

A Distributed Algorithm for Path Restoration in Circuit Switched Communication Networks*

S.Venkatesan, Maulin Patel and Neeraj Mittal

Department of Computer Science and Telecommunication Engineering Program
The University of Texas at Dallas, Richardson, Texas 75083–0688
venky@utdallas.edu maulin@student.utdallas.edu neerajm@utdallas.edu

Abstract

The multi-commodity network flow (MCNF) problem is an important problem with many applications in the areas of routing and telecommunications. While many centralized algorithms have been proposed for solving the MCNF problem, distributed algorithms have received very little attention. This paper presents an online distributed multi-commodity flow approximation algorithm specifically tailored for path restoration. Path restoration is an important approach for building survivable telecommunication backbone networks. Path-based restoration schemes are known their for high restoration efficiencies and their ability to protect against single link, multiple link and node failures. Our algorithm uses $O(|E|diam^2)$ messages and $O(diam^2)$ time in the worst case, and substantially fewer messages and less time in practical networks. When simulated on a sample real-life backbone network similar to those used by the telecommunication service providers, our algorithm finds a solution significantly faster than many published algorithms.

1 Introduction

The multi-commodity network flow (MCNF) problem [2] consists of two or more commodities that have to be sent from their respective sources (or origins) to their respective destinations through a common network using one or more paths. There is a demand associated with every commodity and a capacity associated with every link. The total flow on each link should not exceed its capacity. The MCNF problem has applications in a wide variety of areas including production planning, warehousing, transportation and telecommunication networks [3]. Not surprisingly, it is a well-studied problem and several centralized algorithms have been proposed by many researchers to solve it and its

variants [5, 15, 2]. However, some of the problems in the areas of routing and telecommunication systems, which can be expressed as variants of the MCNF problem, require distributed algorithms. An important example includes traffic restoration in backbone networks.

Online Restoration Problem: Online restoration of disrupted traffic in the event of a failure is a prime concern in self healing (fault-tolerant) mesh telecommunication networks. To ensure that the disruption caused by link or node failures is minimized, efficient algorithms are needed for *real-time* restoration of the disrupted traffic. Many different restoration techniques have been proposed in the literature [12, 8, 16, 9, 4, 14, 22, 17, 11, 19]. We focus on path restoration for mesh-type self-healing networks with shared spare capacity.

The backbone network consists of a set of nodes connected by bidirectional point to point links. Superimposed on this physical network are a set of n paths, each path starting from one source, going through the links of the network, and reaching the corresponding destination. Note that some of the paths may share the same source node and some of the paths may share the same destination node. Each path P_i also has a certain traffic flow requirement or capacities, say, c_i . For a path P_i to go over a link l , link l dedicates c_i units of link capacity to path P_i . Thus, link l dedicates capacities equal to the sum of the capacities of all of the paths using l . This capacity is called working capacity of l and denotes the amount of live traffic carried by l . In addition to the working capacity, a link has spare capacity. The spare capacity of a link does not carry any live traffic.

When a link fails, all the paths using that link are disrupted. In order to survive this link failure, the disrupted paths will have to be rerouted so that they use other links. A path can be rerouted on a link, only when that link has sufficient spares to carry the rerouted traffic (in addition to its live traffic). Spare capacity allocation is done off-line (before failures) during network planning time. The restoration algorithm is run at the time of failures to find alternate paths.

*This work was supported in part by two grants from Alcatel Network Systems.

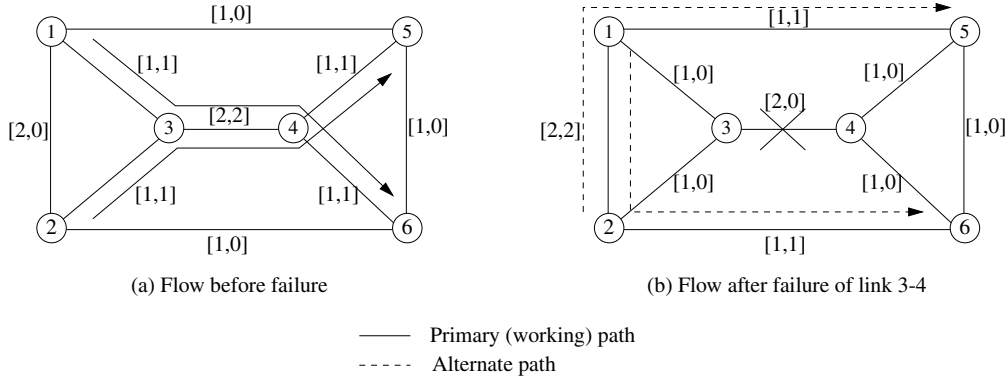


Figure 1. Path Restoration Example.

Path Restoration: Path restoration [12, 9, 4, 14, 11, 22] reroutes each disrupted path independently over one or more alternate paths between the source and the destination of the primary path, using the spare capacity of the network. Therefore determining restoration paths using path restoration scheme can be expressed as a multi-commodity flow problem. Path restoration is known for its high restoration efficiencies and its ability to protect against many kinds of failures including single link, multiple link and node failures. However, the MCNF problem is NP-complete if the flow values are integers [10].

