Improving Recovery in GMPLS-based WSON Through Crank-back Re-routing

Edgard Jamhour, Manoel Camillo Penna
Graduate School of Computer Science
Pontifícia Universidade Católica do Paraná - PUCPR
Rua Imaculada Conceição, 1155, 80215-901, Curitiba, Brasil
e-mail: jamhour@ppgia.pucpr.br, penna@ppgia.pucpr.br

***Abstract*—This paper defines and evaluates an unplanned technique based on Generalized Multi-Protocol Label Switching (re-routing) to restore lightpaths interrupted by failures in wavelength switched optical network. Compared to the pre-planned recovery techniques, the re-routing approach may significantly save network resources, but may suffer from longer recovery times and even fail to recover lightpaths, due to slow convergence of the information transported by Interior Gateway Protocol. To address this issue, we have used the crank-back extensions proposed by IETF, combined with a make-before-break strategy that re-uses resources from the broken lightpath to setup a recovery lightpath. We present an evaluation that permits to conclude about the performance of the proposed approach.**

Keywords-crank-bak re-routingt; make-before-brake; routing and wavelength assignment.

I. INTRODUCTION

Generalized Multi-Protocol Label Switching (GMPLS) defines a set of standards for managing lightpaths in Wavelength Switched Optical Network (WSON). The GMPLS control plane supports the following recovery techniques: protection, restoration, and re-routing [6]. In protection, recovery paths are planned and cross-connected before a failure occurs. It provides fast recovery times, but are costly because backup resources cannot be shared. In restoration, recovery paths are planned and resources are reserved in advance, but recovery paths are cross-connected only when a failure occurs. Restoration is less expensive than protection because multiple recovery paths may share the same wavelengths. However, it is still expensive because shared resources cannot be used by service paths even when there are no failures in the network. Re-routing refers to the unplanned recovery technique, where all the process of defining a recovery path and reserving resources in made after a failure occurs. Compared to the pre-planned recovery techniques, it may significantly save network resources, but may suffer from longer recovery times and fail to recover lightpaths, due to slow convergence of network information transported by Interior Gateway Protocol (IGP).

In a high capacity network, a single failure can interrupt a multitude of lightpaths and trigger a strong competition for resources. The network view will be outdated, and the setup of lightpaths planned with incorrect information will probably fail. To address this issue, we have used the crank-back extensions proposed by IETF [2], which define a flexible way to include additional information in the messages exchanged by the signaling protocol, i.e., Resource Reservation Protocol with Traffic Engineering (RSVP-TE). The extensions permit to include the information required to plan an alternate route in case of failure, and to modify the flow of the signaling messages to contour the parts of the network that are interrupted. Our recovery approach includes the following ideas. First, the nodes adjacent to the failure use the crank-back extensions to inform the ingress nodes about the information required to recover the interrupted lightpaths. Second, the recovery lightpath is planned by the ingress node using a load balance heuristic, which avoids the creation of bottlenecks and favors the reuse of resources. Third, recovery is performed using a Make-Before-Break (MBB) strategy, to reuse as much as possible the resources and cross-connects of the original lightpath that survive the failure. Finally, signaling is performed using a flexible segment re-rerouting strategy, permitting any node along the path to fix the information planned by the ingress node. MBB is pointed as being advantageous to improve the likelihood of a successful recovery (see [6], for example), but no previous work has detailed how it could be implemented in WSON.

The remaining of this paper is organized as follows. In Section 2, we review the WSON literature by focusing on improvements to RSVP-TE and crank-back. Section 3 explains the problems that may rise in an unplanned attempt to recover lightpaths and how we address the pointed issues. Section 4 presents the algorithms that compose our solution. The evaluation of the proposed method is found in Section 5. Finally, Section 6 presents the summary of our most important results and our vision about future research topics related to the subject.

II. RELATED WORK

Some improvements to RSVP-TE have been proposed to increase the likelihood of a successful label suggestion assignment during path creation. Sambo et al. [12] review several strategies that employ the label preference approach. The suggested vector object is introduced by Andriolli et al. [1] for networks with wavelength conversion capability. It collects information about the number of conversions that will be performed by intermediate nodes. This information permits the destination node to select a wavelength from the Label Set that minimizes the number of conversions. The suggested vector approach is further explored by Giorgetti et al. [3] to avoid contention of wavelengths due to outdated information in nodes that receive Path or Resv messages. The proposals previously mentioned improve RSVP-TE but are not comparable to our work because they don´t make use of crank-back.

