of adding a wavelength along the whole ring perimeter.

The following definitions are used in the description of the DRF algorithm. The failure probability of a path is given by the sum of the failure probabilities of the links along the path. The minimum failure probability between two nodes, mfp_{sd}, is the minimum between the failure probability of the clockwise and the counterclockwise path. Let $MFP(c)$ be the maximum failure probability for connection demands in class c.
The algorithm is organized in 7 steps as depicted in Fig. 3.
Step 1. Connection demands are classified using two sets: the set of demands that require some degree of protection and the set of demands that do not require protection. Demand of class c from node s to node d belongs to the former set if $mfp_{sd} > MFP(c)$. It belongs to the latter set otherwise.
Step 2. The working lightpath for each demand in the former set is routed using the shortest path in terms of number of links. The protection lightpath is routed using the opposite direction. Shortest path in terms of number of links is chosen for the working lightpath to yield the maximum number of protection wavelengths in the ring that may be reused by the algorithm at some later step. Notice that these connections have failure probability equal to 0, as they all survive any single fault in the ring.
Step 3. The demands in the latter set are sorted according to increasing values of the difference $X = (MFP(c) - mfp_{sd}) \geq 0$ where c, s and d are, respectively,

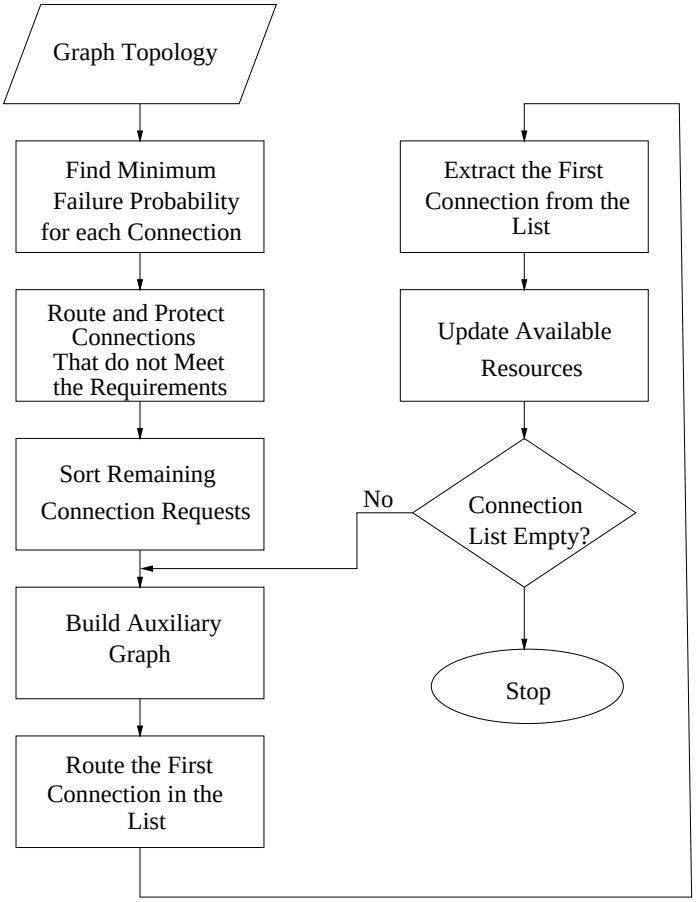

Fig. 3. Algorithm flow chart

the class, the source and destination of the demand. Let S_X be the set of sorted demands. Value X indicates the *excess of reliability* offered to the demand if a newly wavelength was added to each link of the the path with minimum failure probability mfp_{sd}. Since the excess of reliability is not necessary, the algorithm looks for ways to reuse some of the already provisioned protection wavelengths in place of the newly added wavelengths. When a protection wavelength is used in place of a newly added wavelength, the reliability excess of the demand is reduced due to the potential preemption of that wavelength. Notice that the reliability excess must not be reduced below zero, otherwise the required reliability degree requested for the demand is not met.

Intuitively, smaller values of X correspond to demands with lower probability to achieve reuse of protection wavelengths. In order to achieve a more efficient

wavelength reuse — which in turn corresponds to a lower total network cost — demands having smaller X are routed first, hence the name of the algorithm.

