Optical Implementation of Resilient Packet Rings Using Light-trails

Ashwin Gumaste and Si Qing Zheng
Advanced Computer Network and Architecture Laboratory
The University of Texas, Dallas
Email: ashwing@ieee.org

Abstract: *The growth of RPR in metropolitan optical networks, has been due to the requirement of sub-wavelength granularity and dynamic provisioning between nodes in a 2-ringlet optical network. The solution of creating opto-electro-opto conversion at each node and then processing individually the information being sent on each wavelength leads to expensive equipment cost as well as issues such as strict synchronization and limited bit-rate. Light-trails, on the other hand provides a solution that optically is able to broker the bandwidth between multiple nodes on the same wavelength leading to low cost implementations, while guaranteeing the dynamic bandwidth requirements that are a necessary prelude to provisioning revenue bearing services. A light-trail is a generalization of a lightpath [5], such that multiple nodes on the trail can take part in uni-directional communication within the trail without the need for optical switch reconfiguration. Light-trails can be pragmatically implemented [1] and provide for sub-wavelength optical communication by creating time differentiated flows between individual nodes. In this paper we study the implementation of light-trails to facilitate RPR. Through a simulation study we identify the constraints that affect a light-trail network and how it performs as compared to an RPR network. The intuitive benefit in cost savings using a light-trail network is somewhat offset by the need for aggregation of traffic at the data layer (like RPR) as well as the wastage of bandwidth in the broadcast medium (light-trail).*

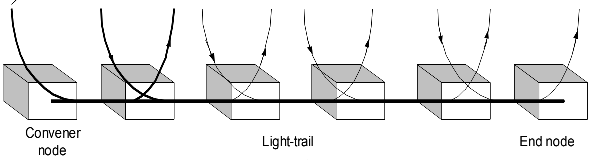

Fig. 1. The concept of multi-point lightpath or simply, a light-trail

I. INTRODUCTION

The surge of IP centric traffic the world over propelled by the rise of the Internet has led to ubiquity of services and an overall explosion of data communications. The distribution of services creates an uncertainty in provisioning bandwidth on a per-demand basis. The optical fiber represents the basis of physical media to meet this huge bandwidth surge. This has led to rise of optical networking in access and metro area in addition to the existing submarine networks. However networking peers in the access and metro area face a severe time-defined uncertainty in determining when and how much bandwidth should be provisioned. This uncertainty is further pronounced by latency and jitter sensitive services. The present form of networks thus resort to a simple over-provisioning model to provision existing services. This form of networking however leads to high network deployment cost. Good statistical multiplexing is the key to being able to provide optimal provisioning and hence obtain low network deployment costs. Legacy networks based on SONET/SDH technology though provide good statistical multiplexing are not efficient for bandwidth-on-demand kind of services typically that are IP-centric. The Resilient Packet Ring (RPR) standard, IEEE 802.17 [11] was proposed to allow dynamic on-demand bandwidth allocation to user stations on an optical ring network. The standard made good use of the optical protection features of providing 50 ms resilience. However the RPR system though able to provide bandwidth-on-demand, requires opto-electronic conversion and regeneration at each node. This makes the system expensive as well as limits the performance in terms of bit-rate. RPR hence poses as a restrictive system in terms of capability to upgrade to higher line rates as well as poses severe cost issues in deployment. The optical backbone on the other hand juxtaposes itself as a high bandwidth medium that provides a means for provisioning huge bandwidth connections. To lower deployment costs and to provide adequate dynamic provisioning (on-demand) we need an optical system that all-optically and using a low cost infrastructure is able to meet the next generation emerging bandwidth and service trends. The light-trail system is then one such network infrastructure that can meet the requirements of next generation networks. The light-trail system is described in [1-4, 8]. In this article we consider light-trails as an alternative for implementing RPR networks. Needless to say that the optical aspect of light-trails makes it unable to meet all the rich features of RPR, but the light-trail system is still able to meet most of the features that are needed for next generation networks.

Section II describes the light-trail solution, Section III proposed light-trails for RPR. Section IV discusses algorithmic issues for light-trails to be able to cater to RPR type provisioning, while Section V discusses results of this paper.