Planning protected lightpaths using a shared protection scheme is discussed by Munoz et al. [10]. The proposed extension of RSVP-TE includes information indicating which wavelengths in a Label Set are already being used by a protection lightpath. The same extensions are employed by Manolova ett al. [9], but considering networks with a limited number of wavelength converters. The paper extends the previous proposal by combining the idea of the suggested vector to reduce the number of wavelength converters along the path. Manolova et al. [7] [8] extend the same approach to include the sharing of optical regenerators. Giorgetti et al. [4] explore the use of suggested vector to improve Resv blocking, and evaluates the strategy in scenarios with or without crank-back attempts. The proposals discussed in this paragraph don't cover segment-based rerouting because the error messages always propagates to the ingress node, which is responsible for generating a new setup attempt.

More recently, some alternative approaches have been proposed. Pavani and Waldman [11] present a Routing and Wavelength Assignment (RWA) strategy with crank-back support based on Ant Colony Optimization (ACO) algorithm. The proposed strategy can be classified as segment-based re-routing, however, instead of using the RSVP-TE extensions, or IGP updates to propagate the crank-back information, the authors assume an ACO based algorithm that updates local state information of different aspects of the routing process. Chen et al. [13] propose a new routing protocol based on the concept of intensity gradient from an information source. It is based on a distance-vector routing scheme that enables the re-routing capability on every intermediate node, which maintains all possible link-disjoint routes to the destination node. The proposal includes a new signaling protocol that implements the information-diffusion-based routing. Because the proposals discussed in this paragraph are based in proprietary protocols, their corresponding approaches require a complete modification of the IGP algorithm and routing information presently used in GMPLS.

To the extent of our knowledge, the literature about the use of crank-back extensions in WSON is still weakly explored. There is nothing in the literature comparable to the study presented in this paper, in terms of exploring the signaling protocol extensions to define a method to recovery of lightpaths using a purely distributed approach that is robust against the problems caused by the slow convergence of IGP information. In a previous study, Jamhour and Penna evaluated [14] eight different network topologies to determine which network features favor the crank-back strategy, considering several network metrics, including some used in Social Network Analysis (SNA), allowing to find the criteria that permits to identify the situations in which the crank-back approach, or other re-route strategy is advantageous. However, the algorithms presented in this paper are totally new.

In special, we define an approach to coordinate the recovery attempts according two strategies: MBB and Break-Before-Make (BBM). In the next section, we present some examples of recovery problems caused by the IGP slow convergence and define the strategies to improve the likelihood of a successful restoration using a combined crank-back and re-routing strategy.

III. PROBLEM FORMULATION

In this section, we show how the MBB approach may be useful in WSON and why it may result in a temporary deadlock in some situations. We propose the use Notify messages to improve the ingress node perception about the possibility of completing a successful MBB recovery. The discussion in this section is based on the scenario in Figure 1. All links are supposed to have only two wavelengths at each direction. There are two uni-directional lightpaths created between the nodes 2 and 6 (represented by the square and circle symbols), and one uni-directional lightpath created between the node 1 and 6 (represented by the triangular symbols). The symbols in the links between nodes represent the direction and the wavelengths used by each lightpath. In this setup, no wavelength converter is used because lightpaths use the same wavelength in all links.

Suppose that link 4-5 fails. The failure is perceived by the adjacent nodes, and interrupts the three lightpaths. A Notify can be used to inform the ingress node of every lightpath affected by a failure (see Figure 2). It is not necessary that both nodes generate a Notify message about the same interrupted lightpath, but according to our recovery method, both nodes send the message.

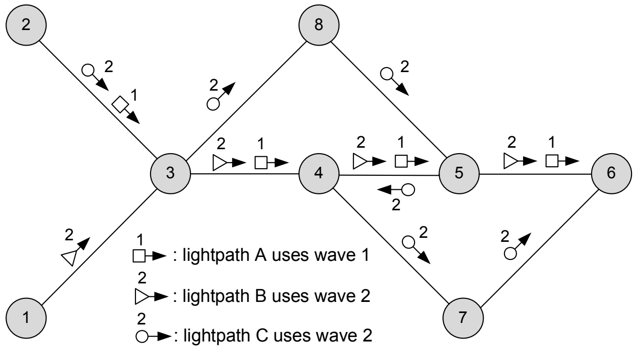

Figure 1. Sample scenario with three lightpaths.

In standard re-routing recovery, the ingress node must setup a recovery lightpath for each broken lightpath reported in the Notify message. The lightpath can be planned by the ingress node, or constructed in a distributed way. Each node in the network has its own view of the availability of resources. Ideally, a node should know about each wavelength available at each link, and the availability of wavelength converters in the nodes. However, flooding information about individual wavelengths is not practical. Moreover, in case of failure, this information is supposed to change very fast because several attempts of lightpath setups will be performed simultaneously. We assume that the only information available is the link state and the number of free wavelengths in each link.

