Restorability versus Efficiency in (1:1)^n Protection Schemes for Optical Networks

David Griffith, Kotikalapudi Sriram, SuKyoung Lee, Stephan Klink, and Nada Golmie

***Abstract*— As network utilization continues to grow in the coming years, there will be increased pressure on network operators to use traffic engineering to provision resources more efficiently. One way to do this is to allow backup paths associated with disjoint working paths to share bandwidth. Increasing the amount of sharing will naturally increase the risk that a failed working path will either be unrecovered or forced to use dynamic recovery mechanisms. To examine the trade-offs between robustness and efficiency and to develop useful performance bounds, we develop theoretical models for (1:1)^n recovery schemes that are independent of the network's topology and management plane. We confirm our results using simulations of uncorrelated failures in a wide-area optical network with various degrees of resource sharing.**

I. Introduction

Optical networks are moving from SONET/SDH to ASON/ASTN architecture and a GMPLS control plane, allowing for more flexible operations but incorporating more complexity. ASON/ASTN networks must be able to recover from failures, hence the emerging standards support 1+1, 1:1, M:N, and other recovery modes for path, subpath, and span recovery. For the sake of efficiency, backup paths can share resources. The IETF introduced (1:1)^n and (M:N)^n notation to describe recovery schemes that are composed of multiple recovery schemes that share backup resources [1], [2]. In this paper, we consider the trade-off between recoverability and efficiency in (1:1)^n schemes. We have already examined the (M:N)^n case in another paper [3].

This paper is organized as follows. In Section II we introduce basic concepts and describe the theoretical framework that we use in our analysis. In Section III we derive lower and upper bounds for the recovery blocking probability for (1:1)^n recovery schemes. In Section IV we present numerical results obtained from network simulations and compare them to the results of our theoretical analysis.

II. System Model

A (1:1)^n recovery scheme comprises n 1:1 recovery groups, each of which consists of a pair of working and backup paths denoted as $(\mathrm{WP}_i, \mathrm{BP}_i)$. Normal traffic is carried on the working path, while the backup path may be used to carry extra traffic that would be preempted if the normal traffic were displaced by a failure. All WPs are mutually disjoint, i.e. $\mathrm{WP}_i \cap \mathrm{WP}_j = \emptyset$ if $i \neq j$.

For all the models that we will consider, we use a Boolean sharing matrix $\mathbf{S}$ to describe the degree of resource overlap among the set of backup paths, where each element s_{ij} is defined as follows.

$$s_{ij} = \begin{cases} 0, & \mathrm{BP}_i \cap \mathrm{BP}_j = \emptyset \\ 1, & \mathrm{BP}_i \cap \mathrm{BP}_j \neq \emptyset \end{cases} \quad (1)$$

For example, a set of n 1:1 protection groups whose backup paths do not overlap has a sharing matrix that is the $n \times n$ identity matrix.

A connectivity graph for a recovery scheme with sharing matrix $\mathbf{S}$ can be created by associating a labeled vertex with index $i \in \{0, 1, \ldots, n\}$ to each 1:1 protection group and drawing an edge between each pair of vertices i and j if $s_{ij} = 1$. If this graph is connected, we say that the associated recovery scheme is *irreducible*. A disconnected connectivity graph corresponds to a *reducible* recovery scheme that can be decomposed into two or more independent irreducible recovery schemes, each of whose size is less than n. The number of irreducible recovery schemes of size n is $1, 1, 4, 38, 728, 26704, \ldots$ for $n = 1, 2, 3, 4, 5, 6, \ldots$ (Sequence A001187 in [4]). We obtain this sequence by taking the inverse Euler transform of the sequence $x_n = 2^{n(n-1)/2}$. From this discussion, it is clear there are many possible sharing arrangements for a recovery system of size n, and analyzing them all is not practical. However, we can generate upper and lower performance bounds for all possible (1:1)^n recovery schemes by considering systems that feature, respectively, the minimum and maximum possible degrees of resource sharing. These bounds will allow a network operator to quickly determine the amount of backup resource sharing that is acceptable given the prevailing conditions in the network.

