An Analytical Approach to Shared Backup Path Provisioning in GMPLS Networks

SuKyoung Lee, David Griffith, and Nah-Oak Song
National Institute of Standards and Technology
100 Bureau Drive, Stop 8920, Gaithersburg, MD 20899

Abstract—**As GMPLS and its supporting set of protocols develop into a viable control plane for optical networks, an important function that they will need to support will be the restoration and protection function that has been a major feature of legacy optical networks. Several models have been proposed for protection with GMPLS using shared backup paths. This previous work has not investigated the effect on recovery time(i.e. service interruption time) critical to the service or the number of backup paths that are required to meet a desired level of performance. Using both recovery time and recovery failure probability, we have developed a new analytic model for GMPLS-based recovery in $M : N$ protection groups.**

I. Introduction

Recovery of traffic is growing in importance and especially fast recovery after failures has become a very important issue at layers above the optical layer. Restoration and protection functionality in both MPLS and GMPLS network, has been pointed out as strong candidate in this area and may be motivated by the notion that there are inherent limitations to improving the recovery time of current routing algorithms. Furthermore, a recovery mechanism using GMPLS could enable IP traffic to be put directly over WDM optical channels, without an intervening SONET layer, while still emulating SONET resiliency features. This would facilitate the construction of IP-over-WDM networks. For recovery in IP over WDM network, even if link-layer restoration such as mesh restoration is recommended to achieve low latencies, IP level recovery, based on GMPLS architecture is employed in the event that link-layer restoration fails.

In GMPLS networks, there are mainly two techniques for recovery: restoation and protection The distinctive difference between protection and restoration is recovery time scale. While restoration is flexible due to its' dynamic path establishment, it takes an order of magnitude longer to restore traffic than protection. Considering one of the most challenging problems is recovery of failures under tight time constraints, protection is preferred for real time traffic in GMPLS networks. There have been many works addressing protection functionality based on GMPLS architecure. (Here some references should be cited.) However, most of the works have focused on $1 : 1$ dedicated protection mechanism, even though GMPLS-based recovery intended different protection modes to provide carriers with the flexibility depending on a spectrum of service levels. In this paper, we propose an analytic model of $M : N$ protection in GMPLS networks.

This work was supported, in part, by the National Institute of Standards and Technology (NIST), the Advanced R&E Activity (ARDA), the Laboratory for Telecommunications Sciences (LTS) MENTER project, the Defense Advanced Research Projects Agency (DARPA) Fault Tolerant Networks (FTN) program, and the National Communications System (NCS).

It is generally desirable to have protection scheme which is resource-efficient. In GMPLS-based recovery, it is important to increase network reliability by providing necessary resources in time as well as enabling a fast response to failures. Based on our proposed model, backup path provisioning mechanism is studied in order to reflect this tradeoff between resource utilization and reliability upon GMPLS-based recovery. In particular, there must be a mechanism to advertise backup path resource and processing rules must be defined for bandwidth accounting when some failure are notified. Therefore, recovery time is also investigated in our model when $M : N$ shared protection is performed.

Furthermore, we analyze the recovery request failure probability numerically for the case that multiple failures occur upon a path in $M : N$ protection with revertive mode.

II. Background

Providing differentiation of services and service guarantees in networks is promising to be a major revenue collector for service providers. This has increased the importance of gaining control over networks via traffic engineering. However, the requested QoS should be guaranteed even in case of network failures. The main objective of any recovery scheme is to operate in a cost-effective manner while minimizing service interruptions to the customer. Providing a high degree of reliability (or equivalently, a low probability of service disruptions) is expensive and tends not to scale well. For this reason, any carrier that operates a wide-area optical backbone network needs to be able to support a variety of service classes in which the degree of protection is tied to the price of the service [1]. For instance, [2] proposed a multi-tiered service model in which the basic (least expensive) service receives no protection support, while more expensive service options feature some various combinations of routing around areas with a relatively high probability of network failure and dedicating backup paths for automatic failover switching of the data stream.