 ISBN: 978-1-61208-318-6

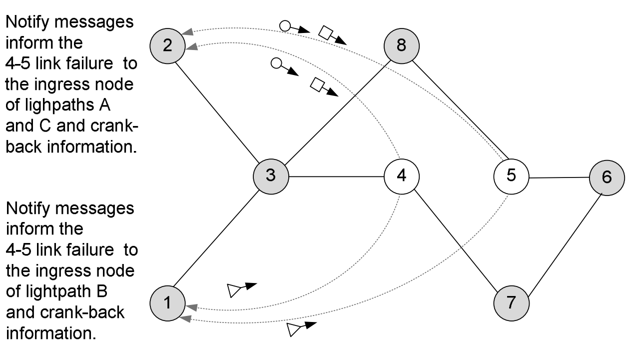

Figure 2. A Notify message is used to report the ingress node of each interrupted lightpath.

Figure 3 shows how the node 2 perceives the recovery options for the lightpath A, in a scenario without wavelength converters. The straight lines in the graph represent the status of the links based on IGP information, and the curved lines additional crank-back information supplied by the Notify messages. Without considering the availability of the wavelengths used by the original path, the recovery is unfeasible. The ingress node cannot perform an immediate recovery attempt without trying to reuse the wavelengths of the broken lightpath. The IGP information indicates that there is one wavelength available at the links 3-8 and 8-5. However, it is not possible to know if the wavelength "1" is available at these links. We use the Notify messages to indicate if the wavelength used by the broken lightpath can be successfully cross-connected to its adjacent edges. The cross-connect is possible if the same wavelength is available, or if it can be converted to an available wavelength. Node 4 informs the ingress node that a recovery attempt using MBB is possible for the link 4-7 and node 5 informs the same for link 8-5. The ingress node, however, does not have additional information about the links 3-8 and 7-6.

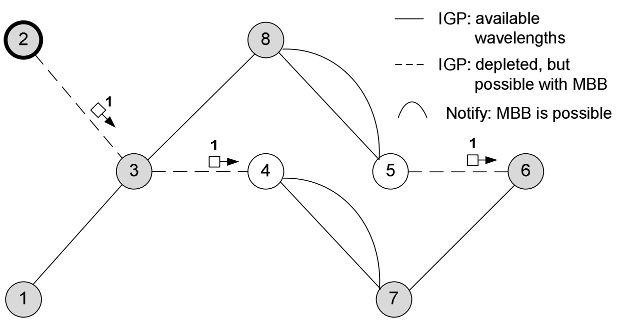

Figure 3. Recovery options for lightpath A perceived by the node 2.

The recovery options for lightpath B are computed by ingress node 1 (see Figure 4). Again, the node has no recovery options for lightpath B without considering the reuse of the wavelengths of the broken lightpath. If the decision is based exclusively on IGP information, node 1 will consider that MBB can be successful. However, the Notify messages sent by nodes 4 and 5 inform that it is not possible to cross-connect the wavelengths of the original lightpath to the links 4-7 and 8-5, because the wavelength is already in use, and the nodes have no wavelength converters. At present situation, the node has no immediate recovery option. Observe in the legend of the figure that we used the term “MBB is unlikely”, instead of “MBB is impossible”. MBB would be impossible if the wavelength required to perform cross-connect belongs to a lightpath that is not interrupted by the failure.

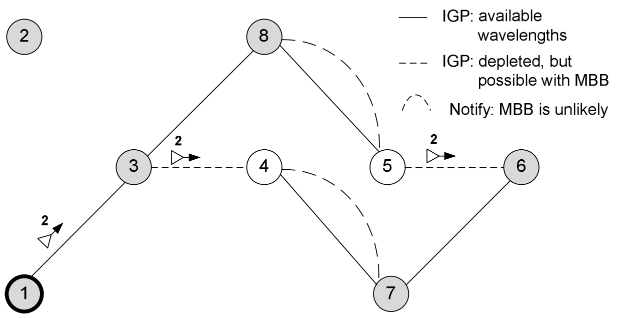

Figure 4. A Notify message is used to report the ingress node of each interrupted lightpath.

As indicated in Figure 5, if node 2 viewpoint of network resources is based solely on the information received by IGP, the recovery of lightpath C will be unfeasible (if the depleted links are removed, nodes 2 and 6 become disconnected). However, the Notify messages will indicate that a recovery may be possible in a near future, because some wavelengths are required to complete a MBB setup belongs to broken lightpaths. The situations of the lightpath B and C are similar, because none of them have an immediate recovery option.

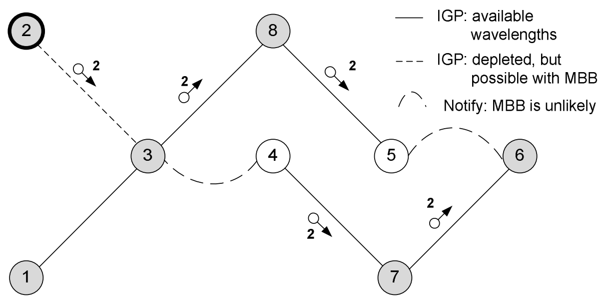

Figure 5. Recovery options for lightpath C perceived by the node 2.