Step 4. The demand in set S_X with the smallest value of X is considered for routing. Prior to routing the demand, an auxiliary graph $G'(\mathcal{N}, E')$ is built. $\mathcal{N}$ is the set of network nodes, E' is the union of set $E \in G$ and set E_p. The latter is the set of links that represent the provisioned protection lightpaths (or segments of them) not yet reused. Set E_p is constructed as follows. Let s be the source node, d the destination node and $I = \{n_1, n_2, \ldots, n_k\}$ the ordered set of intermediate nodes of a protection lightpath. Link (s, d), all links (s, n_i), where $n_i \in I$, all links (n_i, d), where $n_i \in I$, and all links (n_i, n_j), where $n_i, n_j \in I$ and $i < j$ are added to E_p when the associated protection lightpath is added to the ring. Link (s, d) represents the entire protection lightpath. Any other link represents only a segment of the protection lightpath. When multiple protection lightpaths share the same source and destination pair the set of links is derived only once, taking into account the multiplicity of the lightpaths by assigning a capacity greater than one to the links.

Any link in $l \in E'$ is assigned a triple: a cost, a failure probability, and a capacity that indicates the maximum number of working lightpaths one can route on such link. Every link in E' derived from E is assigned its original cost, i.e., the length of its line, its original failure probability, and infinite capacity. Every link in E_p is assigned zero cost as routing on any such link represents reuse of an already provisioned protection lightpath (or segment of it). Each link in E_p is assigned a failure probability that is the sum of two terms: the failure probability of the working lightpath associated with the protection lightpath (or segment of it) represented by the link, i.e., the probability of preemption; and the failure probability of the represented protection lightpath (or segment of it). Each link in E_p is assigned a capacity equal to the number of protection lightpaths (or segment of lightpaths) represented by the link.

Fig. 4(a) represents the auxiliary graph obtained from the ring shown in Fig. 4(b). In Fig. 4(a) the original topology graph with two working connections — the solid lines — is shown. Links are assigned two numbers that represent, respectively, the length of the network line and the network line failure probability. It is assumed that connection demand from node D to node A requires protection, represented by the dotted line in the figure. Connection demand from node D to node C belongs to a lower requirement reliability class and thus it may reuse some of the protection wavelengths provisioned for the former demand. The corresponding auxiliary graph is shown in Fig. 4(b). Set E' consists of the union of link set E and the additional links (C, B), (B, A), and (C, A) (set E_p), that represent, respectively, the possibility to reuse protection wavelengths on the original graph links (C, B), (B, A), and the concatenation of (C, B) and (B, A). The auxiliary graph does not contain link (D, C), since the protection wavelength on that link is already used. Notice that added link (C, A), which represents a segment of the protection lightpath associated with the working lightpath from node D to node A, has failure probability given by the sum of three terms: the probability that the working lightpath from D to A is rerouted

because of a line fault (0.4); the probability that line (C, B) is faulted (0.2); and the probability that line (B, A) is faulted (0.2).

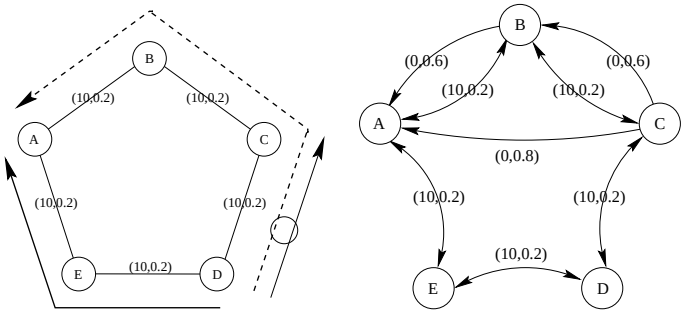

(a) Original graph (b) Auxiliary graph

Fig. 4. New links in the auxiliary graph

Since the auxiliary graph may have a large number of parallel links that may considerably slow down the execution time of the algorithm, a pruning technique is used to reduce the number of links in E', without significantly affecting the algorithm performance. Two steps are used to reduce the number of links in E'. First, links with failure probability greater than the demand MFP are removed from the auxiliary graph. Second, when multiple links with zero cost exist between a node pair, only the link with the smallest failure probability is kept in set E' (Fig. 5), the other zero cost links are not inserted in the auxiliary graph set E'. Notice that links of the auxiliary graph derived from E are never removed from the auxiliary graph as they represent the possibility to add wavelengths to the ring when the reuse of provisioned protection wavelengths is not possible.