For example, consider the small network consisting of six nodes shown in Figure 1(a). For each link, two numbers are shown; the first number denotes the total capacity (working capacity plus spare capacity) and the second number denotes the working capacity. Link (3,4) is used by two paths: path P_1 carrying 1 unit of flow from node 1 to node 6 via nodes 3 and 4; path P_2 carrying 1 unit of flow from node 2 to node 5 via nodes 3 and 4. Only links (1,3), (2,3), (3,4), (4,5), (4,6) carry live traffic. If link (3,4) fails, both P_1 and P_2 are disrupted. The rerouted path for P_1 is (1,2,6) and for P_2 is (2,1,5). Note that both rerouted paths go over links without exceeding spare capacities available in any of the links.

Classification of Restoration Schemes: The restoration schemes for mesh networks can be classified by their computation timing, by their execution mechanisms and by the type of rerouting used [9]. The real time approach [8, 16] finds alternate paths after the failure has been detected, while the preplanned approach [9, 4, 14] precomputes the alternate paths for the given failure scenarios. Real-time approach is suitable when traffic patterns change very frequently. Preplanned approach is faster than real-time approach but suffers from poorer capacity utilization than real-time approach. There are two types of execution mechanisms: Distributed [12, 8, 16, 9] and centralized [4]. The centralized schemes perform the computation at a central

site, where accurate information about the current network state is assumed to be available. Centralized schemes are capacity efficient but have single point of failure, communication overhead and scalability issues.

To minimize data loss and call dropping probability, the restoration should be completed within 2 seconds of failure occurrence and preferably within 1 second as per Bellcore's advisory. [1].

Our Contribution: We propose a new real-time distributed control restoration scheme that works well in real-life backbone networks similar to those used by telecommunication service providers. Simulation results on realistic backbone topologies show that our scheme has the potential to achieve target restoration speed.

2 System Model

The backbone network consists of nodes and bidirectional communication links. In the network, a path starts from a source node, ends at a destination node and consists of a series of neighboring nodes along with the connecting links. (Two nodes are neighbors if they are connected by a single bidirectional link.) For each link, we are given the amount of live traffic and spare capacity available.

2.1 Knowledge of Nodes Before Failures

Each node has local topological information, in addition to the information obtained in the preprocessing step (described later). A node knows (i) the number of links incident on it; (ii) For each link incident on that node, it knows the (a) the id of the node at the other end, (b) working and spare capacities of that link, (c) status of the link—down or up, (d) number of paths going through that link, and (e) path information for the link which includes, for each path going through that link, the identity of the path (id of the source of the path, id of the destination of the path and the capacity of the path). This knowledge is available before any link fails. Also, no node knows the entire topology of the network.

Traditionally, only single link failures have been considered since the probability of multiple simultaneous failures is very small. Thus, when the telecommunication companies design the spare capacity network, only single link failures are assumed to occur. When a link fails, the disrupted traffic is restored on alternate paths. Immediately, steps are initiated to repair the link that failed, and, when the fault is rectified, the disrupted traffic (that was restored) reverts back to the original paths (using the repaired link). Thus alternate paths are used for a short time (few minutes to about half a day in most cases). Our algorithm can cope with (and recover from) multiple link and node failures. An underlying failure detection mechanism exists to detect the failure of a link.

3 Our Algorithm

3.1 Failure Detection and Failure Notification

Failure notification occurs after failure detection. Immediately after detecting the link failure, the two end nodes of the faulty link assemble a *failed* message and send it on all of the outgoing links. The *failed* message contains the number of paths going through the faulty link and the path id of the paths going through the faulty link. For each path, path id consists of the source node id, destination node id, and the amount of traffic going on that path.

When an arbitrary node u receives the *failed* message, it checks if it has already received this message. If so, it discards the message. Otherwise, it records the *failed* message locally and forwards the message (once) on all links except on the link on which it received the *failed* message.

It is possible to release the capacities allocated to paths that get disrupted. These released capacities can be added to the spare capacity on the links.

The two end nodes of the faulty link start the algorithm immediately after detecting the link failures and sending out the *failed* message. Other nodes begin the algorithm immediately after receiving the first *failed* message.

3.2 Maintaining the Nodes in Rhythm

Our algorithm requires that the nodes be kept in “rhythm.” Each node maintains a step number. The step numbers have the following property:

A message m sent by node u to neighboring node v when u 's step number is i will be received by v when v 's step number is $i + 1$. (Note that u and v are directly connected by a link.) One of the three synchronizers α, β, γ proposed by Awerbuch [6] can be used. We use a synchronizer that is simpler and faster than these three synchronizers. We ensure that each node sends exactly one message on each link. Thus, it is easy for a node to check if it can proceed

to the next step: as soon as it receives one message on each link for the current step, it can proceed to the next step.

The following two rules are used to keep the nodes in rhythm: (1) During step i of the restoration algorithm, if u has a message to send to its neighbor v , then u sends it. Otherwise (if u has no message to send to v during step i), u creates a *step_completed* message (which is similar to a null message) and sends the *step_completed* message to v . This step is performed for each neighbor of u . (2) After u receives one message from each neighbor for the current step, node u processes them and proceeds to the next step.

The cumulative step numbers are maintained locally only and are not sent in any of the messages.