The dynamic planning of recovery lightpaths may be done according to MBB or BBM. In the first, the original lightpath is teared down before the setup of the recovery lightpath. In the second, the resources are released only after the setup of the recovery lightpath is confirmed or aborted. If nodes 1 and 2 try MBB recovery, they would be in deadlock state (until the reservation is broken by the soft-state). On the other hand, if they try BBM recovery, it would be possible to recover all lightpaths. However, BBM approach cannot be performed immediately, because the release of the resources of the broken lightpaths is not instantaneous, due to the delay of the tear down message propagation. In addition, nodes will not perceive the new resources instantaneously, because the slow convergence of IGP. Also, a BBM attempt may be slow, because the lightpaths will compete for the same resources, and in this case, the RSVP-TE may block setup attempts.

Our solution is the following: always when possible, the recovery of lightpaths will consider performing MBB. From viewpoint of the ingress node, MBB is possible if the recovery path does not contain any edge assigned by the Notify messages as “MBB unlikely” or “MBB impossible”. The reuse of wavelengths in MBB is not mandatory, if the node finds a more advantageous recovery path that is disjoint with respect to the original path. For the lightpaths that do not satisfy this condition, BBM will be performed, and resources are released immediately after receiving the Notify message. When the ingress node is unable to find a candidate path to perform a recovery attempt, it will wait a random timeout (back-off), with a minimum safeguard time to receive updates from IGP.

The feasibility of MBB may change in a scenario where nodes are capable to perform wavelength conversion. In the scenario of Figure 1, the situation of lightpaths A and C will not change. For lightpath A it would still be possible to perform MBB recovery, and lightpath C would still have no wavelengths available in edges 3-4 and 5-6. However, the situation for lightpath B would change because now it is possible to perform a cross-connect of the MBB wavelengths to the edges 4-5 and 8-5. In this case, lightpaths A and B would perform MBB recovery and lightpath C would perform BBM recovery.

IV. Proposed Solution

To provide some level of load balancing, the ingress node computes an explicit route to the destination using a load balancing heuristic. The most common heuristic consists in assigning a cost to a link that is proportional to the fraction of wavelengths in use with respect to the total number of wavelengths. The Weighted-Shortest-Cost-Path (WSCP) proposed by Hsu et al. [5] follows this strategy. We have modified WSCP to take into account the possibility of performing MBB, as defined in equation (1). A reduction factor (*mbbfactor*) is used to favor routes that reuse the wavelengths of the broken lightpath. In the formula, *hopweight* = *mbbfactor* if the edge contains a reusable wavelength and MBB is possible. Otherwise, *hopweight* = 1. In the expression, P_{sd} is the set of edges e connecting the source node to the destination node, and $\mathrm{fw}(e)$ and $\mathrm{w}(e)$ are, respectively, the number of free wavelengths and the total number of wavelengths in e. In the expression, *bfactor* controls the relative importance between assigning paths that contribute to load balance or are shorter in number of hops. Algorithm 1 is used to determine the recovery route and is responsible to make the decision to use MBB or BBM.

$$\Theta(P_{sd}) = \sum_{e \in P_{sd}} hopweight + \frac{\mathrm{fw}(e)}{\mathrm{w}(e)} \cdot bfactor \cdot \qquad (1)$$

In the following, we show how to build the recovery path using the route computed with the algorithm in Figure 6. According to RSVP-TE, the Path message may include a Label Set (*lset*) that restricts the range of wavelengths that can be selected by the downstream node. Nodes capable of performing wavelength conversion may expand the *lset*. The Path message may also include a Suggested Label (*sl*), a wavelength chosen from the *lset* that is preferentially offered to the downstream node. If the downstream node is able to use *sl*, it performs a cross-connect between the *sl* received from the upstream node and the *sl* offered to the downstream node. Once the Path message is received by the egress node, it selects the Generalized Label (*gl*) and transmits it upstream using the Resv message. If the *gl* is different from the *sl*, a node must remake the cross-connect with the *gl*. The Explicit Route Object (*ero*) permits to define the route and the wavelengths used along the path. The crank-back extensions introduce the possibility to fix blocked setup request without signaling a new setup request from the ingress node. Segment-based re-routing allows any upstream node that receives an error message to make a correction in the setup request through a new Path message.

Begin: The algorithm is triggered by the first Notify message received by the ingress node.

