A New Reconfiguration Protocol for Mesh Optical Networks

A. Shami, *Member, IEEE,* C. Assi, *Member, IEEE,* Nasir Ghani, *Senior Member, IEEE,* and H. T. Mouftah, *Fellow, IEEE*

***Abstract*— Resource efficiency is a key concern in dynamic wavelength-routing networks. In particular, link failure events can yield notable inefficiencies as both failed connections and newly-arriving connections can be routed over less efficient routes. It is here that network reconfiguration after fault repair can yield substantial gains. To study the benefits of this approach, we propose a novel distributed token-based approach to coordinate post-repair reconfiguration. The goal is to minimize overall contention between simultaneous re-routing attempts over the same set of resources. Our findings confirm that the proposed scheme attains good gains in resource efficiency and amenable reconfiguration timescales.**

***Index Terms*— Optical networks, network reconfiguration, restoration, performance evaluation.**

I. Introduction

CONTINUED growth in end-user demands combined with ongoing technological advances in optical transport/switching technologies have opened up new scalabilities and paradigms in bandwidth services provisioning. However, owing to the high dimensionalities involved, optical networks must perform efficient resource management in order to achieve stringent availability for mission critical services. In particular, service disruptions following network failures can result in significant traffic disruption (and revenue loss) in the absence of effective recovery procedures. Now in general, there are two approaches for providing recovery in optical networks: *proactive* methods in which (shared or dedicated) spare capacity is pre-configured and *reactive* methods, in which spare capacity is dynamically provisioned (re-routed) upon failure detection [1]. Typically, the latter types are more resource efficient. Recently in [2], the authors showed that a well-designed reactive restoration schemes can deliver very high availability (over 99.5%) and rapid recovery timescales (order of few tens of milliseconds even with frequent failures). This study considers further extensions to such reactive restoration frameworks.

Under normal conditions, lightpath connections are routed along their optimal routes using a suitable routing metric/objective, e.g., resource minimization (hop count), load balancing, hybrid, etc. Meanwhile, upon a failure event, restoration schemes basically re-route disrupted connections onto alternative backup paths. Now while a network fault is being repaired, it expected that certain incoming connections will not be provisioned along their absolute optimal (e.g., shortest) paths. Hence this will lead to increased resource consumption, as longer routes will usually be selected. In addition, as it has been shown in [3] that the use of longer connection paths can further deteriorate network throughput and possibly lower stability. These issues are of particular concern in operational settings as repair procedures can commonly last days or even weeks [1]. As a result, it is conjectured that dynamic post-repair *reconfiguration* can potentially improve resource utilization.

In this paper we propose a dynamic reconfiguration scheme in which selected connections that have been setup after a failure are re-routed in order to release extra capacity, i.e., upon completion of link repair. Namely, a graceful token-based distributed approach is developed where a token entity is passed between nodes using a pre-determined token walk path. A key property of this scheme is that only a single node is allowed to reconfigure its connections at any given time. This is designed to mitigate multiple simultaneous reconfiguration attempts over the same resource pool, thereby increasing overall reconfiguration stability. The paper is organized as follows. The proposed algorithm is presented in Section II and its related performance is studied in Section III. Final conclusions are discussed in Section IV.

II. Network Reconfiguration

Following a failure, network resource allocation may not be very efficient as existing connections are usually re-routed over longer paths and new connections may be routed along non-optimal paths. Hence upon fault *repair* it is very desirable even expected that the network will return to its normal operating state. This requires that some of the connections provisioned (re-provisioned) during the repair cycle be reconfigured back to their optimal shortest paths, i.e., to free up any excess sub-optimally reserved capacity. This is best illustrated by an example in Fig. 1. Here up to three connections (A-F, F-A, and B-A-F) can be affected by failure of link A-F and hence must be restored onto alternative routes (A-B-C-F, F-C-B-A, and

Manuscript received October 14, 2004. The associated editor coordinating the review of this paper and approving it for publication was Prof. I. Fair. This work was supported in part by the National Science and Engineering Research Council of Canada (NSERC).

A. Shami is with the Department of Electrical and Computer Engineering, The University of Western Ontario, London, ON, N6A 5B9 Canada (e-mail: ashami@eng.uwo.ca).

C. Assi is with the Concordia Institute for Information System Engineering, Concordia University, Montreal, QC, Canada.

N. Ghani is with the Electrical and Computer Engineering Department at Tennessee Tech University, Cookeville, TN, USA.

H. T.Mouftah is with the School of Information Technology and Engineering, University of Ottawa, Ottawa, ON, Canada.

Digital Object Identifier 10.1109/LCOMM.2005.06022.