There are mainly two levels of recovery mechanisms: restoration and protection switching. While restoration is defined as the real-time establishment of appropriate resources to recover affected traffic, protection switching involves the establishment of pre-calculated replacement resources. In the latter scheme, pre-calculation of backup paths differs from the traditional on-demand scheme i.e., restoration. The pre-calculation of backup paths can be carried out under either dedicated mechanism such as $1 + 1$ and $1 : 1$, or under $M : N$ shared mechanism. Thus the problem of QoS supporting with GMPLS-based recovery is translated into not only computing backup path but also select-

TABLE I
PROTECTION LEVEL EXAMPLE

Service level	Protection plan
Gold	Dedicated protection: $1+1$, $1:1$
Silver	Shared protection: $M:N$
Bronze	Rerouting

ing a scheme among $1+1$, $1:1$, and $M:N$ schemes, for every QoS request. In the 1+1 dedicated protection scheme, traffic is passing through both the working and backup paths. Upon detection of a primary path failure, the traffic on the backup path becomes the active traffic. Because of its double-reservation nature, it is the fastest protection switched recovery mechanism, but also the most expensive in terms of resources. In 1:1 protection scheme, the backup path is used only after a failure is detected. To protect against N failures, $M:N$ scheme pre-establishs M(usually less than N) backup paths that are Shared Risk Group(SRG)-disjoint from the working paths. M protection entities are shared among N working resources. Thus, a protection priority could be used as a differentiating mechanism for premium services that require high reliability as can be seen in Table. I.

As a control signaling to differentiate the protection level of each path, the field *Service Type*(8 bits) in Generalized Label Request can be used[5] that is similar to *Service Type* defined in [6]. Accordingly, a carrier may specify a range of different classes of service (e.g. gold, silver, bronze) with different types of recovery plans where there could exist no recovery, dedicated protection, shared protection and etc. as can be seen the protection level example in Table I.

The capacity reserved for protection could be considered as the main cost of recovery QoS from the fact that protection capacity along backup as well as primary paths then provides QoS guarantees. The amount of capacity reserved in the network depends on the number of backup paths and on the protection scheme used. Especially, for $M:N$ case, the full reservation of all M backup paths will result in wasting the resources in network. Hence at other times the backup paths can be used to carry lower priority traffic even when the resources for backup paths are preallocated, or the backup paths can be reserved on demand. If lower priority traffic is using resources along the backup paths, the edge nodes may need to be notified of the failure in order to complete the switchover [3],[4]. Therefore, even if the backup path is pre-signaled, it takes time to switch the traffic to the backup path allowing preemption. In this paper, based on this protection switching time, namely, recovery time, we propose a new approach to provision M backup paths for $M:N$ shared protection. We analyze recovery failure probability on the basis of recovery time. According to recovery failure probability and recovery time, the protection manager can allocate different capacity for each protection group, that results from provisioning different number of backup paths. In other words, proper provisioning of backup paths leads to a more efficient capacity usage. Besides efficient capacity utilization over the backup paths, it is important to note that the number of backup paths passing through a link also indicates label space(e.g. number of lambdas) and label table sizes that are generally limited in the amount.

Our approach could provide different levels of guaranteed services for various $M:N$ backup groups in terms of both the number of backup paths and service interruption time, around network failures. Thus even shared protection groups that are M backup groups as Silver service in Table I, can be differentiated according to restorability requested by each service class (backup group).

III. Backup Path Provisioning

In protection, network can quickly utilize pre-provisioned backup resources for recovery from a resource failure along the working (primary) path. Meanwhile, at the time when a failure occurs, the network state is not static, i.e. the number of occupied backup paths and the number of faults are different. Actually, some amount of protocol signaling is required at the time of failure. This varies from simply propagating the error from the point of detection to the point of recovery, to the full signaling of the backup path. Thus, it is usually difficult to predict how many backup paths will be necessary for the shared backup path case. In spite of this difficulty, it is not desirable to use real-time approach like restoartion for some high priority services since the approach requires time to compute the alternate path after failure is detected and hence is likely to be slower. In consideration of the tradeoff among recovery time and pre-provisioned resource, we will analyze the recovery time to provision the shared backup path efficiently before a failure happens.

In this section, we investigate the number of enough backup paths to recover the data on the working paths based on a model for the recovery signaling time. The number of attempts depends on current network status, that is, how many backup paths are used and if the resources are available in the backup path.