1. Save the information received in the first Notify message in: “failed”, “MBB likely”, “MBB unlikely” and “MBB impossible” edge sets.
2. Wait for the second Notify message. If message arrive before timeout, go to Step 3, otherwise go to Step 4.
3. Update the information received in the edge sets.
4. Create a graph including the edges that are not in the “failed” edges set; the edges not depleted according to IGP; and the edges depleted but with wavelengths used by the original broken lightpath. Depleted edges usable only with MBB waves are called “MBB only”.
5. Compute a list of candidate paths by considering the k-shortest paths with respect to the number of hops.
6. Eliminate from the list of candidate paths all paths that include at least one “MBB only” edge and at least one “MBB unlikely” or “MBB impossible” edge.
7. If the remaining set of candidate paths is not empty, go to Step 9.
8. Tear down the original lightpath to free its resources, and perform a recovery attempt without explicit routes after a back-off timeout. Terminate the algorithm.
9. Select the best route among the candidate paths according to the cost function in equation (1). If the best route does not contain any edge with a MBB wavelength, tear down the original lightpath to free its resources. Perform a recovery attempt using the best route as an explicit route. Terminate the algorithm.

Figure 6. Algorithm 1: Determine the recovery route.

Figure 7 illustrates the sequence of messages required to setup a recovery lightpath using our method. The Path message is generate using an *ero* with explicit labels, associated to each hop in the *ero* object (see the *lero* sub-object in the figure). The ingress node fills the *lero* with the MBB wavelengths that are present in the *ero*. New links in the recovery path have no explicit label (they are indicated as "0" in the *lero* sub-object). The ingress node also indicates the lightpath is being recovered using the Association Object (*ao*), informing that for the broken lightpath, a cross-connect can be undone and the reserved wavelengths can be released.

The Path planned by the ingress node may be unfeasible. In the example in Figure 7, the ingress node knows there is a free wavelength at the link 4-7. However, it does not know which wavelength it is, neither if the node 4 may perform a conversion to this wavelength. Because of this, node 4 generates an error, sending a PathErr message back to node 3. Instead of forwarding the error upstream, node 3 computes a new path to the destination excluding node 4. The new path is included in a new Path message as an explicit route and sent to the node 8 (re-routing). The Path message may also carry an excluded route object (*xro*), in order to inform to the downstream nodes known blocking resources. This is necessary if another node is required to solve a blocking by performing another segment re-route.

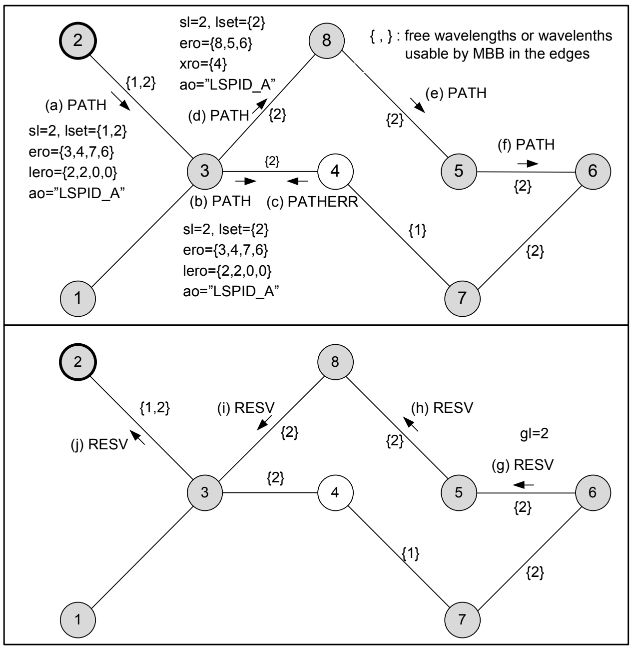

Figure 7. A Notify message is used to report the ingress node of each interrupted lightpath.

The algorithm in Figure 8 describes the procedure to determine the next node to forward a Path message. This procedure can be triggered by a Path message received from the upstream node, or a by PathErr message received by the downstream node. In the first case, some information such as the upstream label set must be retrieved from the node state database. The set of adjacent edges (local ports) with no wavelengths that satisfy the label set (*blockedPorts*) is computed using the local information of the node. The *xro* and the *nex* objects are specific crank-back information created by the node or received from the upstream node.

V. EVALUATION

We have developed a simulator for the GMPLS control plane using Wolfram Mathematica. The RSVP-TE messages propagate as individual packets and are delayed by the transmission rate, link propagation and queuing in the input and output ports of the nodes. The Reconfigurable Optical Add-Drop Multiplexer (ROADM) is able to process only one RSVP-TE message at a time. We have included in the simulator all elements required to estimate the setup time. The control plane messages propagate as individual packets and are delayed by the transmission rate (1Gbps), link propagation and queuing in the input and output ports of the nodes. The ROADM nodes are modeled as single processor entities, i.e., each node is able to process only one RSVP-TE/SDN protocol message at a time. Incoming messages are queued and processed sequentially in a FIFO. An optical cross-connect (i.e., the creation a flow in a WOFS) is the most timing consuming operation. The time to perform an optical cross-connect is 10 ms and to release a cross-connect, 5 ms. The time consumed to process Path and Resv messages is 2 ms. The time consumed to process PathErr, ResvErr, PathTear is 1 ms. Lightpaths are teared down explicitly.