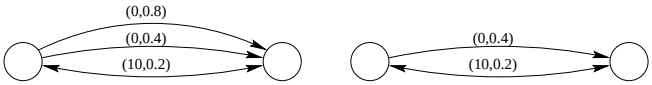

Fig. 5. Pruning of parallel multiple links with zero cost

Step 5. The demand under consideration is routed using the Dijkstra shortest path algorithm [13] applied to the pruned auxiliary graph. The link weight l_{ij}^w used by the shortest path is computed as a linear combination of the link length (l_{ij}) and the link failure probability $(P_f(i,j))$. The linear combination is controlled by parameter a as follows: $l_{ij}^w = a \cdot l_{ij} + (1-a) \cdot P_f(i,j)$. When $a = 1$

the algorithm returns the shortest path in terms of mileage, when $a = 0$ the algorithm returns the shortest path in terms of failure probability, i.e., the most reliable path.

By varying a, a dichotomic search is performed. Let $a_{max} = 1$ and $a_{min} = 0$. The shortest path algorithm is run using the mean value $a = a_{int} = (a_{max} + a_{min})/2)$. Let $pf_{a_{int}}$ be the failure probability of the path so found. If $pf_{a_{int}}$ is greater than the MFP requested for the demand a_{max} is set equal to a_{int}. If $pf_{a_{int}}$ is smaller than the MFP, a_{min} is set equal to a_{int}. These steps are iterated until $pf_{a_{max}} \leq MFP$.

Once the path is found, the capacity of the path is computed as the minimum capacity of all links that belong to the path. Working lightpaths of the demand are then routed along the path, up to its capacity. If the capacity of the shortest path is not large enough to accommodate all lightpaths of the demand, Step 5 is repeated until all lightpaths are routed.

Step 6. The wavelengths assigned to the demand under analysis are provisioned and the corresponding demand is removed from set S_X.

Step 7. The status of the provisioned protection wavelengths is updated taking into account possible reuse. Notice that such update may result in changes of the auxiliary graph. The algorithm goes back to step 4 until set S_X is emptied.

3 Performance Results

The DRF algorithm described in section 2.2 is tested using a ring topology that comprises 20 nodes connected by equal length lines of 10 miles. It is assumed that line failure probability is uniformly distributed among the network lines, i.e., $P_f(i,j) = 1/20 \forall (i,j) \in E$. Connection demands are divided into three classes.

- Class 1 demands are uniform, requiring one lightpath between every node pair (s,d). $MFP(1) = p_1$.
- Class 2 demands are uniform, requiring two lightpaths between every node pair (s,d). $MFP(2) = p_2$.
- Class 3 demands are uniform, requiring three lightpaths between every node pair (s,d). $MFP(3) = p_3$.

The total number of lightpath requests is 2280. Presented results are obtained setting p_1 and p_2 to a fixed value, while p_3 takes on values in the interval $[0,1]$. With the given ring and connection demands, the DRF algorithm runs on a Linux-PC pentium $450MHz$, requiring a computational time in the order of tens of seconds.

Figs. 6, 7 and 8 show, respectively, the total λ-mileage (working and protection), the total protection λ-mileage, and the total protection λ-mileage reused by the DRF algorithm.

As expected, the required total λ-mileage decreases as the reliability requirement of class 3 becomes less stringent, i.e., p_3 increases. This is due to the combined effect of two factors: 1) the protection λ-mileage necessary to fulfill