II. THE LIGHT-TRAIL PARADIGM

A *light-trail* is a generalization of a *lightpath* (optical circuit) in which data can be inserted or removed at any node along the path. Light-trails are a group of linearly connected nodes capable of achieving dynamic provisioning in an optical path through an out-of-band control channel (overlaid protocol). This leads to multiple source-destination pairs being able to establish time differentiated connections over the path while eliminating the need for high-speed switching. A light-trail is characterized by a linear segment of nodes that allow communication in one direction. A node in a light-trail employs the feature of drop-and-continue architecture. In addition to the drop-and-continue philosophy, a node is also able to add data into the light-trail without switching. This allows nodes to communicate to one another through time-non-overlapping connections and without optical switching. The switch-less aspect makes a light-trail analogous to an optical bus, but a light-trail due to its out-of-band protocol enhances the known properties of an optical bus. Shown in [1,3,4] are methods to implement light-trails in optical networks for both ring and mesh topologies.

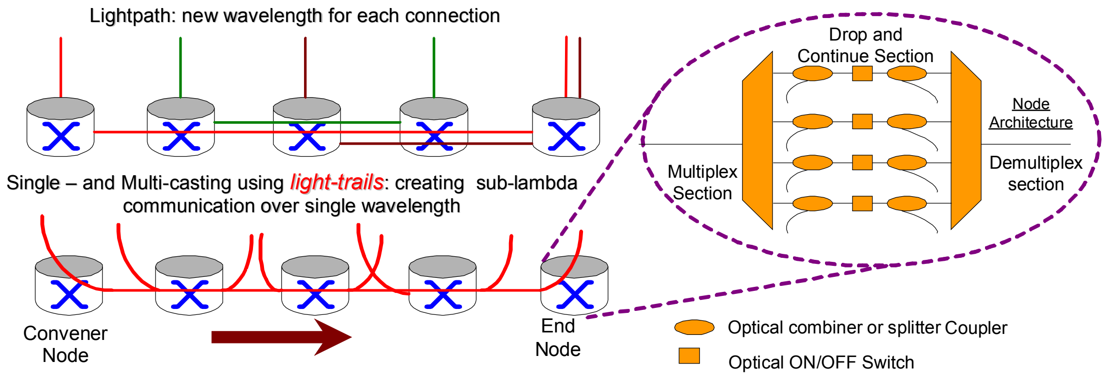

Fig. 2. Conceptual difference between a lightpath and a light-trail and architecture of a light-tralil node

Shown in Fig. 2 is the conceptual difference between a lightpath and a light-trail. The first node in a light-trail is called the *convener node*, while the last node is called the *end node*. The light-trial which essentially resides on a wavelength is optically switched between these two nodes. Multiple light-trails can use the same wavelength as long as the wavelengths do not overlap, thereby leading to spatial reuse of the wavelength. Light-trails present a suitable solution for traffic grooming. It was shown in [3] and [4] that multiple nodes can share an opened wavelength in an optimum way so as to maximize the wavelength's utilization. In this paper we will further strengthen those results to make light-trails more pragmatic – to enable light-trails to create RPR like provisioning model.

The control channel has two primary functions:

(1) creation and deletion of light-trails (*macro-management*) and,

(2) creation and deletion of connections within light-trails (*micro-management*).

The macro-management function of the control channel is responsible for setting up, tearing down, and dimensioning of light-trails. Dimensioning of light-trails means growing or shrinking light-trails in order to meet the requirements of a virtual embedded topology. Macro-management involves switching of a wavelength at the convener and end nodes in order to create the optical bus. Macro-management is a simple procedure and is somewhat static in time hence seldom used. Micro-management on the other hand is more dynamic. It is invoked whenever two nodes communicate to one another using an existing or preset light-trail. Hence this procedure does not require switching. Through micro-management connections can be set up/torn down or QoS needs can be met as desired purely using software control. The overlaid control layer actively supports both forms of light-trail management. Nodes arbitrate bandwidth through the control layer. In this article a scheme for bandwidth arbitration for RPR type traffic in light-trails at the optical layer will also be discussed. Since at a given time, only one connection can reside in the light-trail, the chosen connection must meet requirements of fairness by allowing other nodes to take part in a timely and fair manner.