We now provide some definitions that will be used in the theoretical development. A *symmetric* recovery scheme is one in which the n 1:1 recovery groups can be labeled such that all the elements of each diagonal of the resulting sharing matrix $\mathbf{S}$ have the same value, i.e. $s_{ij} = s_{kl}$ if $|i - j| = |k - l|$. A *minimally connected* symmetric scheme is characterized by a sharing matrix $\mathbf{S}$ whose elements are given by

$$s_{ij} = \begin{cases} 1, & |i - j|_{\mathrm{mod}\, n-1} \leq 1 \\ 0, & \text{else.} \end{cases} \quad (2)$$

Conversely, a *maximally connected* symmetric scheme is characterized by a sharing matrix in which $s_{ij} = 1$ for all i and j. The corresponding resource sharing graph is fully connected, i.e., each vertex has degree $n - 1$.

D. Griffith, K. Sriram, S. Klink, and N. Golmie are with the National Institute of Standards and Technology (NIST), Gaithersburg, MD 20899 USA (contact e-mail: david.griffith@nist.gov).

S. Lee is with Sejong University, Seoul, Korea.

This research was partially supported by the National Communications System (NCS).

We provide an example to illustrate these concepts. In Fig. 1(a) we show an optical network in which four working paths are organized into a (1:1)^4 recovery scheme that is symmetric and minimally connected. For example, consider recovery group 1, which consists of a working path whose route is AB and a backup path whose route is $AEFB$. Group 1 shares recovery resources AE and FB with recovery groups 2 and 4, respectively. These sharing relationships are depicted in the resource sharing graph shown in Fig. 1(b). From the form of the graph, we see that the recovery scheme depicted in Fig. 1(a) is a symmetric minimally connected scheme. The corresponding sharing matrix is

$$\mathbf{S} = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \end{bmatrix}. \tag{3}$$

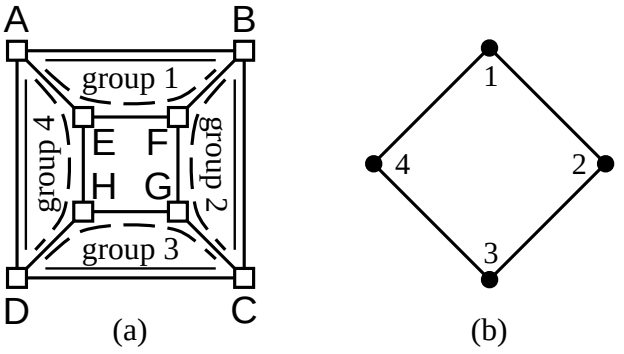

Fig. 1. **(a)**: An example optical network supporting a (1:1)^4 recovery scheme. **(b)**: The associated resource sharing graph.

To model a general (1:1)^n recovery scheme, we define λ to be the rate of working path failures over the entire network, and we define μ to be the rate of repair, so that the mean time to repair a failed working path is $1/\mu$. Failures and repairs are characterized by Poisson point processes. We assume that sufficient resources are available to the network operator so that failed WPs can be repaired in parallel, i.e., the mean repair time is not state-dependent. The state of the system is given by a Boolean vector $\mathbf{w}$ of length n, where the i^{th} element w_i is 1 if WP$_i$ is in a failed state and its traffic is using BP$_i$ and is 0 otherwise. Thus, if all n working paths are operational, the state vector has value $\mathbf{w} = \mathbf{0}$. We define $F(\mathbf{w})$ to be the number of failed WPs when the system is in state $\mathbf{w}$. It can be computed as

$$F(\mathbf{w}) = \sum_{i=1}^{n} w_i. \tag{4}$$

If a working path has failed and other backup paths share resources with its backup path, the ability of those working paths to recover from failures can be compromised. Consider an arbitrary working path WP$_j$. The set of working paths whose respective backup paths share resources with BP$_j$ is $\{\text{WP}_i : s_{ij} = 1\}$, so we can define $v_j = \sum_{i=1}^{n} w_i s_{ij}$ to be the number of failed working paths, including WP$_j$, that are using resources that are shared by BP$_j$. Using this definition, it follows that the row vector $\mathbf{v} = \{v_j\}$ can be computed as $\mathbf{v} = \mathbf{w}\mathbf{S}$.