3.3 High Level Description of Algorithm

The restoration algorithm proceeds in *iterations*. Each iteration consists of two phases—an explore phase and a return phase. The explore phase is started by the source nodes of the disrupted paths. The return phase is started by the destination nodes of the disrupted paths.

Consider the i^{th} iteration. (Initially, $i = 1$.) Each phase (explore phase and return phase) of the i^{th} iteration consists of exactly $i + 1$ steps. Thus, the i^{th} iteration consists of $2 \times (i + 1)$ steps.

A Step: During each step (except the first step of the explore phase and the first step of the return phase), a node waits for a message (*explore* or a *step_completed* message in the case of *explore* phase and *return* or *step_completed* message in the case of the return phase) from each neighbor, processes the received messages and sends one message (*explore* or a *step_completed* message in the case of *explore* phase and *return* or *step_completed* message in the case of the return phase) to each neighbor. Each source S_j finds the subnetwork that consists of all the nodes (and the links) that are reachable from S_j within $i + 1$ hops. Using this subnetwork only, S_j locally finds the maximum amount of traffic that can be restored between S_j and D_j and restores the maximum possible traffic. If all of the nodes are in rhythm, then it is easy to see that by the end of the i^{th} iteration, all paths of hopcount less than or equal to $i + 1$ from S_1 to D_1 , from S_2 to D_2 , ..., S_n to D_n are considered. We next describe how the subnetwork is obtained at each source. Determining the maximum amount of traffic on the subnetwork is described next.

Explore Phase: The i^{th} iteration starts with the explore phase. The source nodes send *explore* messages. The *explore* messages traverse links looking for the correct destination nodes. To control the flow of *explore* messages, a hop count is given. For the i^{th} iteration, the hop count is set to $i + 1$ at the source nodes. When an *explore* message traverses a link, the hop count field is decremented. When hop

count reaches zero, the *explore* message is not propagated further. When an *explore* message reaches the intended destination node, the *explore* messages are not propagated further. Instead, we wait till the end of the explore phase and then start the return phase.

Return Phase: During the return phase, the destination nodes of the disrupted paths assemble a *return* message if an *explore* message sent by the corresponding source node has reached the destination node during the explore phase of the present iteration. (All other nodes send *step_completed* messages.) The *return* messages are sent by the destination nodes to their neighbors and are propagated. The *return* messages also acquire spare capacities on each link traversed by them. Contention for spares of the links is resolved in this phase. (Details in § 3.4.4.) When a *return* message reaches the source node, the received information is locally stored. Recall that the *return* message has both the source and destination ids in its field. The *return* message contains the following information: (1) source id, (2) destination id, (3) amount of traffic to be restored, and (4) information about the subnetwork traversed by the *return* messages.

At the end of the return phase (at the end of the $i + 1^{st}$ step of the return phase), each source node locally has the subnetwork that consists of the nodes and the links (along with the spares acquired for restoring the disrupted traffic) that have been explored by the *explore* messages. Using the subnetwork, each source node locally runs the max flow algorithm and determines the maximum amount of traffic possible between the source and the destination using the acquired spare capacities as edge capacities. Any of the (sequential) max flow algorithms can be used. If the computed maximum traffic is less than the original amount of traffic carried by the path, then more iterations are needed and the source node proceeds to the next iteration. Otherwise, the source node knows that one path has been restored and the source node sends a *path_restored* message. It also releases the capacities on links that have been allocated for restoring this path but are not needed.

When a *path_restored* message is sent by all of the n source nodes (where n is the number of disrupted paths), the nodes receiving these n distinct *path_restored* messages terminate the restoration algorithm.

3.4 Algorithm Details

3.4.1 Algorithm Initiation Phase

The algorithm starts with the algorithm initiation phase. This phase is started when a link failure is detected. The *failed* message is assembled and broadcast to all of the nodes.

3.4.2 Explore Phase

Actions of node v at the beginning of an iteration

As noted earlier, the i^{th} iteration consists of exactly $2 \times (i + 1)$ steps. Also, the i^{th} iteration begins at a node only when the previous iteration ends. Thus, a node can locally determine when an iteration begins and when an iteration ends.

Consider an arbitrary node v . If v is a source node of some path P_j and if P_j has not been fully restored, then it begins the explore phase of the i^{th} iteration (after it ends the $i - 1^{st}$ iteration). At this time, node v assembles an *explore* message. The hop count field of the *explore* message is equal to $i + 1$. The field “amount of traffic to be restored” is initialized to the difference between the path’s original traffic and the amount of traffic that has been restored for that path in all of the iterations so far.

If v is the source of more than one disrupted path, then an *explore* message is assembled for each path for whom the source is v .

If node v is not a source node of any of the disrupted paths, then node v assembles a *step_completed* message locally.

After assembling either *explore* message(s) or a *step_completed* message, node v sends the assembled message(s) on all of its outgoing links. Node v then waits for a message (*explore* or a *step_completed* message) from each neighbor. Note that the message (that v is waiting for) may be sent by the neighbor either autonomously or in response to the message v had sent earlier. If node v has assembled several *explore* messages (because it is the source of several disrupted paths), then all of those messages are bundled and sent as a single (long) message.

Actions of node v at other times

If it is not the beginning of an iteration, then node v waits for a message from each neighbor and processes these messages. (*step_completed* messages are discarded after receiving them.)