A. Recovery Time Analysis

The time taken from the instant a link fails to the instant the backup path of a connection traversing the failed path is enabled, could be defined to be the protection-switching time for the connection. Our recovery time analysis concentrates on this protection-switching time. As soon as a failure is detected on a working path, an attempt will be made to restore the working path. We assume that the control network is reliable, i.e., does not incur message losses.

Assume that there is an infinite number of feasible backup paths $\{P_1, P_2, \ldots\}$ for attempts. The backup paths will be attempted until the recovery is successfully made. For the i^{th} attempt to a backup path P_i, it takes time t_i to check if the path P_i is available for the recovery function. And assume that these times $t_1, t_2, \ldots$ are independent and identically distributed (i.i.d.) random variables having a distribution $F_t(t)$.

Let a path with a failure need K attempts until the recovery is successfully made. That is, the first $K-1$ attempts find that the paths $P_1, P_2, \cdots P_{K-1}$ are not available but the K^{th} attempt finds that the path P_K is available for recovery. Then the recovery time T_r, which is required for finding an available path to restore a working path with failure, is

$$T_r = t_1 + t_2 + \cdots + t_K \qquad K \geq 1 \tag{1}$$

It is also assumed that each attempt is successful with probability p, that is, each backup path is available for recovery process with probability p. Thus, the expected number of attempts that will be required to activate a backup path is

$$\begin{aligned} E[K] &= \sum_{K=1}^{\infty} K(1-p)^{K-1}p \\ &= \frac{1}{p}. \end{aligned} \quad (2)$$

Since $t_1, t_2, \cdots$ are i.i.d. random variables with finite expected values and K is a stopping time for $t_1, t_2, \cdots$, we can apply renewal theory to Eq. 1. Then, we have

$$E[T_r] = E[K]E[t], \quad (3)$$

where $E[t] = \int_0^\infty t\, dF_t(t)$. For the case where t_K is exponentially distributed with mean $1/\mu$, $E[T_r]$ becomes $\frac{1}{p\mu}$.

Each traffic flow will have its own expected recovery time limit. The network QoS manager could use the result from Eq. 3 as a constraint on the requested recovery time. The average recovery time is indicative of the expected amount of data lost during a failure. That is, during the time required to activate the backup path and switch the traffic over to it, the affected connection will experience data (and revenue) losses. For example, a sudden disconnect during an active transaction in a network of ATM machines or other systems can cause uncertain states from which the end application may not recover, causing failure of the transaction. Thus, it is imperative to facilitate seamless handover of data so that information loss is minimized.

B. Number of Backup Paths

To prevent excessive resource usage for backup paths, and to meet the implicit service provider requirement of improving network resource utilization so as to increase the number of potential future demands that can be used for protection, it is important to determine the appropriate number of backup paths to be shared.

When a failure occurs, up to k attempts will be made to find a backup path. If the k attempts fail, then the recovery attempt is considered to have failed and a new working path must be created for the customer. Thus, regardless of whether the recovery attempt succeeds, the system will spend T_r^k units of time trying to set up a backup path, where

$$\begin{aligned} E[T_r^k] &= pE[t] + 2(1-p)pE[t] + \cdots \\ &\quad +(k-1)(1-p)^{k-2}pE[t] + k(1-p)^{k-1}E[t] \\ &= \frac{1-(1-p)^k}{p}E[t]. \end{aligned} \quad (4)$$

Suppose that as part of the Service Level Agreement(SLA) that a carrier has with the customer, there is an upper limit ϵ on the expected recovery time. This would be requested by a service class with shared backup protection (e.g. Silver class in Table I). Thus the expected recovery time must satisfy

$$E[T_r^k] \le \epsilon. \quad (5)$$

Substituting Eq. 4 into InEq. 5 results in

$$\frac{(1-(1-p)^k)}{p}E[t] \le \epsilon. \quad (6)$$

The above InEq. can be expressed as

$$\frac{\ln(1-\frac{\epsilon}{E[t]}p)}{\ln(1-p)} \ge k \quad (7)$$

From InEq. 7, the maximum number of shared backup paths can be computed satisfying the requested recovery time of the service class.