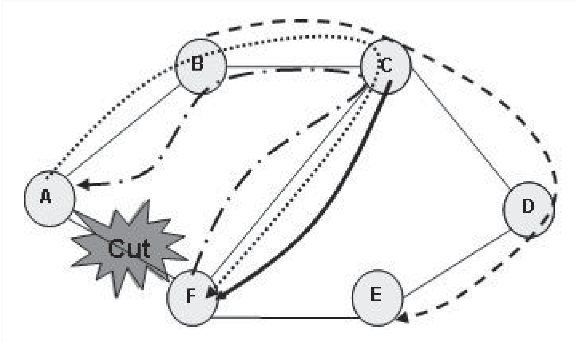

Fig. 1. Sample illustrative network.

A-B-C-D-E-F, respectively). It is evident that this resulting configuration consumes more network resources versus the original working configuration, e.g., 11 wavelengths versus 5 wavelengths. Clearly, reconfiguring these affected connections back to their original routes after repair of link A-F will improve resource efficiency (6 wavelengths) and lower future request blocking. However, a key challenge here is the proper coordination of reconfiguration activity. In particular if many affected connections simultaneously attempt reconfiguration, resulting contention can adversely impact overall performance, i.e., increased reconfiguration blocking.

This paper proposes a novel scheme for graceful *post-repair* connection re-routing (reconfiguration). The solution uses a token-based approach in which a token is passed between all nodes using a pre-determined walk path. Here, a node can only capture the token if it has connections to re-route, otherwise the token bypasses the node until it is finally discarded by the last node on the walk path. Hence when a link has been repaired, a token is created by one of the end nodes of the link (e.g., the one with greater address) and launched into the network with its walk-path is appended. Now in order to minimize overall reconfiguration times, a shortest token walk path must be chosen. This can be done using a *minimum spanning tree* (MST) approach is used to resolve the walk path using a depth-first search heuristic [4], complexity $O(N\ log(N))$, where N is the number of nodes. To further improve responsiveness, it is also assumed that all token walk paths are computed and stored in an offline manner. Hence a node only need to re-compute a walk path if the stored route contains a failed link (or node). Note that the use of fixed MST token walk paths implies that the reconfiguration procedures will follow a predetermined sequence for a given link failure. However, differing link failures will likely have different MST walk path sequences.

Now consider per-node token processing. Upon receipt of a token, a node must first search its table of active *sourced* connections to check whether any them require re-routing, i.e., if the original shortest path route includes the repaired link. For each such connection, the node runs an appropriate *routing and wavelength assignment* (RWA) algorithm to compute a (new) shorter path and subsequently attempts to reserve resources using distributed signaling [5]. If this setup is successful, the source switches all affected traffic to the shorter path and releases resources along the previous path. Note that the actual RWA algorithm is outside the scope of this paper, and many possibilities exist. However, in this study the scheme proposed in [5] is adopted, i.e., distributed routing with minimum hop count metric (i.e., resource minimization) with *first-fit* wavelength assignment [1]. Furthermore, in order to improve the timeliness of the approach (i.e., minimize token holding times), the each node receiving a token performs *parallel* connection reconfiguration. Namely, it can re-compute (via RWA) the routes of all affected connections and concurrently initiate connection switchover signaling, i.e., ensuring unique wavelength assignments over any common connection links.

Now a key performance metric for any optical layer scheme provisioning, recovery, reconfiguration is its operational timescale. Here, consider the specific case of reconfiguration, in which existing (sub-optimally routed) connections are re-routed after repair of the failed link. It is important to note that prior to reconfiguration, all of these connections are already in operational mode, i.e., non-failed state. As such, the re-routing of these connections (upon fault repair) need not be governed by stringent protection/restoration timescales, which pertain more to failed connections. Instead, it is widely expected that larger timescales will be more amenable here since reconfiguration is expected to be applied over timescales on the order of connection duration [1]. As a result, the choice of serial reconfiguration (via a token walk) is very much acceptable, as the broader focus is more upon resource efficiency and overall network stability. Hence, reconfiguration timescales will be on the order of number of connections that may be affected by a failed (repaired) link and their associated path lengths, i.e., $O(NC)$, for N nodes with C connections, $O(N^3)$ in worst case if C is $O(N^2)$. Furthermore, performing parallel connection reconfiguration at each node improves the timeliness of the approach.

III. Results and Discussions

The performance of the proposed reconfiguration scheme is studied via discrete event simulation on a 16-node NSFNET network topology. Here, dynamic traffic models are used in which call requests arrive at each node according to a Poisson process with a network arrival rate λ and session holding times are exponentially distributed with mean $1/\mu = 60\ minutes$. Furthermore, it is assumed that link faults occur at a rate of 0.0015 cuts/second and associated link repair times exponentially distributed with mean 300 minutes (5 hours). Carefully note that link failure inter-arrival and repair times are chosen to be much smaller than those in real-world settings. Additionally, average repair times are set significantly larger than average fault inter-arrival times. This is done in order to stress the performance of the proposed scheme. Furthermore, for comparative reasons, both path-based and link-based restoration schemes are considered [5].