Begin: the algorithm is triggered by a Path or a PathErr message.

1. If the node has no wavelength converter, determine *blockedPorts* as the set of adjacent edges with no free wavelengths included in *lset*. Otherwise, set *blockedPorts*= Ø.
2. If the procedure has been triggered by a Path message, and it includes an *ero*, determines next hop from it. If the edge connecting to the next hop does not belong to *blockedPorts*, go to step 8.
3. Determine the set of edges that has no more wavelengths available: *depletedEdges*.
4. If the procedure has been triggered by a PathErr message and *nex* is present, set *xro* with the nodes in *nex*. If *nex* is not present, set *xro* with the node that generated the PathErr message.
5. Build a graph excluding the edges in *depletedEdges* and *blockedPorts* and the *xro* nodes.
6. If the graph is connected, go to Step 7. Otherwise, go to Step 9.
7. Compute the recovery path from the current node to the egress node using the metric given by Equation (1), with *hopweight* = 1. Set *ero* (with no explicit labels) with the new recovery path.
8. Update the *ls* with the wavelengths that can be cross-connected from the incoming *ls* to the local port connecting to the next hop. If the next-hop has an explicit label in the *ero*, set the corresponding wavelength as the sl. Otherwise, selects a random wavelength from the ls as the next sl. Send the Path message to the downstream node, and terminate.
9. Send a PathErr message to the upstream node including itself in *nex*, and terminate.

Figure 8. Algorithm 2: Build and forward a Path message.

The topologies of the control plane and the data plane are identical. All links have 32 wavelengths and all nodes have a shared converter pool with capacity to perform 8 wavelength conversions. The parameters *mbbfactor* and *bfactor* in Equation (1) are set to 0.5 and 4, respectively.

 ISBN: 978-1-61208-318-6

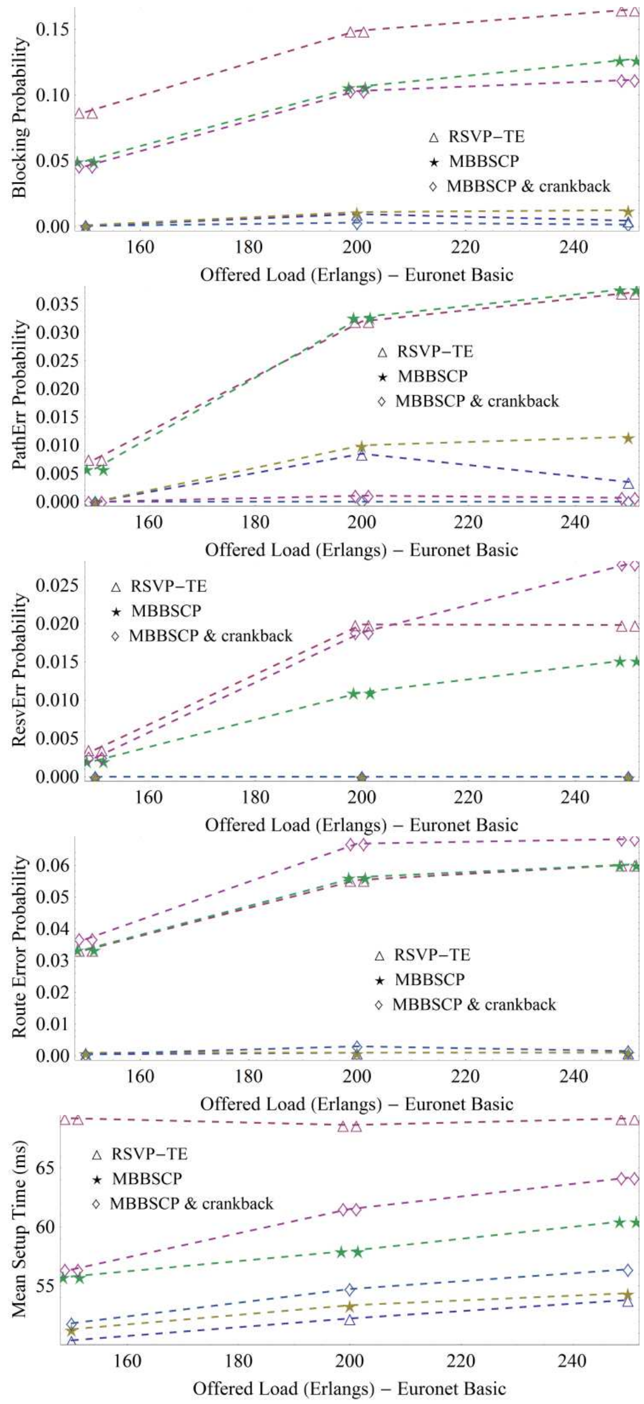

Figure 9. Evaluation of the Euronet basic topology.