For premium services, the network operator may also want to guarantee a certain probability of recovery success in the event of a failure. In other words, we may demand that the probability of recovery failure after k attempts does not exceed some limit, δ. So we require

$$P[\text{failure}] = (1-p)^k \le \delta, \quad (8)$$

which implies that

$$k \ge \frac{\ln(\delta)}{\ln(1-p)}. \quad (9)$$

From InEq. 9, we can derive the minimum number of shared backup paths.

In accordance with the grade of service survivability, the carrier could determine the minimum or the maximum numbers of shared backup paths. If δ or ϵ is given according to the requested QoS, the other limit could be also determined such that

$$\frac{1-\delta}{p}E[t] \le \epsilon, \qquad \text{equivalently} \qquad 1 - \frac{p\epsilon}{E[t]} \le \delta. \quad (10)$$

Then, as soon as the QoS limits are determined, the carrier could restrict k to lie within a range of values given by

$$\frac{\ln(\delta)}{\ln(1-p)} \le k \le \frac{\ln(1-\frac{\epsilon}{E[t]}p)}{\ln(1-p)}. \quad (11)$$

As can be seen in Fig. 1, some carriers can refer the above range in accordance with the requested QoS for recovery time and recovery failure rate. Normally, if the customers' traffic is so critical, then one would (to meet the SLA) assign a separate (or at least shared) backup path for this particular Label Switched Path (LSP). If the network is properly designed and used, the situation where no backup LSP is available, when the primary LSP fails, should not arise. In the event a new service request comes in and a backup cannot be found (and reserved) due to bandwidth exhaustion or for whatever reason, then the request (with protection LSP) should be denied. If the customer agrees to an unprotected LSP service, then depending upon the SLA, "best effort" service in the event of a node/link failure could be provided. If the unprotected LSP service cannot be provided also, then the request for this service is also denied, and depending upon the SLA only "best effort" service may be provided.

IV. $M : N$ Shared Protection with Revertive Mode

The applications requesting high reliability, began to require a variety of failures to be taken into account. Among the failures,

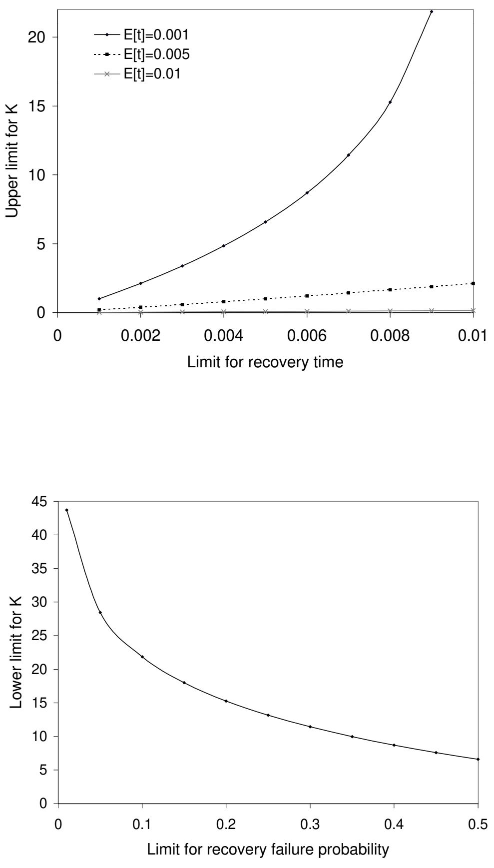

Fig. 1. Range for the number of backup paths

our analysis have focused on $M : N$ protection with revertive mode [8] in GMPLS network since it is generally desirable that the alternate path can be switched back to the original working path once the failure is repaired in order to assure an optimized survivable network architecture. In $M : N$ protection, up to N working paths are protected using M protection paths which should be diversely routed. Under this protection mode, we investigate the effective failure occurrence for non-simultaneous multiple failures. In our model, we assume that the following two paths cannot be restored to another backup path for next failure before switching back to its original working path: First one is the path which has been using a protection path since previous fault, and the second one is the path which is already in the recovery operation due to previous failure. For the both cases, a higher-layer rerouting mechanism will be used to set up an alternate connection path as can be seen in Fig. 2, where the procedure is associated with the activation of backup path. This approach is slower than the protection switching mechanism and so we use it only as a last resort.