What makes light-trails so unique is their ability to meet the emerging demands such as optical RPR due to rich features such as optical multicasting and dynamic provisioning, while maintaining low implementation cost. Moreover the light-trail system being completely optical based is oblivious to line rates, hence allows for multi-rate transmission. This becomes particularly attractive with tunable lasers and multi-rate receivers becoming fast popular in service provider networks. Besides these, the light-trail solution provides an *opportunistic* mechanism that couples the data and optical layer through a *control scheme*. This control scheme can be implemented in several ways and two such implementations are discussed in [5,9]. It is the control software that couples the two layers together, but this cannot happen without a hardware that 'allows' itself to be configured. The combination of the light-trails solution – hardware and the software – creates a dynamically provisionable network. This combination potentially solves the uncertainty equilibrium between switching and transport layers by optimized provisioning (provides bandwidth whenever needed). If we compare the light-trail solution to a solution consisting of WDM add-drop multiplexers and overlaid control, then the latter would not be able to provide the necessary dynamism. The obvious hindrance would be in-line optical switching which is somewhat slow (MEMS being the most prolific in today's service provider networks) and suffers from impairments such as cross-talk and extinction ratio. Besides the switching another hindrance in conventional schemes is the requirement of signaling. However this is clearly and cleanly defined in light-trails. The light-trail node architecture however removes these obstacles by deploying the drop-and-continue methodology. It then provides for the ability to provision connections (micro-management) using pure software (signaling) methods, eliminating optical switching all together from micro-management feature of light-trails.

The light-trail system presents itself as an opportunistic medium for nodes that reside on a trail. Such a system allows nodes to 'pitch' in their data whenever possible in the best possible trail or therein without switching. The dynamic nature of communication within a light-trail fathoms a need for optical components such as lasers and detectors that can be switched ON and OFF dynamically. While these 'burst-mode' technologies have reasonably matured [2], the light-trail system (along with PON) effectively uses such technologies very effectively. Burst mode transmitters and receivers that enable dynamic communication carry out the function of micro management in light-trails – setting up and tearing down connections as desired. The maturity of these technologies shown by their prominence in consumer centric markets like PON also means that there is not much of a cost difference viz. conventional Continuous Wave (CW) lasers and detectors. Reference [2-3] gives the reader a more detailed explanation of the light-trail optical performance as well as an in-depth analysis of the technologies used.

III. LIGHT-TRAILS AND RPR

Resilient Packet Ring Standard IEEE 802.17 is defined for ring networks with OEO stations that enable stations or nodes to send packet trains through the network. The RPR standard from a conceptual level has two key principles – optimal statistical multiplexing i.e. grooming, and spatial reuse protocol. To implement RPR optically, we propose light-trails with some modifications. The light-trail system presents itself as an opportunistic medium such that the utilization of the medium is always maximized subject to load. What this means is that the light-trail system allows for the implementation of near optimal grooming strategies leading to very good (tight) statistical multiplexing. While doing so, the light-trail system is absolved of any need for optical switching or electronic processing. In this way, the light-trail system is lower in cost as the electronics is relegated to the edge of the ring and not 'in-line'. However, a light-trail by itself does not support spatial reuse. In a light-trail of n nodes we can have up to n(n-1)/2 connections between the n(n-1)/2 source-destination pairs possible. However, no two connections can co-exist even if they are graphically apart in the same light-trail. This is due to the broadcast nature of the optical bus like properties exhibited by the light-trail system. From a network perspective the light-trail system however does exhibit spatial reuse – that of two or more light-trails on the same wavelength as long as they are graphically disjoint. If we then compare the light-trail system to the RPR system, then obviously the RPR system has better spatial reuse properties as the granularity exhibited through reuse is that of a packet slot. The light-trail system in contrast, exhibits spatial reuse that is much defined by the virtual topology of light-trails that are mapped on to the WDM ring network. So from an absolute comparison perspective, a light-trail system provides for the higher degree of statistical multiplexing and a somewhat lesser degree of spatial reuse. The higher degree of statistical multiplexing is brought about by the ability to provision the light-trail on a per-demand basis. This absolves the need for slotted system and hence caters to bursty traffic [7] flows seamlessly. However what makes light-trails still effective is their overall deployment cost. The architecture for light-trails shown in Fig. 1 details the simplicity of implementation as well as allows a light-trail to function at high bit-rates 10 Gb/s and even beyond only to be limited by the transmission characteristics of the optical media (pass-band narrowing etc.). The RPR system however is severely impaired by the electronic processing speeds. In the pursuit to provide fine granularity the RPR system compromises on costs. It requires high speed electronics and tight slot synchronization for effective performance. The light-trail system does not need any data channel synchronization for providing dynamic provisioning of granularities which are very comparable to RPR. Further, while fairness is a virtue that most shared mediums desire, the fairness allocation in RPR is somewhat complicated. The light-trail control channel is open to any fairness allocation algorithms. We have for example implemented auction based fairness algorithms for light-trail networks that are extremely fair in terms of bandwidth allocation. In this paper to show light-trails for RPR we have also implemented a simple equally fair scheme in which for an n node light-trail any k nodes can be pre-selected at any time and have equal transmission rights of the channel.