The traffic load is generated as in most papers discussed in our review section. The network is submitted to a load of setup requests uniformly distributed among all pairs of nodes. The requests are controlled by two exponential variables: average interval among requests $1/\lambda$ and average duration of the lightpaths $1/\mu$. The total setup request load is measured in Erlangs λ/μ. To generate a variable setup request load we set $1/\mu$ = 2400h, and we vary the value of λ. For each load scenario we performed the simulation for 2000 setup requests. The number of failures is variable, because they are based on the exponential failure rates computed taking into account the length of the links, and the number of optical amplifiers in the spans. Failures of nodes are perceived as multiple link failures by the adjacent nodes. Depending on the failure, or the importance of the link or node that failed, dozens of simultaneous recovery attempts may be performed simultaneously. In general, the number of the recoveries in the evaluated scenarios varied between 2000 and 3000. The amount of simulated connections resulted in a small standard deviation, of the order of 10^{-4} for the average probabilities and of the order of 10^{-1} for the mean setup time.

The following methods are evaluated: (i) RSVP-TE: standard distributed RWA. (ii) MBBSCP: uses explicit routes to support MBB, according to Algorithm 1. (iii) MBBSCP & crank-back: uses explicit routes to support MBB, and the crank-back re-routing, according to Algorithm 2. We present the results obtained for two distinct topologies, based on variations of the Pan-European network (basic and large) and the NFS network. For all networks, we have assumed that the topologies of the control plane and the data plane are identical. The control plane uses a reserved wavelength in all links with a throughput of 1 Gbps.

Figure 9 shows the obtained results for the Pan-European network basic topology. The performance metric is the average blocking probability. The blocking probability of the first setup is indicated as single plot markers, and the recovery blocking probability is indicated as double plot markers. In all scenarios, the best performance is obtained by the MBBSCP & crank-back approach, followed by the MBBSCP approach, indicating that the major influence results from the coordination of the recovery attempts. Blockages caused by the exhaustion of drop ports are not considered because it cannot be controlled by any of the methods. The reason for exhaustion is the variable load that can saturate drop ports on the ends of the connections. However, the failure to consider this effect does not affect the results, because it occurs in all methods evaluated.

There are basically three main reasons for a setup attempt not to be completed: (i) A PathErr, caused when a node is not able to find a wavelength in the downstream port that satisfies the incoming label set restrictions. (ii) A ResvErr, caused when the label selected by the downstream node cannot be used, because this label has been assigned to another lightpath since the Path message was forwarded. (iii) A route error, caused when a node cannot find a route to the egress node (caused by depleted edges or failed edges or failed nodes). The MBBSCP is expected to reduce the number of route errors. Crank-back is expected to reduce the number blocking caused by PathErr.

The number of recovery attempts varied from 2248 (lowest load scenario) to 2914 (highest load scenario). The number of recovery setups completed with the help of crank-back re-route increases consistently with the load of the network. It is insignificant for the lowest load (150 Erlangs), but achieves 3.5% at 200 Erlangs and 4.2% at 250 Erlangs.

In the highest load situation, crank-back re-route is required even to help completing the setup of 1.65% of the lightpaths when they are first provisioned. At the highest load, the proposed method has reduced to almost zero the number of blocked recoveries caused by PathErr messages. The setups are blocked mainly due to route error (6.8%) and Resv error (2.8%). The crank-back slightly increases the recovery setup time, because requires a higher number of messages to complete the setup.

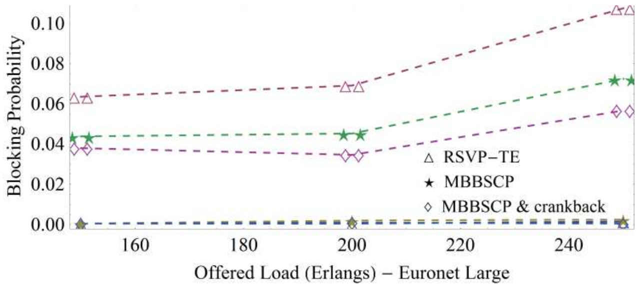

Figure 10. Evaluation of the Euronet large topology.

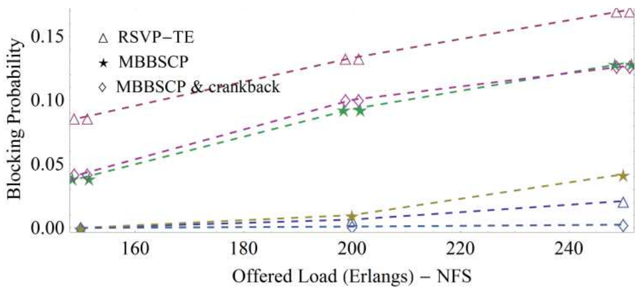

Figure 11. Evaluation of the NFS topology.