If the (1:1)^n system is in state $\mathbf{w}$ with $F(\mathbf{w}) = m$, then $n-m$ WPs are operational, but some of them may not be able to recover in the event of a failure. We say that a working path WP$_i$ is *blocked* if there exists at least one backup path BP$_j$ such that $s_{ij} = 1$ and BP$_j$ is being used. Blocked WPs are assumed to be recovered by using dynamic rerouting while keeping the same BP. We can determine, for a given state $\mathbf{w}$, how many functioning WPs are blocked as follows. We define $B(\mathbf{w})$ to be the number of functioning WPs that will be blocked if they fail. A working path WP$_i$ is functioning if $w_i = 0$ and blocked if $v_i > 0$, so $B(\mathbf{w})$ is given by

$$B(\mathbf{w}) = \sum_{i=1}^{n} I_{\{0\}}(w_i) I_{\mathbb{Z}^+}(v_i), \tag{5}$$

where $\mathbb{Z}^+$ is the set of positive integers and where $I_A(x)$ is an indication function that returns a Boolean value based on the membership of the argument element x in the set A:

$$I_A(x) = \begin{cases} 0, & x \notin A \\ 1, & x \in A. \end{cases} \tag{6}$$

Because $\mathbf{w}$ is a Boolean vector, $I_{\{0\}}(w_i) = 1 - w_i$ is a valid indication function. Likewise we can use $I_{\mathbb{Z}^+}(v_i) = 1 - \delta(v_i)$, where $\delta(v_i)$ is the Kronecker delta function. This leads to the following expression:

$$B(\mathbf{w}) = \sum_{i=1}^{n} (1 - w_i)(1 - \delta(v_i)). \tag{7}$$

Similarly, for a state $\mathbf{w}$, we define $U(\mathbf{w})$ to be the number of functioning WPs that will not be blocked if they fail (because no BPs that share resources with their BPs are being used). A 1:1 group meets this criterion if its WP failure indicator w_i is zero and if its BP usage measure v_i is also zero. Thus $U(\mathbf{w})$ is

$$U(\mathbf{w}) = \sum_{i=1}^{n} I_{\{0\}}(w_i) I_{\{0\}}(v_i) = \sum_{i=1}^{n} (1 - w_i)\delta(v_i). \tag{8}$$

By combining (7) and (8) and using (4), we obtain

$$B(\mathbf{w}) + U(\mathbf{w}) = \sum_{i=1}^{n} (1 - w_i) = n - F(\mathbf{w}), \tag{9}$$

showing that $B(\mathbf{w}) + U(\mathbf{w}) + F(\mathbf{w}) = n$ for any state $\mathbf{w}$.

For the discussion that follows, it is useful to partition the universal set of all allowable states into subsets based on the number of active backup paths $F(\mathbf{w})$. If we define $\mathcal{W}$ to be the set of all states $\mathbf{w}$, we can define

$$\mathcal{W}_m = \{\mathbf{w} : F(\mathbf{w}) = m\} \tag{10}$$

for $m = 0, 1, 2, \ldots, n$. For example, $\mathcal{W}_0 = \{\mathbf{0}\}$. Clearly $\{\mathcal{W}_m\}_{m=0}^{n}$ is a cover of $\mathcal{W}$, i.e. $\mathcal{W} = \bigcup_{m=0}^{n} \mathcal{W}_m$.

We can illustrate the partitioning of the state space using the following example. Consider the state transition diagram shown in Fig. 2. Here four 1:1 protection groups are organized into a minimally connected symmetric recovery scheme of the type depicted in Fig. 1, with $\mathbf{S}$ given by (3). Fig. 2

shows transition rates between states; for instance, the balance equation for state $[0\,0\,0\,1]$ can be written by inspection as

$$(\lambda+\mu)p_{[0\,0\,0\,1]} = \lambda p_{[0\,0\,0\,0]} + \mu p_{[0\,1\,0\,1]}, \tag{11}$$

where $p_{\mathbf{w}}$ is the steady state probability that the system is in state $\mathbf{w}$. The dashed lines in the figure indicate the partitioning of $\mathcal{W}$ into subsets $\mathcal{W}_0$, $\mathcal{W}_1$, and $\mathcal{W}_2$, based on the value of $F(\mathbf{w})$ associated with each state. Note that $\mathcal{W}_3 = \mathcal{W}_4 = \emptyset$. In fact, for any symmetric irreducible recovery scheme of size n, $\mathcal{W}_m = \emptyset$ if $m > \lfloor n/2 \rfloor$, where $\lfloor x \rfloor$ is the largest integer i that satisfies $i \le x$.