IV. ALGORITHMS FOR IMPLEMENTING RPR TYPE FUNCTIONALITY OVER LIGHT-TRAIL NETWORKS

In this section we shall show how the light-trail system can be adapted for RPR like performance. Consider a ring network of R nodes and let the ring be defined by a 2-fiber WDM system with each fiber ring-let transmitting in a direction of communication that is opposite from the other. Let us now define a light-trail pair of n nodes such that each of the two counter propagating ringlets has a light-trail of n nodes in it. The convener node of the light-trail in the CW ring corresponds to the end node of the light-trail in the CCW ring and vice versa. We shall now show a simple algorithm for implementing fixed slot size communication in these two light-trails.

In generic light-trail networks as explained in [1-4] whenever a source node desires to establish a connection with some destination node through an existing light-trail it does so through micro-management. Here the source node listens to its upstream neighbors and once it ascertains the channel to be free it sends a control packet to *all* its downstream neighbors in the light-trail. This control packet is called communication control packet. It then waits for a finite amount of time defined as the connection provisioning time and then sends its data through the light-trail. In course of transmission if some upstream neighbor of this incumbent source node requests for channel bandwidth, then the incumbent source halts its transmission and allows the upstream neighbor to establish its connection. In this manner the light-trail system is not fair – nodes that are closer to the convener get a higher share of the network bandwidth as compared to the nodes that are closer to the end node. Fairness in generic light-trail networks decreases from the convener to the end node.

We however propose a simple alteration of this algorithm to implement optical RPR. The alteration involves a very minor hardware change at the control layer. Since in WDM ring networks we always assume a 2-fiber ring (2 fibers due to resiliency), let us also assume that the unidirectional light-trail is supported by a bi-directional control channel – one in each fiber. Further let us assume that for any given light-trail the end node acts as universal light-trail arbiter, i.e. the end node decides transmission times for each of the n-1 light-trail nodes as well as transmission duration. Our assumption has good reason because if the light-trail is in CW direction then the connection requests reach the end node through the CW control channel. Likewise the end-node can communicate back to any of the light-trail nodes indicating connection grants etc through the CCW control channel. The end node thus can run a fairness algorithm that is centralized, and ensure that nodes within a light-trail receive equal fairness or receive fairness according to some service level agreement that a node is running. The dual control channel then provides as a simple mechanism but has one serious disadvantage: it exemplifies a *busy* system. The system becomes very complicated with a lot of back and forth interactions between nodes and the end node varying linearly with the decrease in granularity. That is, for fine granularity of bandwidth provisioning the interactions between light-trail nodes and the end node are intense. We hence propose a simpler modification of the same scheme. In our proposed scheme we assume that a node is intimated of the service level agreements it would have to cater to slightly ahead of the time (typically a few milliseconds) before the service is actually provisioned. Once a SLA request arrives at a node, the node selects a light-trail from the preset set of existing light-trails, or establishes a new light-trail depending on factors such as location of the destination node(s), fairness values of existing light-trails, delay profiles of existing light-trails and SLA latency and jitter needs. Assuming that it finds an existing light-trail say L, it then sends a request to the end node of the existing light-trail request for addition of its SLA into the already provisioned SLAs of the existing light-trail. In doing so the node intimates the end-node of the SLA details such as latency needs, transmission bandwidth (flow) and jitter. The end node then examines the SLAs provisioned for nodes in L and checks to see if the new required SLA can be provisioned in L. In doing so, it has to ensure two factors:

1. The summation of all the sub-lambda time-averaged flows from all the nodes in L has to be lesser than C, the bit-rate exhibited at the optical layer.