Let $\{m_1, m_2, \dots, m_x\}$ be the *explore* messages received by v . First node v “remembers” on which link it received the *explore* messages. Consider message m_j ($1 \leq j \leq x$) received by node v sent by node u on the link $u \rightarrow v$. Let $m_j = \text{explore}(S_{j_1}, D_{j_1}, c_{j_1}, h_{j_1})$. It first decrements h_{j_1} . Node v locally stores the entry $\{\text{sender}=u, \text{source}=S_{j_1}, \text{destination}=D_{j_1}, \text{traffic to be restored} = c_{j_1}, \text{and decremented hop count} = h_{j_1}\}$. This entry signifies that v has received an *explore* message from u with parameters $S_{j_1}, D_{j_1}, c_{j_1}$, and h_{j_1} (h_{j_1} is the decremented hop count). If v has already received an *explore* message with the same source and destination ids in the same iteration (but at an earlier step), then m_j is not propagated further. If m_j and m_k have the same source and destination identities (for

D_i	h_i	Action
$\neq v$	$= 0$	Ignore m_i
$\neq v$	> 0	Store m_i locally and propagate m_i if necessary
$= v$	$= 0$	Assemble a <i>return</i> message and send it to u
$= v$	> 0	wait for $2 \times h_i$ “steps”; assemble a <i>return</i> message and sent it to u

Figure 2. Summary of actions needed to process an *explore* message m_i .




Figure 3. Node v at step 5 (iteration 2).

some $1 \leq k \leq x$), then only one of them will be forwarded. Also, if the destination D_{j_1} is equal to v then the *explore* message is not forwarded.

To summarize, node v processes message m_j as follows: (Note that h_{j_1} is the decremented hop count in the following description.)

- If $h_{j_1} > 0$, $v \neq D_{j_1}$ and v has not received an *explore* message with the same source and destination id earlier in the same iteration, then v forwards the *explore* message after decrementing hopcount field and storing the *explore* message locally.
- If m_j and m_k have the same source and destination identities, then at most one of them will be forwarded.
- If $v = D_{j_1}$, then the *explore* message is not forwarded and v assembles a *return* message. If (decremented) hop count $h_{j_1} = 0$ then a *return* message is generated and sent back. If $h_{j_1} > 0$, then we do not start the return phase immediately. Instead, we wait for $2 \times h_{j_1}$ “steps.” (Waiting is necessary to make sure that if some other path is contending for a link that will be traversed by this *explore* message, then the contention must be known before assigning spares to the path $S_{j_1} \rightarrow D_{j_1}$.) This is equivalent to the *explore* message traversing an imaginary path of length h_{j_1} links. Node v




Figure 4. Node v at step 6 (iteration 2).

records the information $\{\text{pending}, \text{explore}(S_{j_1}, D_{j_1} = v, c_{j_1}, h_{j_1}), \text{currentstepnumber} + h_{j_1}, u\}$.

Forwarding an explore message

Node v , in response to the *explore*($S_{j_1}, D_{j_1}, c_{j_1}, h_{j_1}$) message, checks if $h_{j_1} > 0$ and this *explore* messages is to be forwarded. If so, it sends an *explore*($S_{j_1}, D_{j_1}, c_{j_1}, h_{j_1}$) message to all neighbors from which it has not received an *explore* message with the same source and destination id. All of the messages are processed in this manner.

If v does not send an *explore* message on a link, it sends a *step_completed* message on that link. The actions are summarized in Figure 2.

An Example: Consider an arbitrary node v of a network as shown in Figure 3. Assume that it is the second iteration now and one step has elapsed. Since this is the second iteration, we are trying to find paths of hop count 3 or less. Note that all paths of hop count 3 or less include paths of hop count 2. Those paths are included again since some spare capacities that were allocated to the other paths may have been released if those spares were not used by the other paths. After one step, the hop count field of the *explore* message is 2. Let nodes a, b, c, d, e, f and g be the 7 neighbors of node v . Assume that nodes e, f and g are not neighbors of any of the sources and the other nodes (a, b, c and d) are neighbors of one of the sources. Let $m_1 =$

$explore(w, x, 5, 2)$ be the message sent by nodes a and b to v . Assume that c sends $explore(p, q, 2, 2)$ to v and d sends the message $\{explore(w, x, 5, 2), explore(p, q, 2, 2)\}$ to v . The other nodes e, f and g send a *step_completed* message to v . Node v stores $\{explore(w, x, 5, 2), explore(p, q, 2, 2)\}$ locally. Node v decrements the hop count of all of the *explore* messages. The *explore* messages sent by a and b (m_1 and m_2) will be sent to all neighbors except a, b and d . (Recall that d had sent m_4 to v and $m_4 = \{m_1, m_3\}$.) The other *explore* messages are propagated similarly and the results are shown in Figure 4.

3.4.3 Return Phase

The return phase starts after the explore phase ends. In the return phase, the paths traversed by the *explore* messages are retraced by the *return* messages¹. The return phase ends when the *return* messages reach the sources. In the return phase, the contention for spare capacity is resolved.

Origination of the return messages:

When node v receives an $explore(S_i, D_i, c_i, h_i)$ message from node u , it first decrements h_i and then checks if $D_i = v$. If $D_i = v$ (the *explore* message has reached the destination), then it is not forwarded, but a *return* message is generated.