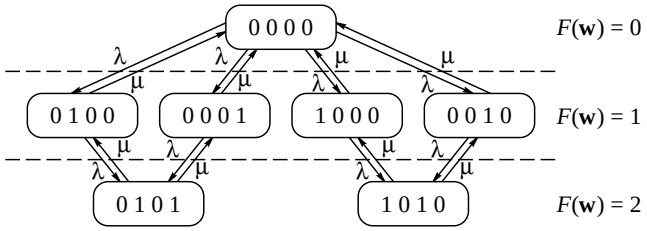

Fig. 2. State space $\mathcal{W}$ for a minimally connected recovery scheme where $n = 4$, showing transition rates and partitioned into subsets $\mathcal{W}_m$ based on the value of $F(\mathbf{w})$.

In order to compute performance metrics for different (1:1)^n recovery schemes, we need to determine the steady state probabilities $p_{\mathbf{w}}$ for all $\mathbf{w} \in \mathcal{W}$. The solution for a (1:1)^n recovery scheme with an arbitrary sharing matrix $\mathbf{S}$ is given by the following theorem.

Theorem 1: Given a recovery scheme with sharing matrix $\mathbf{S}$. For any state $\mathbf{w} \in \mathcal{W}$, if $F(\mathbf{w}) = m$, then $p_{\mathbf{w}} = r^m p_{\mathbf{0}}$, where $r = \lambda/\mu$ and $p_{\mathbf{0}}$ is the steady-state probability of being in state $\mathbf{w} = \mathbf{0}$.

Proof: Because the arrival and departure processes that characterize the system are Poisson, it is not possible to have multiple simultaneous failure or repair events, since these occur in an interval of length Δt with probability $\mathrm{o}(\Delta t)$. For any state $\mathbf{w} \in \mathcal{W}_m$, the only allowable transitions to or from other states involve states $\mathbf{w}'$ whose Hamming distance from $\mathbf{w}$, $d(\mathbf{w},\mathbf{w}') = \sum_{i=1}^{n} |w_i - w_i'|$, is unity. Since $F(\mathbf{w}) = m$, the states to which the system can transition from $\mathbf{w}$ can be grouped into two sets, $\mathcal{A}_{m-1} \subseteq \mathcal{W}_{m-1}$ and $\mathcal{A}_{m+1} \subseteq \mathcal{W}_{m+1}$, where $\mathbf{x} \in \mathcal{A}_{m-1}$ if $F(\mathbf{x}) = m-1$ and $d(\mathbf{w},\mathbf{x}) = 1$, and similarly for $\mathbf{y} \in \mathcal{A}_{m+1}$. $F(\mathbf{w})$ and $U(\mathbf{w})$ are the size of sets $\mathcal{A}_{m-1}$ and $\mathcal{A}_{m+1}$, respectively. The portion of the state transition diagram consisting of the states in $\mathcal{A}_{m-1} \cup \{\mathbf{w}\} \cup \mathcal{A}_{m+1}$ is shown in Fig. 3. Other transitions from the states in $\mathcal{A}_{m-1}$ and $\mathcal{A}_{m+1}$ may exist, but we do not show these for the sake of clarity. The global balance equation for $\mathbf{w}$ can be written as

$$0 = -(U(\mathbf{w})\lambda + F(\mathbf{w})\mu)p_{\mathbf{w}} + \lambda \sum_{\mathbf{x}\in\mathcal{A}_{m-1}} p_{\mathbf{x}} + \mu \sum_{\mathbf{y}\in\mathcal{A}_{m+1}} p_{\mathbf{y}}. \tag{12}$$

This equation can be decomposed into a set of local balance equations as follows, where $r = \lambda/\mu$:

$$p_{\mathbf{w}} = r p_{\mathbf{x}} \quad \forall \mathbf{x} \in \mathcal{A}_{m-1} \tag{13}$$

$$p_{\mathbf{y}} = r p_{\mathbf{w}} \quad \forall \mathbf{y} \in \mathcal{A}_{m+1} \tag{14}$$

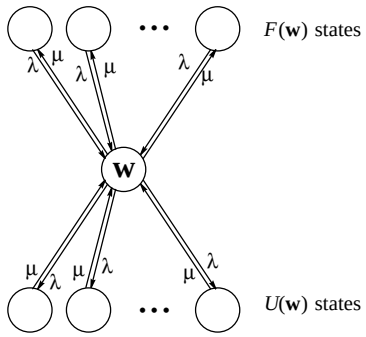

Fig. 3. Probability flux into and out of state $\mathbf{w}$. The states at the top compose set $\mathcal{A}_{m-1}$ while those at the bottom compose set $\mathcal{A}_{m+1}$.