2. The SLA quality levels such as latency and jitter can be met by the light-trail L.

The first factor is fairly simple to understand from the law of conservation of energy or a flow model. The total incoming flows into the light-trail have to be lesser than the total available light-trail bandwidth else some of the flows would never be provisioned leading to high blocking irrespective of what provisioning algorithm is deployed. The second factor is of a critical nature. When we map flows into a light-trail trying to make it exhibit RPR like properties, we provide somewhat loose transmission slots to each of the nodes. Slots for each nodes are interleaved in time, somewhat SONET like. However the slots we design are far simpler than SONET payloads and easier to operate due to their dependence on optics rather than electronics for processing.

The end node which receives a request for the new SLA, then computes the slot duration for transmission for each of the nodes in L and communicates back to nodes in L through the reverse control channel (reverse because this time it is opposite to the data flow). The nodes then change their transmission durations as well as connection provisioning times thereby accommodating the new SLA. The control channel which runs through each node in this case is assumed to be synchronous. The data channel is not synchronized to avoid high costs and to facilitate both synchronous and asynchronous services to be provisioned.

V. FUNCTIONING OF THE ALGORITHM:

V.A. Light-trails for RPR

Assume the same light-trail L having *n* nodes as in the previous section. Let k of these n nodes have some SLAs running at their client interfaces with the destinations in the network part of the same light-trail (downstream). In this section we shall examine how the light-trail performs RPR like time-interleaved, non-synchronous, bursty communication. There are two kinds of bandwidth allocations that we shall consider – equal fairness to each of the k-contending nodes and on-demand fairness to nodes on demand. The second case is a special case of the first one, whereby we can simply scale nodes with higher fairness levels averaged over nodes that have a lower fairness level. Let us now consider the method to provision an *n*-node light-trail with *k* of these *n* nodes wanting to equally share the channel. For this system let us assume a buffer size of B (bits) and a light-trail transmission rate of C bits/sec. Our desire then is to find the relation between the buffer size B, the number of contending nodes k and the average slot duration T_{slot} that a node is given. Further, we also are interested in finding out the maximum latency in the network D_k.

The critical value of T_{slot} can be found through the following discussion on connection provisioning time. We assume that once a node is given transmission rights of the channel it would be given transmission time proportional to the amount of time it requires to empty its buffer, regardless of whether the buffer is full or not. Further a node gets transmission rights every k-1 slots, each of duration T_{slot}. Now further if we assume a super-slot T, of duration T_{slot} + connection provisioning time, then this super-slot T should be equal to the buffer size B divided by the line speed C over k nodes.

The connection provisioning time is dependent on the propagation delay of the control packets, the control packet processing time, the time to switch ON/OFF a laser and finally the burst mode receiver locking time. We assume components that are contemporary and available off-the-shelf. Hence it is possible to obtain better performance with next generation components with somewhat lesser technological maturity.

We shall now calculate connection provisioning time of the light-trail as follows:

Let T_{pd} be the propagation delay between any two adjacent nodes.

Let T_{on} be the laser turn ON/OFF time.

Let T_{bs} be the burst mode receiver lock on time.

Let T_{slot} be the time allocated for transmission to a particular node under fairness constraints that each node has equal transmission rights.

Let B be the buffer size at a node allocated for any light-trail. B is measured in bits.

Let C be the line speed of light-trail in bits/sec.

For a given light-trail L, let k<L be the number of nodes (out of the L nodes in a light-trail) that are 'active', i.e. contending for transmission rights of the light-trail channel.

Our aim is to calculate the relation between B and T_{slot} for a given k and L and establish a recurrence relation that generically enables network operators to provision the light-trail.