It is assumed that a mechanism for detecting and isolating multiple failures is in place in the network. In general, failures are detected by lower mechanisms. The lower mechanism passes up an alarm to an GMPLS control entity as soon as a node detects a failure. To do the analysis, we can use some of the theoretical framework developed in [9] for detecting and isolating multiple failures in WDM networks.

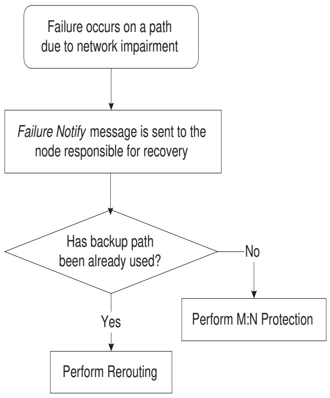

Fig. 2. Procedure for activation of backup path

A. Recovery Blocking Rate

In this analysis, it is assumed that there are N working paths and $M < N$ backup paths in a $M : N$ protection domain. Let λ be the failure occurrence rate in a working path. The time for traffic to revert from a backup path to its original working path is exponentially distributed with rate μ. And let π_i be the steady state probability that i backup paths are used. In the state diagram as inFig. 3, state i corresponds to i backup paths being in use, and a transition from state i to state $i+1$ occurs with rate $(N-i)\lambda$ for $i < m$. Let n_f be the number of recovery requests by a failure occurrence upon a working path, n_r be the number of recovery completions (the number of accepted recovery requests), n_a be the number of recovery blockings because the traffic on the working path is already using a protection path, and n_b be the number of recovery blockings because no backup path is available. It is clear that

$$n_f = n_r + n_a + n_b. \tag{12}$$

From the assumption for revertive mode, the effective failure occurrence rate per working path can be defined as

$$\lambda_f = \frac{n_f - n_a}{n_f}\lambda. \tag{13}$$

This λ_f is used to determine the number of necessary backup paths, not λ. Let p_f be the recovery blocking probability and p_f^* be the blocking probability that excludes the blocked recovery requests due to using a protection path. We have

$$p_f = \frac{n_b}{n_f}, \qquad p_f^* = \frac{n_b}{n_f - n_a}. \tag{14}$$

If $p_f = p_f^*$ then the system can be described using the Erlang distribution, while $p_f^*(\neq p_f)$ becomes π_m in our system model.

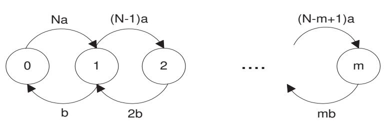

Fig. 3. State diagram for multiple failures(a=λ, b=μ)

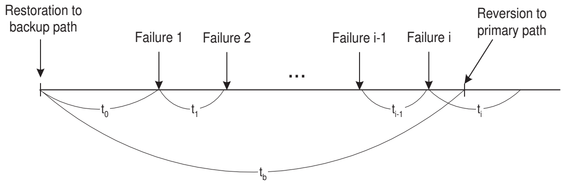

Fig. 4. Time model for multiple failures

We derive the probability p_f^* from the state diagram in Fig. 3. For $1 \leq i \leq m$, from [7],

$$\pi_i = \binom{N}{i} \left(\frac{\lambda}{\mu}\right)^i \pi_0. \tag{15}$$

Using the above Eq. 15 and the fact that $\pi_0 + \pi_1 + \cdots + \pi_m = 1$, the probability p_f^* can be expressed as

$$p_f^* = \pi_m = \frac{\binom{N}{m}\left(\frac{\lambda}{\mu}\right)^m}{\sum\limits_{0 \leq i \leq m} \binom{N}{i}\left(\frac{\lambda}{\mu}\right)^i} \tag{16}$$

If the system does not consider the reversion, where the system can be described using the Erlang distribution, then we can compute $p_f = p_f^*$, which is the probability that an Erlang system with m states is in State m:

$$p_f = \pi_m = \frac{\rho^m/m!}{\sum\limits_{n=0}^{m} \rho^n/n!}, \tag{17}$$

where $\rho = \lambda/\mu$. For the non-revertive mode in Eq. 17, depending on the configuration, the original working path may, upon being repaired, become the protection path, or it may be used for new working traffic. However, it is desirable to move the traffic to the original working path that is calculated based on network topology and network policies, gaining optimal network performance. Thus, we have more focused on the revertive mode developing expressions for some of the other probabilities related to the system in revertive mode.