Case (i): $D_i = v$ and $h_i = 0$:

The *explore* message has reached the destination node from the source S_i . Thus, node v immediately generates a *return* message and sends it to node u . The *return* message contains the source id, the destination id, amount of traffic to be restored, and the subnetwork field. Initially, the subnetwork field is empty.

Case (ii): $D_i = v$ and $h_i > 0$.

Node v records the information $\{\text{pending}, explore(S_i, D_i, c_i, h_i), 2 \times h_i + c, u\}$ locally.

Let $count$ be the value of the local counter of node v . Node v checks if it has an entry $\{\text{pending}, explore(S_i, D_i, c_i, h_i), count, y\}$. If so, then a *return* message is generated and sent to node y . The *return* message sent by node v to node u is $return(S_i, v, c_i, \phi)$. This is similar to case (i) discussed above.

Processing return messages:

Let $m = return(S_i, D_i, c_i, L)$ be the message received by node v from node u on the link (u, v) . L is the set of links and their spare capacities. L will be used in restoring traffic between S_i and D_i . Spare capacity is given to the links using the contention resolution steps described in § 3.4.4.

Node v receives a *return* message from u only if v had sent an *explore* message to u earlier in the same iteration. Clearly, for node v to send an *explore* message, (a) it must have either received an *explore* message from one of its

¹using the information stored by the nodes when they received the *explore* messages

other neighbors earlier or (b) v is the source of the disrupted path $P_i = (S_i, D_i)$ for some P_i . If case (a) applies, then node id S_i found in the *return* message is not equal to v . In such a case, node v forwards the *return* message. For the $return(S_i, D_i, c_i, L)$ message, node v checks its locally recorded information. It must have received an *explore* message with the same source and destination identities. Consider only those *explore* messages received and recorded by node v such that the source and destination ids are equal to the source and destination ids of the *return* message received. Among these *explore* messages, let $m' = explore(S_i, D_i, c_i, h_i)$ be the *explore* message, recorded locally at v , with the maximum hop count h_i . Assume that the sender of m' is w . This message (m') is the earliest *explore* message received by v with source S_i and destination D_i . Now, c' , the spare capacity of the link (u, v) that can be assigned to the path $P_i = (S_i, D_i)$, is found by the contention resolution rule and assigned. After this step, node v sends the message $return(S_i, D_i, c_i, L \cup \{l = (u, v), c'\})$ to its neighbor w . (Recall that w is the node that sent message m' to v earlier.) For each neighbor $x \neq w$ from which v had received an *explore* message with the same source and destination ids, node v sends a $return(S_i, D_i, c_i, \phi)$ message to neighbor x . The subnetwork information is sent in one *return* message only and is not duplicated on every *return* message. If case (b) is applicable, then v is the source of the path $P_i = (S_i, D_i)$. (That is, $v = S_i$.) Now, node v locally saves the list of links and spares on those links that are part of the *return* message received. During the last “step” of the current iteration, the max flow algorithm is run by the source node.

3.4.4 Resolving Contention

Often, there will be contention by several disrupted paths for the available spare capacity of a link. Resolving contention is important. Contention resolution is performed during the return phase.

Consider an arbitrary node v and its neighbor u . Assume that v receives a *return* message on the link (u, v) sent by node u . Let sp be the amount of spares currently available in the link (u, v) .

Assigning spare capacity of link (u, v) to path $P_i = (S_i, D_i)$: When a $return(S_i, D_i, c_i, L)$ message is received by v on the link (u, v) , node v determines how much of the spare capacity sp of link (u, v) is to be assigned for restoring P_i . Let m_1, m_2, \dots, m_a be the *explore* messages, which correspond to the paths P_1, P_2, \dots, P_a with unique source-destination ids, received by v during the current iteration. Clearly, for every pair m_j, m_k , $1 \leq j, k \leq a$ and $j \neq k$, m_j and m_k are *explore* messages with sources S_j and S_k and destinations D_j and D_k such that (i) $S_j \neq S_k$ or (ii) $D_j \neq D_k$ or (iii) $S_j \neq S_k$ and $D_j \neq D_k$. Now, spare capacity sp of link (u, v) is uniformly distributed among the




Figure 5. A Sample Subnetwork.

contending paths in proportional to their demands.

3.4.5 Max Flow Algorithm

The source node stores the subnetwork in the form of a series of links along with their spares. Next we use a max flow algorithm. Any of the algorithms from the literature for the sequential model (centralized) is sufficient.

If the total flow obtained from the max flow algorithm in the current and previous iterations is sufficient to restore the traffic in that disrupted path, then the source sends a *path_restored* to all of the nodes by a broadcast. In that case, the source does not participate in the exploration phase of any higher hop counts for that path id. If the traffic needs of the disrupted path are not completely met at the end of this iteration, the source proceeds to the next iteration (with next higher hop count).

3.4.6 Terminating the Algorithm

When a path is fully restored by the source node, the source node sends a *path_restored* message (in a manner similar to *failed* messages). This message is sent to all the nodes (including itself). When a node receives this *path_restored* message for the first time, it decrements its variable NUM_PATHS. When NUM_PATHS reaches the value 0, the node terminates the restoration algorithm. Also, note that when the number of iterations is equal to $O(diam)$ where *diam* is the diameter of the network, each source would have explored the complete topology and all nodes terminate the algorithm.