From the above convention, the time required to set up a connection is calculated as:

$$T_{connection} = \{(k-1)(T_{pd}) + T_{BS} + 2T_{ON}\} \tag{1}$$

Further the slot duration at each of the k 'contending' nodes in an N node light-trail with a buffer of size B is:

$$T_{slot} = \left[\frac{B}{KC} - \{(k-1)(T_{pd}) + T_{BS} + 2T_{ON}\}\right] \tag{2}$$

From (2) we can obtain the end to end queueing delay cross a light-trail as:

$$D_k = \left(T_{slot} + T_{connection}\right)(k-1) \tag{3}$$

The end-to-end queueing delay has specific significance for light-trail networks especially when catered to RPR kind of traffic. To provision services over a light-trail network the end-to-end delay is particularly important because it determines the maximum amount of time a packet of a particular service type would spend through the light-trail system. Our calculation (in the next section) shows that the light-trail system produces latency that is acceptable for most service types. Further, there is a strong relation between maximum latency experienced through the light-trail system and the utilization or efficiency of the light-trail. We calculate utilization or efficiency of the light-trail system as the ratio of the time the light-trail is busy in transmitting data to the sum of time it transmits data as well as provisions connections. Then, it can easily be understood that as long as the arrival rate at a node satisfies the relation in (4) the latency can easily be reduced. In (4) is shown a relation that explains how the arrival rate of packets is related to the buffer size:

$$\left(\frac{B}{C} + T_{connection}\right)(k-1) = \lambda \tag{4}$$

Where λ is the arrival rate of packets into the buffer.

V. B Allocating Bandwidth on Demand to Nodes in a Light-trail

In this section we will consider how to dynamically allocate bandwidth to nodes and yet maintain good statistical multiplexing properties. For dynamic allocation the light-trail medium is an ideal system for adaptation. Multiple algorithms have been proposed for dynamic light-trail allocations and are beyond the scope of this article. However, in our view an algorithm that is highly parallel in nature suites the application of providing dynamic bandwidth allocation. One such algorithm that becomes a natural contender is the auction algorithm [10]. The auction algorithm due to its greedy, opportunistic and immensely parallel character can provide a wide framework for dynamic allocation in light-trails.

To adapt a dynamic allocation algorithm with the constraint of still being able to provide fairness requires the light-trail framework to be somewhat re-engineered. In this re-engineering process we assume the light-trail system to have a bi-directional control channel that essentially allows bi-directional control messages to be exchanged between the nodes of a light-trail.

However the requirement of an auction algorithm is the need for a central arbiter. For this central arbiter we assume the end node of a given light-trail. The reason we choose the end-node is because the end node has access to all the forward messages that any node sends to its downstream neighbor.

The working of the auction algorithm is described in [11]. The nodes poll to the end node their buffer or SLA states. The end node runs the auction algorithm with some user-defined alterations. It then allocates bandwidth to nodes that need them the *most*. The allocated node then establishes connections while all other nodes keep vying for the underlying bandwidth by continuing to poll their buffer and SLA states to the end node, with a view of securing the light-trail bandwidth. The node which wins the auction and establishes its connection does so without optical switching meaning absolving the need for high provisioning time.

VI. SOME RESULTS

To validate our hypothesis that light-trails serve as an effective alternative to RPR networks we conducted a simulations study. In this simulations study our primary aim was to evaluate light-trail performance from network parameters like delay, buffer size and utilization.