Letting $\mathbf{w} = \mathbf{0}$ produces a set of $U(\mathbf{0}) = n$ local balance equations whose solution is $p_{\mathbf{y}} = r p_{\mathbf{0}}\ \forall \mathbf{y} \in \mathcal{A}_1 = \mathcal{W}_1$. If we assume that there exists a positive integer k such that $p_{\mathbf{w}} = r^m p_{\mathbf{0}}\ \forall \mathbf{w} \in \mathcal{W}_m$, where $m = 0, 1, \ldots, k$, then using the local balance equations for each $\mathbf{w} \in \mathcal{W}_{k+1}$ gives $p_{\mathbf{w}} = r(r^k p_{\mathbf{0}}) = r^{k+1} p_{\mathbf{0}}$. So by induction we have $p_{\mathbf{w}} = r^m p_{\mathbf{0}}$ $\forall \mathbf{w} \in \mathcal{W}_m$. Furthermore, inserting this result into the global balance equation (12),

$$\begin{aligned} 0 &= -(U(\mathbf{w})\frac{\lambda^{m+1}}{\mu^m} + F(\mathbf{w})\frac{\lambda^m}{\mu^{m-1}})p_{\mathbf{0}} \\ &\quad + F(\mathbf{w})\lambda\frac{\lambda^{m-1}}{\mu^{m-1}}p_{\mathbf{0}} + U(\mathbf{w})\mu\frac{\lambda^{m+1}}{\mu^{m+1}}p_{\mathbf{0}}, \end{aligned} \tag{15}$$

produces an expression that holds for all $\mathbf{w} \in \mathcal{W}$. ■

Applying the normalization condition $\sum_{\mathbf{w}\in\mathcal{W}} p_{\mathbf{w}} = 1$, we have

$$1 = p_{\mathbf{0}} + \sum_{m=1}^{n} \sum_{\mathbf{w}\in\mathcal{W}_m} p_{\mathbf{w}}. \tag{16}$$

Using the result of Theorem 1 in (16), we obtain the following expression for $p_{\mathbf{0}}$:

$$p_{\mathbf{0}} = \left(1 + \sum_{m=1}^{n} r^m |\mathcal{W}_m|\right)^{-1}. \tag{17}$$

III. Recovery Blocking Probability Bounds

We define the recovery blocking probability P_B to be the probability that an arbitrary working path failure is blocked due to a lack of backup resources. Consider the i^{th} 1:1 recovery group in a (1:1)^n recovery scheme. The group can be in one of two states, as shown in Fig. 4. The group transitions from state 0 to state 1 only if its backup path is available; otherwise it uses dynamic recovery to reroute the working path. Once the group is in state 1 it waits an average of $1/\mu$ units of time before the working path is repaired and then transitions to state 0. The steady-state probability that the group's working path is using its backup path is

$$P_1 = \frac{\lambda(1-P_B)}{\lambda(1-P_B)+\mu}. \tag{18}$$

The recovery blocking probability P_B is therefore

$$P_B = \frac{\lambda - P_1(\lambda+\mu)}{\lambda(1-P_1)} = \frac{r - P_1(r+1)}{r(1-P_1)}. \tag{19}$$

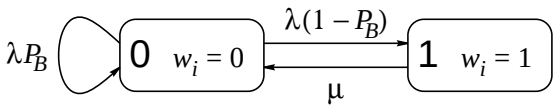

Fig. 4. State diagram for the i^{th} 1:1 recovery group in a $(1{:}1)^n$ recovery scheme.

In order to find P_B, we must first compute P_1. Since the recovery scheme is symmetric, P_1 is the same for all 1:1 recovery groups (WP$_i$,BP$_i$), $1 \leq i \leq n$. P_1 is found by summing the probabilities of all states $\mathbf{w}$ where $w_i = 1$:

$$P_1 = \sum_{\substack{\mathbf{w}\in\mathcal{W} \\ w_i=1}} p_{\mathbf{w}}. \tag{20}$$

We are now in a position to derive upper and lower bounds for the recovery blocking probability, which are respectively associated with maximally and minimally connected symmetric recovery schemes. We examine the maximally connected case first.