3.5 Algorithm Novelties

This algorithm has some novelties.

1. We keep the nodes in “rhythm.” It may appear that keeping the nodes in “rhythm” slows the algorithm. For example, as per § 3.2, a node waits till it receives a message from each neighbor. Thus, if one neighbor is slow to start the algorithm or one of the links is slow, then we may be slowing down the node. However, keeping nodes in “rhythm” ensures that nodes carefully select paths for restoration. For example, consider the sample network shown in Figure 5.

Assume that link (a, d) , carrying 16 units of a single path $P = (a, d)$ of traffic, fails. If the restoration path chosen is (a, b, f, d) , then at most 8 units of traffic can be restored. This situation may occur during the execution of many of the published algorithms. In our algorithm, during the second iteration, all paths of hop count 3 between a and d are found. For this example, all the three paths, path (a, b, c, d) , (a, e, f, d) , and (a, b, f, d) are examined by the node a at the same time. By using the max flow algorithm, we will choose the paths (a, b, c, d) and (a, e, f, d) instead of the path (a, b, f, d) . Note that since our algorithm is a heuristic, this may not always result in fast or complete restoration, but this has been very useful in practical networks.

2. Contention resolution is handled in a uniform way in our algorithm. In many of the published algorithms, the first explore message traversing a link may preemptively use all of the spare capacities leaving nothing for the explore messages of the other paths. At a later time, if the first explore message does not use all of the spare capacities obtained, it will release them so that the released spares may be used by others for restoring some other path. However, this process of reserving and then releasing later on consumes time. Also, this process may happen in a cascaded manner. Our algorithm avoids this problem to a large extent by resolving contention in a uniform way.

As a downside of resolving contention in this manner, our path restoration algorithm may produce flows that are fragmented in the sense that a flow is distributed over many paths. This is not an issue, however, because, in practice, the failed link is repaired by the service provider in a few hours. After that fault has been repaired, flows revert back to their original (non-fragmented) paths.

3.6 Preprocessing Steps

The algorithm can be improved by some one-time preprocessing. We suggest the following preprocessing:

Consider node with id, say u . It has an array, called min-hops: $minhops[i]$ is the number of hops (number of links) of the shortest path from itself (node u) to the node with id i . Clearly, this array can be constructed at each node before the restoration algorithm starts. For example, the algorithm of [20] can be used. This is done before failures.

This array is used in deciding when to propagate *explore* messages. Let the node with id u receive an *explore* message with source S and destination D . Node u , on receiving this *explore* message, checks if $minhops[D]$ is greater than the hop count allowed (this number is part of the *explore* message). If $minhops[D]$ is less than or equal to the hop

count allowed, then the *explore* message is propagated (provided node u has not already propagated an *explore* message with the same source and destination ids in the current iteration). On the other hand, if $\text{minhops}[D]$ is greater than the hop count allowed, then it is clear that the *explore* message cannot reach the destination within the hop count allowed. In such a case, the *explore* message is not propagated but is discarded.

Although the *minhops* array of a node is determined before the link failure, the link failure does not reduce the minimum number of hops needed between any two nodes; instead, the link failure may increase the minimum number of hops. Thus, when a node decided to drop an *explore* message, the node will not be making a mistake. Some of the *explore* messages may be sent unnecessarily since the minimum number of hops may have increased because of the link failure, but this does not affect the correctness of the algorithm. An alternate method is to compute the minimum number of hops from each node to all of the other nodes at run time (immediately after the link failure). Note that in the first iteration, knowledge about 2-hop neighbors is needed (and knowledge about nodes that are more than 2 hops away is not needed in the first iteration). To gain this knowledge, each node sends its list of 1-hop neighbors on all links in the beginning. Subsequently, k -hop neighborhood information ($k \geq 2$) can be sent to all neighbors (by each node) by piggybacking this information on the *explore/return/step_completed* messages. Thus, there is a delay of one step, but the *minhops* array will have the accurate value.

3.7 Complexity Analysis

When the current iteration count is equal to $diam$, where $diam$ is the diameter of the network, an *explore* message sent by a source node reaches all the nodes by the end of the explore phase. Therefore, $O(diam)$ iterations are sufficient to explore the entire network. Iteration i consists of $2 * (i + 1)$ steps. During each phase, exactly one message is sent on each link (in each direction). The number of message sent per step is $O(|E|)$ where $|E|$ is the number of links in the network. The total number of messages sent is $O(diam^2|E|)$. Note that if the *minhops* array is computed on the fly (and not as a preprocessing step), one additional step is needed. The time complexity is $O(diam^2)$ since each step takes one time unit.

4 Simulations

The performance of the algorithm was tested through simulation. Simulations were done on the network specified in Bellcore advisory [1], which represents a reasonable metropolitan size network of 15 nodes and 28 links. The advisory also reports the simulated performance of three on-line restoration algorithms and we compare the performance




Figure 6. The test network.

of our algorithm with the performance of three algorithms reported in [1]. The working and spare capacities were designed in [1] to guarantee 100% restoration for any single link failure. The performance figures are obtained by simulating the failure of each of the 28 links, one failure at a time.