Defining x to be the expected number of failures that occur while the working path is still using the protection path,

$$n_a = x n_r. \tag{18}$$

This follows from an examination of Fig. 4, which shows a scenario in which the interarrival time between failures is less than the average time required to allow traffic to revert to the original working path. From the figure we see that x is the mean number of failure events per recovery period. If the failure occurrences form a Poisson process with rate λ and the backup path holding times for each failure are exponentially distributed with mean $1/\mu$,

$$\begin{aligned} x &= \sum_{i=1}^{\infty} i Pr[t_b \leq t < t_b + t_i] \\ &= \sum_{i=1}^{\infty} i \int_{t=0}^{\infty}\int_{t_b=0}^{t}\int_{t_i=t-t_b}^{\infty} \mu e^{-\mu t}\frac{(\lambda t_b)^{i-1}}{(i-1)!} \\ &\quad \lambda e^{-\lambda t_b}\lambda e^{-\lambda t_i} dt_i dt_b dt \\ &= \sum_{i=1}^{\infty} \frac{i\lambda^i \mu}{(\lambda+\mu)^{i+1}} \\ &= \frac{\lambda}{\mu}, \end{aligned} \tag{19}$$

where $t_b = t_0 + t_1 + t_2 + \cdots + t_{i-1}$ when i failures occur while the connection is using the backup path, as can be seen in Fig. 4.

Using Eq.s 12, 14, and 18, we obtain the following probabilities. The recovery failure probability, accounting for failures that occur while traffic is on a backup path, is

$$p_f = \frac{p_f^*}{1 + (1 - p_f^*)\frac{\lambda}{\mu}}. \tag{20}$$

The probability of recovery request acceptance can be computed as

$$\begin{aligned} p_r &= \frac{n_r}{n_f} \\ &= \frac{1 - p_f^*}{1 + (1 - p_f^*)\frac{\lambda}{\mu}}, \end{aligned} \tag{21}$$

and the probability of recovery failure resulting from using a protection path is found in a similar manner to be

$$\begin{aligned} p_a &= \frac{n_a}{n_f} \\ &= x p_r \\ &= \frac{(1 - p_f^*)\frac{\lambda}{\mu}}{1 + (1 - p_f^*)\frac{\lambda}{\mu}}. \end{aligned} \tag{22}$$

From the above Eq. 22, we can get the effective failure occurrence rate as

$$\begin{aligned} \lambda_f &= \lambda(1 - p_a) \\ &= \frac{\lambda}{1 + (1 - p_f^*)\frac{\lambda}{\mu}}. \end{aligned} \tag{23}$$

This effective failure occurrence rate is informative in utilizing backup LSPs, because most carriers prefer to make the LSP to revert back to its original working path. Usually, the routing of the protection path may not be as efficient as the original one.

B. Performance Evaluation

The performance of the proposed analytical model is analyzed by considering recovery blocking rate, i.e. we characterize optical data network services by restorability. It is assumed that a failure occurs with exponential distribution (mean is 10) and recovery time is 1(simulation time unit) in the simulation test. After setting up not only 50(100 working paths in the other test) but 10 backup paths between ingress node and egress node, we generated failures over the working paths. Since these working paths are randomly chosen for each failure, some working paths could have multiple failures. It is assumed that all paths