A. Maximally Connected Symmetric Recovery Schemes

Recall from our definition that a maximally connected scheme is characterized by a sharing matrix in which $s_{ij} = 1$ for all i and j. In such a scheme, the failure of a single working path will cause the system to enter a state belonging to $\mathcal{W}_1$. Now, $|\mathcal{W}_1| = n$, and $\forall \mathbf{w} \in \mathcal{W}_1$, $F(\mathbf{w}) = 1$ and $B(\mathbf{w}) = n-1$. In other words, if any one working path fails the remaining $n-1$ working paths will be blocked. Thus the state space is $\mathcal{W} = \{\mathbf{0}\} \cup \mathcal{W}_1$, and $\mathcal{W}_m = \emptyset$ for $m > 1$. From Theorem 1, $p_{\mathbf{w}} = r p_{\mathbf{0}}$ for all $\mathbf{w} \in \mathcal{W}_1$. We use (17) and find that $p_{\mathbf{0}}$ is

$$p_{\mathbf{0}} = \frac{\mu}{\mu + n\lambda} = \frac{1}{1+nr}. \tag{21}$$

Hence $P_1 = r/(1+nr)$. Using this fact in (19), we get

$$P_{BU} = \frac{(n-1)r}{1+(n-1)r} \tag{22}$$

as the upper bound for the recovery blocking probability.

B. Minimally Connected Symmetric Recovery Schemes

Next we consider the minimally connected case. We obtain the recovery blocking probability by computing (20) and inserting the result into (19). We have

$$\begin{aligned} P_1 &= \sum_{m=1}^{\lfloor n/2 \rfloor} \sum_{\substack{\mathbf{w}\in\mathcal{W}_m \\ w_i=1}} p_{\mathbf{w}} \\ &= p_{\mathbf{0}} \sum_{m=1}^{\lfloor n/2 \rfloor} r^m |\{\mathbf{w} : \mathbf{w} \in \mathcal{W}_m, \;\; w_i = 1\}|. \end{aligned} \tag{23}$$

In order to evaluate the above expression, we need to compute $p_{\mathbf{0}}$ and $|\{\mathbf{w} : \mathbf{w} \in \mathcal{W}_m, \;\; w_i = 1\}|$. Examining (17), we find that meeting the first objective requires that we determine $|\mathcal{W}_m|$ for $m = 1, 2, \ldots, \lfloor n/2 \rfloor$, which we do in the following theorem.

Theorem 2: For a minimally connected symmetric recovery scheme of size n, the number of elements in the set $\mathcal{W}_m$ is

$$|\mathcal{W}_m| = \frac{n}{m}\binom{n-m-1}{m-1}. \tag{24}$$

Proof: Each element of the set $\mathcal{W}_m$ is a state $\mathbf{w}$ where $F(\mathbf{w}) = m$ and no more than one working path uses a given shared backup path. In the minimally connected case, the degree of each vertex in the sharing graph is 2. Thus, if WP$_i$ has failed, WP$_{i-1}$ and WP$_{i+1}$ (modulo n) cannot be in a failed state since BP$_i$, which is in use, shares resources with BP$_{i-1}$ and BP$_{i+1}$. We can represent the configuration of WP failures associated with each state in $\mathcal{W}_m$ by modeling the sharing graph as a ring of consecutively numbered urns $\{1, 2, \ldots, n\}$. A failed working path WP$_i$ is denoted by placing a single ball into the corresponding urn i. Thus a state $\mathbf{w} \in \mathcal{W}_m$ can be represented by placing m indistinguishable balls into the urns so that there is no more than one ball in any urn and there is at least one empty urn between any two urns with balls in them.