The percentage of traffic restored (PTR) relative to the traffic disrupted for each link failure event is used as a performance metric. Another performance metric measured is the restoration time (RT) which is the maximum time required to restore all the disrupted paths completely from the time the algorithm is initiated.

4.1 Simulation Parameters

The simulation methodology accurately measures the local time at each node and hence it permits the calculation of the total running time of the algorithm. To simulate the real performance of the network, we need to add time delays at various points. To compute these delays, the following parameters are used: (a) Internal data transmission rate for the links is set to 64 Kbits/sec, (b) Data transfer rate between the I/O port and CPU is 128 Kbits/sec, (c) The CPU is assumed to have a performance of 1 Million Instructions Per Second.

The test network of [1] is shown in Figure 6.

4.2 Results

Figure 7 shows four plots of PTR versus RT on the test network. Plot 1 represents the simulation results of our algorithm. The result shown is the average value assuming that each link failure is equally likely. Plots 2, 3 and 4 are from [1] which represent the simulation results of the distributed restoration algorithms of [16, 21, 18]. The input parameters used in our simulations are identical to the input parameters used by Bellcore in their simulations [1].




Figure 7. Relative performance of the algorithm.

The best, average and worst case time taken for restoration by our algorithm is shown in the Figure 8.

Based on simulation, our restoration algorithm is fast and achieves 100% restoration. Figure 7 shows that improvement in the restoration time is significant.

Although only single link failures were simulated and results are shown in the paper, multiple link failures and node failures can also be handled with minimum amount of changes. For double link failures, for example, two *failed* messages will be broadcast. Node failures can be treated as the failure of all of the links incident on that node. For node failures, the paths for which either the source or the destination is the faulty node will not be restored. The other paths will be restored. No other changes are needed.

5 Comparison with Related Work

Our work is most closely related to the work of [12, 16, 21]. A main characteristic of all of these algorithms is the following: Each affected source node u , on learning about the disruption of the path for which u is the source node, constructs a packet containing the source id, destination id, amount of traffic to be restored, and a hop count, and floods the network with this packet. When an arbitrary node receives such a packet, it sets the amount traffic to be restored (on the packet) to the smaller of the value of the spare capacity of the link on which this message was received and the amount of flow to be restored (available in the packet itself) assuming that all the spare capacity is available for restoring this path. The recipient decrements the hop count and adds its own id to the packet (so that the packet contains the path traversed so far). When a packet reaches the intended destination, a path from the source to the corre-

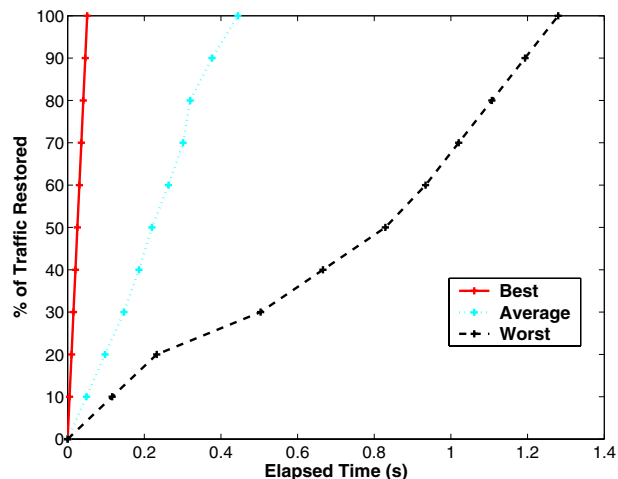


Figure 8. The best, average and worst case restoration time.

sponding destination has been found. Details vary from one algorithm to the other in terms of when the spare capacities are allocated to a path; some algorithms allocate in the forward phase (source to destination) and others allocate in the return phase. The methods create situations where one or more paths reserve (allocate) needed spare capacities too soon, only to find that these paths do not need them or cannot use them (because a path is unable to reserve spares on all the links of the path; only on some) because reserving is done in an uncoordinated way. This can lead to a cascaded allocation and deallocation scenarios.

It is possible to construct scenarios where our algorithm chooses better paths than above-mentioned algorithms for path restoration, and vice versa. Therefore our algorithm is no better and no worse than these path restoration algorithms *in general*. However, as indicated by simulation experiments, it significantly outperforms the above algorithms on networks used by telecommunication service providers.

Awerbuch and Leighton [7] present approximation algorithms for the multi-commodity flow problem. Their approach is based on assigning “potentials” to the nodes and balance the potentials using links. For example, at the source add a potential equal to the amount of flow that is disrupted. This potential is divided uniformly among all the incident links and moved to the 1-hop neighbors of the source, if possible, subject to a local optimization rule and a rule to ensure that high potential nodes do not overtake link capacities too soon. The number of steps (rounds) needed is $O(|E|^2 K^{\frac{5}{2}} L \epsilon^{-3} \log K)$ where K is the total number of commodities, L is the length of the longest path and $0 < \epsilon$ is a parameter chosen. The advantage enjoyed by their algorithm is that if there exists a feasible solution in the current network when the demand is increased by a fac-

tor of $(1 + \epsilon)$, their algorithm finds a solution for the needed demand. (The network is “over designed” by a factor of $(1 + \epsilon)$.) The algorithm of Kamath *et al* [13] is an improvement over the algorithm of [7]. The algorithm of [13] runs in $O(|E|^2 \epsilon^{-2} \log n)$ steps (rounds) and produces flows that are $(1 - \epsilon)$ fraction of each demand, provided feasible flows exist with the current spare capacities for the required flows. In comparison, our algorithm uses $O(\text{diam}^2)$ steps (rounds), but may need more spares.