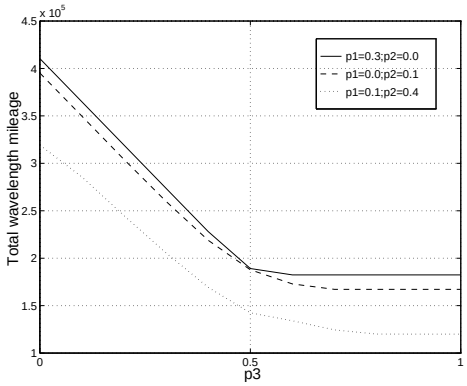

Fig. 6. Total wavelength mileage

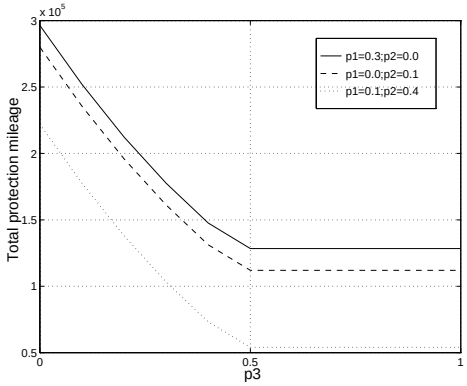

Fig. 7. Total protection mileage

the requested reliability degree is reduced (Fig. 7); 2) protection wavelength reuse is improved (Fig. 8).

It is interesting to note that when $p_1 = 0.1$, $p_2 = 0.4$ and $p_3 \to 1$, the total λ-mileage found by the DRF algorithm equals the same λ-mileage that would be required in the same ring, with the same connection demands, had conventional shortest paths been used to route the working lightpaths and no protection lightpaths been provisioned at all. While costing the same, the two approaches are however significantly different. With the shortest path approach the requested demand reliability is not always guaranteed and shorter working lightpaths have better reliability degree than longer lightpaths have. With the DRF algorithm, demands in the three classes always meet their reliability degree, without regard of the distance between the source and the destination.

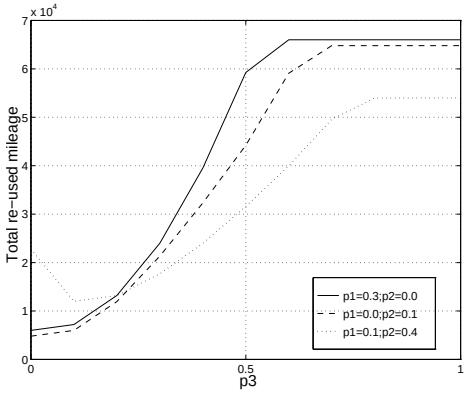

Fig. 8. Total reused mileage

Figs. 9, 10, and 11 numerically demonstrate the above conjecture by plotting the distribution of the connection failure probability for the demands in the three reliability classes when $p_1 = 0.1$, $p_2 = 0.4$, and $p_3 = 0.8$. The results obtained with the shortest path algorithm and the DRF algorithm are reported. The total λ-mileage is the same in both cases. The shortest path routing yields the same distribution of connection failure probability for every reliability class, irrespective of the reliability requirement of the class. (The distribution is actually scaled by a factor equal to the number of lightpath requests in each class.) The DRF algorithm trades higher connection failure probabilities for class 1 traffic, with lower connection failure probabilities for the other two classes. By so doing, the requested class reliability degree is met for all connection demands.