For $m = 1$, there are n possible outcomes (ball in urn 1, ball in urn 2, etc.), so $|\mathcal{W}_1| = n$. For $m \geq 2$, there are n places to put the first ball. Once the ball is placed into an urn, no balls can be put into that urn or the two urns on either side of it. We can therefore remove those three urns from the ring, leaving a chain of $n-3$ urns into which we must put the remaining $m-1$ balls such that the two above conditions are satisfied. The $m-1$ balls, if placed in such a manner, will partition the chain into groups of empty urns. We use the symbol $\boxtimes$ to represent an urn that contains a ball, and we define x_i to be the number of empty urns in the i^{th} group of empty urns in the chain, where the first group is located at the left end of the chain. If a ball is placed in the leftmost urn, then $x_1 = 0$. Likewise, if a ball is placed in the rightmost urn, then $x_m = 0$. Thus we can represent the chain of urns as follows:

$$x_1 \boxtimes x_2 \boxtimes x_3 \boxtimes \cdots \boxtimes x_{m-1} \boxtimes x_m$$

where $x_1 \geq 0$, $x_2 > 0$, $x_3 > 0$, …, $x_{m-1} > 0$, and $x_m \geq 0$.

Now define $y_1 = x_1 + 1$, $y_2 = x_2$, $y_2 = x_3$, …, $y_2 = x_{m-1}$, and $y_m = x_m + 1$. The ball placement problem is isomorphic to determining that number of ways that y_1, y_2, …, y_m can be added to produce $n-m$. It is well known (see for instance [5]) that there are ${}_{N-1}C_{R-1}$ length-R vectors of positive integers whose elements can be added to produce N. Thus there are ${}_{n-m-1}C_{m-1}$ ways that $m-1$ failures can be arranged over a chain of $n-3$ nodes. Since the balls are indistinguishable, we must divide this term by m, since we used one ball to break the circle. We must also multiply this term by n because there are n urns into which we can place the first ball. This gives (24), completing the proof. ■

This leads immediately to the following corollary.

Corollary 3:

$$|\{\mathbf{w} : \; \mathbf{w} \in \mathcal{W}_m, \;\; w_i = 1\}| = \binom{n-m-1}{m-1}. \tag{25}$$

Proof: Let the first ball be dropped into urn i. From the proof of Theorem 2, the number of ways to arrange the remaining $m-1$ balls is ${}_{n-m-1}C_{m-1}$. ■

Using the results of Theorem 2 and Corollary 3 in (23), we obtain

$$P_1 = \frac{\sum_{m=1}^{\lfloor n/2 \rfloor} r^m \binom{n-m-1}{m-1}}{1 + \sum_{m=1}^{\lfloor n/2 \rfloor} \frac{nr^m}{m} \binom{n-m-1}{m-1}}, \tag{26}$$

which, using (19), gives

$$P_{BL} = 1 - \frac{\sum_{m=1}^{\lfloor n/2 \rfloor} r^{m-1} \binom{n-m-1}{m-1}}{1 + \sum_{m=1}^{\lfloor n/2 \rfloor} \frac{(n-m)r^m}{m} \binom{n-m-1}{m-1}} \tag{27}$$

as the lower bound on the recovery blocking probability.

In Fig. 5 we plot values of P_{BU} and P_{BL}. The upper limit on r corresponds to a mean repair time that is 1% of the mean time between failures, e.g. a mean repair time of 4 hours while a failure occurs every 17 days, on average. All of the lower bound curves nearly overlap, so that they are not distinguishable in the figure. If r is small, then a high degree of sharing is possible; if we impose a constraint that $P_B < 0.01$, we can use maximal sharing for $n < 8$ if $r < 0.0015$. Minimal sharing works for $r < 0.005$ for all n; this suggests that we can chain together large numbers of backup paths without penalty as long as the probability is small that a WP's neighbors in the sharing graph fail while the WP is failed.

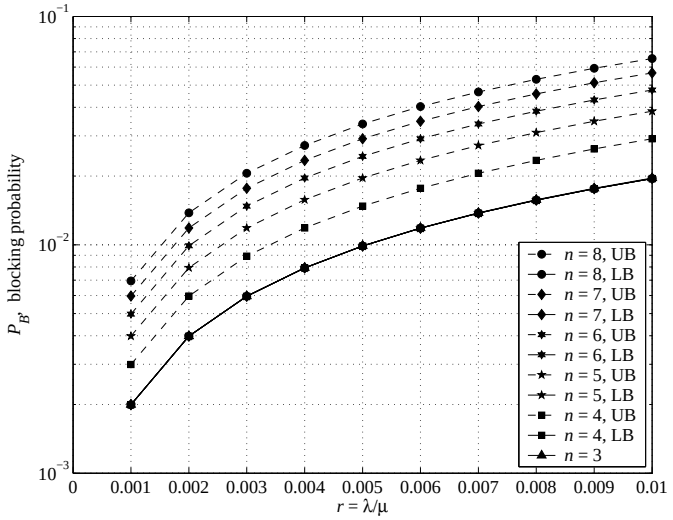

Fig. 5. Theoretical upper and lower bounds on P_B for $n = 3, 4, \ldots, 8$.