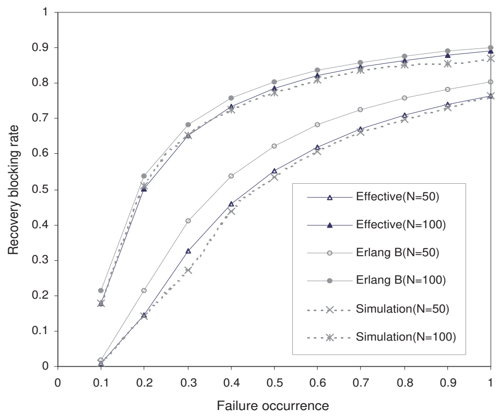

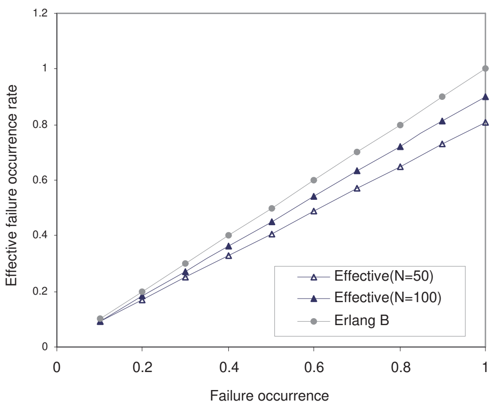

Fig. 5. Impact of multiple failures

are pre-calculated and wavelengths are pre-assigned to working and backup paths. Fig. 5 illustrates the impact of the multiple-failure effect comparing our model with the Erlang. In these graphs, $m = 10$ and two sets of curves are considered where one is $N = 50$ and the other is $N = 100$. The first graph in Fig. 5 indicates that our model is consistent with the simulation test. We observe that when N is small, the Erlang model is not appropriate to predict recovery blocking probability (restorability) for a GMPLS network with a lower number of failures. As for the second graph, when the number of failures in a network is small, each working path with failures is likely to send current traffic on a backup path and the subsequent failures are unlikely to get the recovery service under revertive mode. Thus, effective failure occurrence rate per working path becomes smaller than the failure occurrence rate which does not reflect reversion. When N is large, it is more likely that a failure is unable to use a backup path because there is no free backup path. These graphs show that our model is a good approximation that can be used for protection group sizing.

V. Conclusion

In this paper, we prooposed a new analytical model for shared backup path provisioning in GMPLS networks. In our model, protection bandwidth capacity was considered as the main cost of recovery QoS, with the result that different amount of backup resources could be assigned to services with different levels of protection. We have reviewed some of the ways that GMPLS, in combination with other QoS mechanisms, can be used to allow service providers to offer customized levels of protection to their customers. To determine the optimum size of a $M : N$ protection group given QoS constraints, we have developed a model that predicts the amount of time required to establish a backup path. We have also developed a model of effective failure occurrence rate for $M : N$ protection with reversion. The examination of our simulation results demonstrated that shared protection groups can be sized so that the probability that a backup path is unavailable is less than a desired threshold.

Finally, future work is to expand on this work by analyzing the effect of network topology on the probability of multiple failure events and by studying switchover delays in more detail. In particular, we are examining the behavior of several restoration signaling algorithms in a variety of failure scenarios.

References

[1] O. Gerstel and R. Ramaswami, "Optical Layer Survivability: A Services Perspective", *IEEE Communications Magazine*, vol. 38, no. 3, pp. 104-113, March 2000.

[2] H. Ishimatsu et al., "Carrier Needs Regarding Survivability and Maintenance for Switched Optical Networks", *Internet draft*, draft-hayata-ipo-carrier-needs-00.txt, Nov. 2000.

[3] P. Smith, et al, "Generalized MPLS Signaling - RSVP-TE Extensions", *Internet Draft*, draft-ietf-mpls-generalized-rsvp-te-06.txt, Nov. 2001.

[4] P. Smith, et al, "Generalized MPLS Signaling - CR-LDP Extensions", *Internet Draft*, draft-ietf-mpls-generalized-cr-ldp-01.txt, Mar. 2001.

[5] Y. Xu, et al, "Modification and Reorganization to GMPLS Signaling Functional Specification", *Internet Draft*, draft-xu-ccamp-gmpls-sig-reorg-00.txt, Mar. 2001.

[6] Many, "OIF UNI Signaling Specification", OIF2000.125.3, Feb. 2001.

[7] L. Kleinlock, "Queueing Systems: Theory, vol.I", *John Wiley & Sons*, NewYork, 1975.

[8] V. Sharmai et al, "Framework for MPLS-based Recovery", *IETF Draft*, draft-ietf-mpls-recovery-frmwrk-03.txt, Jul. 2001.

[9] C. Mas and P. Thiran, "An Efficient Algorithm for Locating Soft and Hard Failures in WDM Networks", *The IEEE Journal of Selected Areas in Communications*, vol. 18, no. 10, pp. 1900-1911, October, 2000.