Fig. 2. shows the network performance (for both link and path rerouting) with and without reconfiguration; as expected, network configuration helps freeing some of the reserved extra resources that will be available to accommodate future connections which in turns yields a better blocking probability. Also, simulations show that the token-based scheme yields

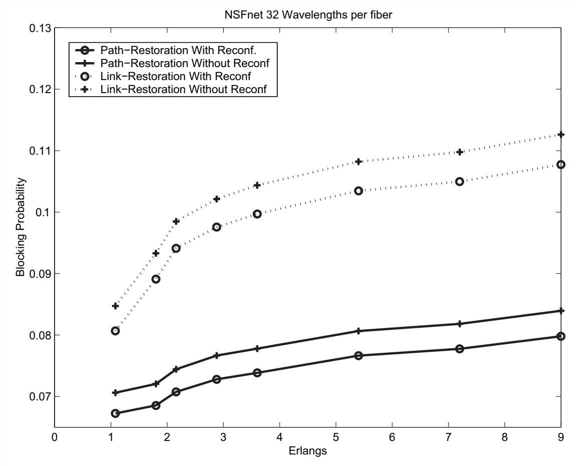

Fig. 2. Blocking probability versus load.

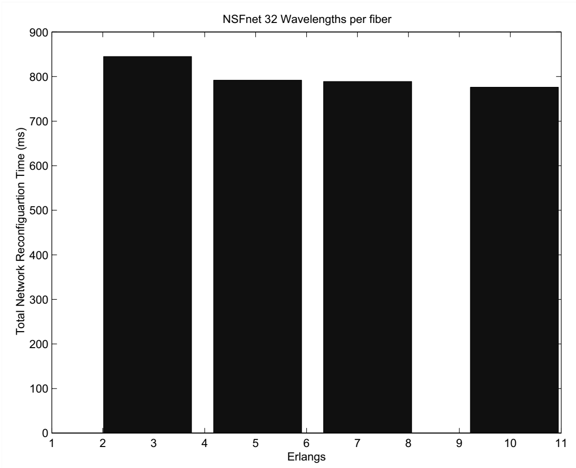

Fig. 3. Network reconfiguration times.

minimal resource contention, with less than 1% of the re-routing attempts experiencing contention. Note that token based reconfiguration may still experience some contention in the case of multiple simultaneous fault repair events (each of which will use separate MST walk sequences).

Further, the operating timescales of the proposed scheme are shown in Fig. 3. Here, the total reconfiguration time is measured as the time it takes for a token to visit all the nodes and includes all nodal processing/reconfiguration times. These results show that the total global times for NSFNET topology stay well below 1 second (850 ms to be exact), and do not vary widely with input loadings. Finally, simulation results show that the proposed reconfiguration scheme adds minimal additional control bandwidth overhead, e.g., under 2% increase. In all, these findings confirm the viability of the proposed reconfiguration scheme for improving long-term resource provisioning in optical networks.

IV. Conclusion

Distributed reconfiguration provides an important means of increasing resource efficiency in optical networks. Here, we propose a novel scheme to achieve graceful post-repair connection re-routing in order to free network capacity and help accommodate more requests. The scheme uses a token passing framework with a pre-determined MST-based walk path. The proposed approach is evaluated using detailed simulation studies and is shown to yield notable gains in resource efficiency and minimal reconfiguration contention/overhead. Moreover, related re-provisioning timescales fall well within the expected range for reconfiguration-type schemes, indicating good practical applicability.

References

[1] N. Ghani *et al.*, "On IP over WDM integration," *IEEE Commun. Mag.*, vol. 38, pp. 72-84, March 2000.

[2] J. Wang *et al.*, "Path vs. subpath vs. link restoration for fault management in IP-over-WDM networks: performance comparisons using GMPLS control signaling," *IEEE Commun. Mag.*, vol. 40, pp. 80-87, November 2002.

[3] E. Wong *et al.*, "Analysis of rerouting in circuit–switched networks," *IEEE/ACM Trans. Networking*, vol. 6, pp. 419-427, June 2000.

[4] S. Skiena, *The Algorithm Design Manual Book*. Springer-Verlag Publishers, 1997.

[5] A. Shami *et al.*, "Connection management protocols for dynamic lightpath provisioning in future WDM networks," *Photonic Network Communications*, vol. 6, pp. 25-32, July 2003.