6 Guidelines for Future Work

The on line restoration can find alternate paths well below the 2 seconds time limit. However, significant amount of cross connections will have to be made at all of the nodes that are in the restored paths. At present, as the Bellcore study [1] indicates, the cross connect times are very significant resulting in total restoration times exceeding the 2 second time frame. One of the problems for future study is to find the earliest times, during the time that the restoration algorithm is running, when some part of the cross connections can be performed. This will ensure that nodes do not wait for the restoration algorithm to terminate before starting cross connection. Also, new architectures that use parallel processing are needed to speed up the cross connect operations at all of the nodes. With such changes and enhancements, it is hoped that restoration can be achieved in under 2 seconds.

Acknowledgement

The problem formulation was done in cooperation with researchers at Alcatel Network Systems and parts of the simulation study was done with their support. The algorithm described in this paper has been patented by Alcatel Network Systems. Alcatel’s help is gratefully acknowledged.

References

- [1] Digital cross-connect systems in transport network survivability. Special report SR-NWT-002514, issue 1, Bellcore, January 1993.
- [2] R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows*. Prentice Hall, New Jersey, 1993.
- [3] I. Ali, D. Barnett, K. Farhangian, J. Kennington, B. Patty, B. Shetty, B. McCarl, and P. Wong. Multicommodity network problems: Applications and computations. *IIE Transactions*, 16(2):127–134, 1984.
- [4] J. Anderson, B. Doshi, S. Dravida, and P. Harshavardhana. Fast restoration of ATM networks. *IEEE Journal on Selected Areas in Communications*, 12(1):128–138, Jan. 1994.
- [5] A. Assad. Multicommodity network flows: A survey. *Networks*, 8:37–91, 1978.
- [6] B. Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804 – 823, Oct. 1985.
- [7] B. Awerbuch and T. Leighton. Improved approximation algorithms for the multi-commodity flow problem and local competitive routing in dynamic networks. In *Proc. of ACM Symposium on the Theory of Computing*, pages 487–496, Montreal, Quebec, Canada, May 1994.
- [8] E. C. Chow, J. Bicknell, S. McCaughey, and S. Syed. A fast distributed network restoration algorithm. In *Proc. of Computers and Communications*, pages 261–267, Tempe, AZ, Mar. 1993.
- [9] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, and Y. Wang. Optical network design and restoration. *Bell Labs Technical Journal*, 4(1):58–84, Jan./Mar. 1999.
- [10] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multi-commodity flow problems. *SIAM Journal of Computing*, 5(4):691–703, 1976.
- [11] W. D. Grover, R. R. Iraschko, and Y. Zheng. *Comparative methods and issues in design of mesh-restorable STM and ATM networks*, chapter 10 in Telecommunication Network Planning, pages 169–200. Kluwer Academic Publishers, 1999.
- [12] R. Iraschko and W. Grover. A highly efficient path-restoration protocol for management of optical network transport integrity. *IEEE Journal of Selected Areas in Communications*, 18(5):779–793, May 2000.
- [13] A. Kamath, O. Palmon, and S. Plotkin. Simple and fast distributed multicommodity flow algorithm. Unpublished Manuscript.
- [14] R. Kawamura, K.-I. Sato, and I. Tokizawa. Self-healing atm networks based on virtual path concept. *IEEE Journal on Selected Areas in Communications*, 12(1):120–127, Jan. 1994.
- [15] J. Kennington. A survey of linear cost multicommodity network flows. *Operations Research*, 26:209–236, 1978.
- [16] H. Komine, T. Chujo, T. Ogura, K. Miyazaki, and T. Soejima. A distributed restoration algorithm for multiple-link and node failures of transport networks. In *Proc. of IEEE GLOBECOM*, volume 1, pages 459–463, San Diego, CA, Dec. 1990.
- [17] W. Lai and D. McDysan. Network hierarchy and multilayer survivability. RFC 3386, Nov. 2002.
- [18] K. Miyazaki, T. Chujo, H. Komine, and T. Ogura. Spare capacity assignment for multiple-link failures. In *Proc. of International Workshop on Advanced Communications and Applications for High Speed Networks*, pages 191–197, Mar. 1992.
- [19] M. Patel, R. Chandrasekaran, and S. Venkatesan. A comparative study of restoration schemes and spare capacity assignments in mesh networks. In *Proc. of IEEE 12th International Conference on Computer Communications and Networks*, pages 399–404, Dallas, 2003.
- [20] K. Ramarao and S. Venkatesan. On finding and updating shortest paths distributively. *Journal of Algorithms*, 13(2):235–257, 1992.
- [21] H. Sakauchi, Y. Nishimura, and S. Hasegawa. A self-healing network with an economical spare-channel assignment. In *in Proceedings of IEEE GLOBECOM*, pages 438–443, Dec. 1990.
- [22] Y. Xiong and L. G. Mason. Restoration strategies and spare capacity requirements in self-healing ATM networks. *IEEE/ACM Transactions on Networking*, 7(1):98–110, Feb. 1999.