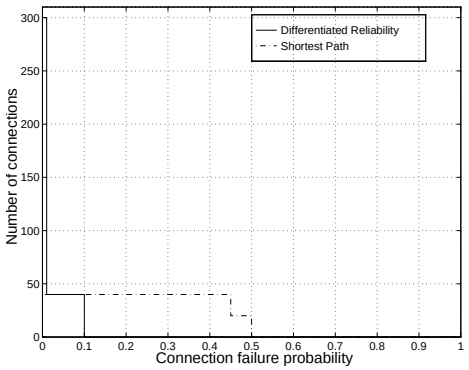

Fig. 9. Class 1 connection failure probability distribution

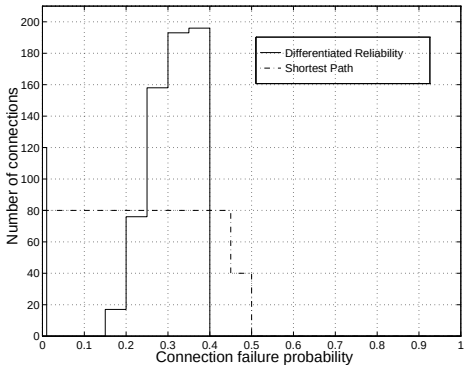

Fig. 10. Class 2 connection failure probability distribution

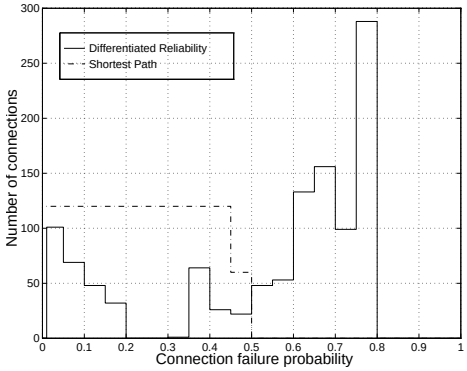

Fig. 11. Class 3 connection failure probability distribution

From a revenue viewpoint, the DiR approach allows to accommodate all connection demands with the requested degree of reliability. The conventional shortest path approach, on the contrary, imposes a de-classification of a large portion of the demands in class 1 and class 2 below the requested reliability degree, thus reducing the potential overall revenues, yet requiring the same λ-mileage of the DiR approach.

4 Summary

The paper presented the novel concept of Differentiated Reliability (DiR) classes of traffic, according to which, connections are differentiated based on their

requested individual degree of reliability. The proposed concept enables to differentiate traffic flows in classes without regard for network topology, equipment MTBF, and most importantly connection span, both in terms of line hops and mileage.

An algorithm was proposed to sub-optimally design DiR-based optical rings. Presented results demonstrated the potential advantages of the proposed concept, in terms of both overall network costs and guaranteed reliability degree for every traffic class. Beside the optical layer discussed in the paper, the DiR concept may find other interesting applications in a variety of network layers, including the Multi Protocol Label Switching (MPLS) layer.

References

1. IEEE Communications Society, *IEEE Communication Magazine*, vol. 37, No 8, August 1999, Special Issue on Survivable Communication Networks.
2. S. Makan, V. Sharma, K. Owens, and C. Huang, "Protection/restoration of MPLS networks," draft-makam-mpls-protection-00.txt, Expiration day April 2000, IETF, October 1999.
3. R.W. Stevens, *TCP/IP Illustrated, Volume 1*, vol. 1, Addison-Wesley Professional Computing Series, 1994.
4. P. Demeester, M. Gryseels, A. Autenrieth, C. Brianza, L. Castagna, G. Signorelli, R. Clemente, M. Ravera, A. Lajszczyk, D. Janukowicz, K. Van Doorselaere, and Y. Harada, "Resilience in multilayer networks," *IEEE Communications Magazine*, vol. 37, No 8, August 1999.
5. P. Demeester, M. Gryseels, K. Van Doorselaere, A. Autenrieth, C. Brianza, G. Signorelli, R. Clemente, M. Ravera, A. Jajszcyk, A. Geyssens, and Y. Harada, "Network resilience strategies in SDH/WDM multilayer networks," in 24^{th} *European Conference on Optical Communication (ECOC) '98*, Madrid, 1998, vol. 1, pp. 579–80.
6. D. Cavendish, "Evolution of optical transport technologies: From SONET/SDH to WDM," *IEEE Communications Magazine*, June 2000.
7. M. Ajmone Marsan, A. Bianco, E. Leonardi, A. Morabito, and F. Neri, "All-optical WDM multi-rings with differentiated QoS," *IEEE Communications Magazine*, vol. 37, No 2, February 1999.
8. A. Jukan, A. Monitzer, and H. R. Van As, "QoS-restorability in optical networks," in *Proceedings of* 24^{th} *European Conference on Optical Communication (ECOC'98)*, 1998, vol. 1, pp. 711–712.
9. IEEE Communications Society, *IEEE Network*, vol. 13, No 5, September/October 1999, Special Issue on Integrated and Differentiated Services for the Internet.
10. N. Golmie, T.D. Ndousse, and D.H Sue, "A differentiated optical services model for WDM networks," *IEEE Communications Magazine*, vol. 38, No 2, February 2000.
11. T.H. Wu, *Fiber Optic Survivability*, Artech House, 1992.
12. R. Ramaswami and K. N. Sivarajan, *Optical Networks: a pratical prospective*, Morgan Kaufmann Publishers, Inc., 1998.
13. M. Gondran and M. Minoux, *Graphs and algorithms*, Wiley Interscience, 1979.