IV. Simulation Results and Analysis

We used the GLASS [6] simulation tool to perform discrete event simulations for the case of minimally and maximally connected symmetric recovery schemes for values of $n = 3, 4, \ldots, 8$. In each case we simulated the behavior of a single recovery scheme of size n; because the current practice is to use dedicated backup resources for each new working path, it is unlikely that $(1{:}1)^n$ schemes will incorporate more than eight shared backup paths in the near future. Each of the n working paths was independently transitioned between failed and healthy states with exponentially distributed durations of $1/\mu = 4$ hours and 400 hours $\leq 1/\lambda \leq 4000$ hours, respectively. We computed the recovery blocking probability by taking the ratio of the number of blocked failures to the total number of failures (50,000 and 250,000 for the upper and lower bounds, respectively) in each simulation run. In Fig. 6, we plot the recovery blocking probabilities averaged over three simulation runs. The error bars in the figure indicate one standard deviation in the data. 40% of the data deviated from the theoretical results in Fig. 5 by less than 1%; 10% deviated by more than 3.1% and the greatest error was 5.7%.

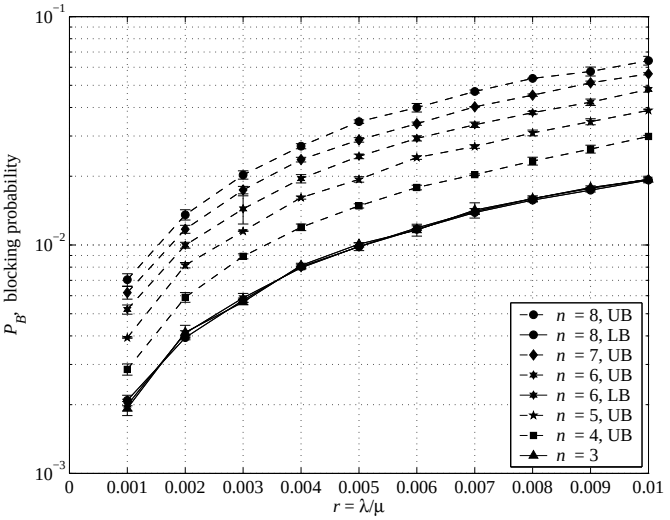

Fig. 6. Upper and lower recovery blocking probability bounds obtained from simulation for $n = 3, 4, \ldots, 8$.

V. Summary

We have presented analytical models for $(1{:}1)^n$ shared protection schemes, considering both maximally and minimally shared schemes. We have validated the results of the analytical models with the help of network simulation results. The analysis and results reported in this paper are useful for quantifying the trade-offs between the resource sharing benefits vs. performance in terms of protection blocking. Network operators can use these results to determine the level of resource efficiency they can achieve for a specified connection availability (or robustness) requirement.

References

[1] "Recovery (Protection and Restoration) Terminology for GMPLS," E. Mannie and D. Papadimitriou, Eds., IETF Internet draft.

[2] "Analysis of Generalized MPLS-based Recovery Mechanisms (including Protection and Restoration)," D. Papadimitriou and E. Mannie, Eds., IETF Internet draft.

[3] K. Sriram, D. Griffith, S. Lee, and N. Golmie, "Backup Resource Pooling in (M:N)^n Fault Recovery Schemes in GMPLS Optical Networks," in Proceedings of SPIE Vol. 5285 *OptiComm 2003: Optical Networking and Communications*, edited by Arun K. Somani, Zhensheng Zhang, (SPIE, Bellingham, WA, 2003), pp. 185–196.

[4] N. J. A. Sloane and S. Plouffe, *The Encyclopedia of Integer Sequences*, Academic Press, San Diego, 1995.

[5] S. Ross, *A First Course in Probability*, Macmillan Publishing Company, New York, 1988.

[6] Oliver Borchert, Richard Rouil, "The GMPLS Lightwave Agile Switching Simulator - An overview," www.antd.nist.gov/glass