The results for the other topologies are similar to the Euronet basic, so we are going to present them briefly. For the Euronet large (see Figure 10) and for the NFS basic topology (see Figure 11), it can be observed that the recovery blocking probability steadily increases with the offered. Because these networks provide many recovery alternatives, the advantage of the proposed methods is more visible. Observe that the distance between the RSVP-TE approach and our proposed extensions increases with higher loads.

VI. CONCLUSION AND FUTURE WORK

In optical networks, the unplanned recovery technique based on re-routing poses a number of difficulties that are not observed in packet switched networks. In this paper, we have proposed a method to improve the robustness of lightpath setup performed in a distributed scenario. The proposed approach takes use of the flexibility provided by the Notify messages and crank-back extensions introduced by GMPLS. Our evaluations showed that our approach can significantly reduce the blocking probability of recovery attempts. However, there is still room for improvement. In special, with regard to the method proposed in this paper, the ResvErr messages have not been handled by crank-back. ResvErr occurrence is insignificant when the network is in a normal state of operation, but it is an important factor to be addressed during restoration, because the concurrence for resources may prevent a node to honor the wavelengths offered by the Path messages. We intend to address this issue by improving the crank-back re-route logic to also take into account this effect. We also intend to develop a method for dimensioning the network to give a degree of assurance about the success of lightpaths restorations.

REFERENCES

[1] N. Andriolli, J. Buron, S. Ruepp, F. Cugini, L. Valcarenghi, and P. Castoldi, “Label preference schemes in gmpls controlled networks,” Communications letters, IEEE, vol. 10, December 2006, pp. 849–851.

[2] A. Farrel, A. Satyanarayana, A. Iwata, N. Fujita, and G. Ash, “Crankback Signaling Extensions for MPLS and GMPLS RSVP-TE,” RFC 4920 (Proposed Standard), July 2007.

[3] A. Giorgetti, N. Sambo, I. Cerutti, N. Andriolli, and P. Castoldi, “Label preference schemes for lightpath provisioning and restoration in distributed gmpls networks,” Journal of lightwave technology, vol. 27, March 2009, pp. 688–697.

[4] A. Giorgetti, N. Sambo, I. Cerutti, N. Andriolli, and P. Castoldi, “Suggested vector scheme with crankback mechanism in gmpls-controlled optical networks,” In Optical Network Design and Modeling (ONDM), 14th Conference on, February 2010, pp. 1-6.

[5] C. Hsu, T. Liu, and N. Huang, “An adaptive routing strategy for wavelength-routed networks with wavelength conversion capability,” In Communications, ICC 2002, IEEE International Conference on, vol. 5, April 2002, pp. 2860–2864.

[6] J. Lang, B. Rajagopalan, and D. Papadimitriou. Generalized multiprotocol label switching (gmpls) recovery functional specification. RFC 4426 (Proposed Standard), March 2006.

[7] A. Manolova, et al., “Distributed sharing of functionalities and resources in survivable gmpls-controlled wsons,” IEEE/OSA Journal of Optical Communications and Networking , vol. 4, March 2012, pp. 219 –228.

[8] A. Manolova, et al., “Wavelengths and regenerators sharing in gmpls-controlled wsons,” In Global Telecommunications Conference (GLOBECOM 2010), IEEE, December 2010, pp. 1-5.

[9] A. Manolova, et al., “Shared path protection in gmpls networks with limited wavelength conversion capability,” In High Performance Switching and Routing (HPSR), 2010 International Conference on, June 2010, pp. 203–208.

[10] R. Munoz, R. Casellas, and R. Martinez, “An experimental signalling enhancement to efficiently encompass wcc and backup sharing in gmplsenabled wavelength-routed networks,” In Communications, ICC’08 IEEE International Conference on, May 2008, pp. 5401–5406.

[11] G.S. Pavani and H. Waldman. “Routing and wavelength assignment with crankback re-routing extensions by means of ant colony optimization,” IEEE Journal on selected areas in communications, vol. 28, May 2010, pp. 532 –541.

[12] N. Sambo, N. Andriolli, A. Giorgetti, and P. Castoldi, “Wavelength preference in gmpls-controlled wavelength switched optical networks,” Network protocols and algorithms, vol. 3, July 2011, pp. 110–125.

[13] Yue Chen, Nan Hua, Xiaoping Zheng, and Chunming Qiao, "Experimenting with immediate re-routing on an information-diffusion-based routing test-bed," Optical Fiber Communication Conference and Exposition (OFC/NFOEC), and the National Fiber Optic Engineers Conference , March 2011, pp.1–3.

[14] E. Jamhour, and M. Penna, "Evaluation of segment-based crankback re-routing for GMPLS-based WSON," Telecommunications (ICT), 20th International Conference on , May 2013, pp.1–5.