We consider a 10 node light-trail and vary the number of contending nodes from 1 to 8. In Figures 3a through 3c we observe the T_{slot} variation for different buffer sizes. We make the following assumptions: the connection set up time is 100 micro-seconds comprising mainly of propagation delay. Note that propagation delay can be neglected if we have a highly parallelized control algorithm. Burst mode transceivers are known to lock on in 400 ns (spec) but we assume 1 micro-second for convenience. In Fig.3a we observe T_{slot} for 256, 128 and 64 Mb buffer sizes. The line rate is fixed at 1 Gb/s. It can be seen that the T_{slot} is somewhat same for 4-7 nodes, which represents a typical light-trail system. However when there are only 1-2 nodes contending in the light-trail system then T_{slot} is much higher. The same results are obtained for 32 and 8 Mb buffers in Fig 3b. In Fig. 3c we simulate for 4 Mb and 2 Mb buffers. We observe T_{slot} to be significantly low. However so is the end-to-end delay or latency as shown in Fig. 4a to 4c. The end-to-end latency for light-trail networks varies with buffer size. We see that for smaller size buffers the end-to-end delay is lower because the time required to empty out the buffer is lesser than that for larger buffer sizes. Intuitively it can be understood that the network utilization or efficiency is quite opposite to the results obtained in Fig. 3 and Fig. 4. In Fig. 5 we plot the utilization (light-trail) curve as a function of number of nodes contending for bandwidth of the 10-node light-trail for different buffer sizes. What is critically important is that as the buffer size increases so does the utilization. The reason is straight-forward. For large buffer sizes the time required to empty out the buffer is large, signifying T_{slot} to be large. Connection provisioning being constant then ensures the efficiency to be high. For smaller buffer sizes however as shown in Fig. 5b and Fig. 5c. the utilization is not much lesser than that seen for larger buffers. But correspondingly as can be seen in Fig. 4b, 4c and Fig. 3b 3c, the end-to-end delay as well as T_{slot} size is small. The lower latency ensures the ability to provision next generation services while still deploying a shared wavelength medium, void of any optical switching. This is the fundamental contribution of light-trail networks as an alternative to optical RPRs, in being able to meet the dynamic bandwidth guarantees as well as provide a platform for shared access of the channel to multiple nodes in a low-latency and highly resilient [12] system.

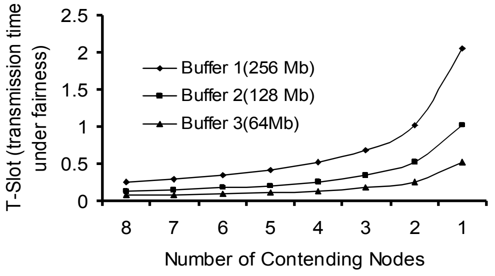

Fig. 3a

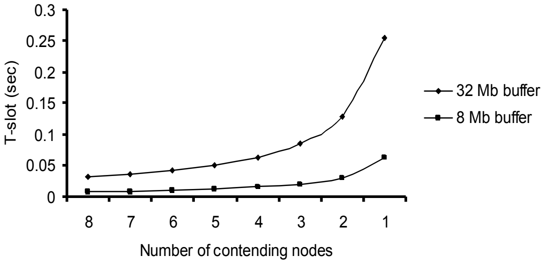

Fig. 3b

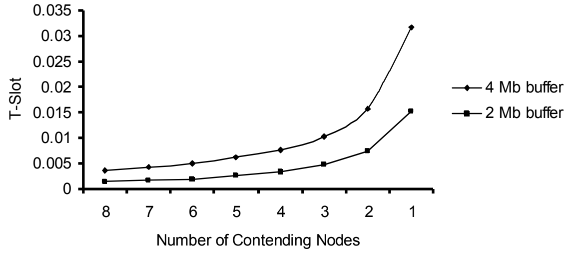

Fig. 3c.

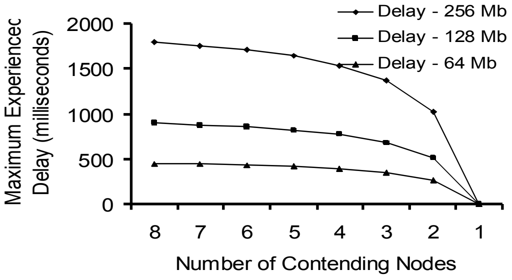

Fig. 4a

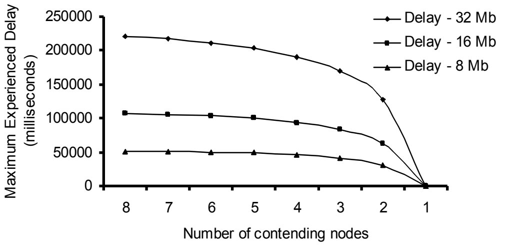

Fig. 4b

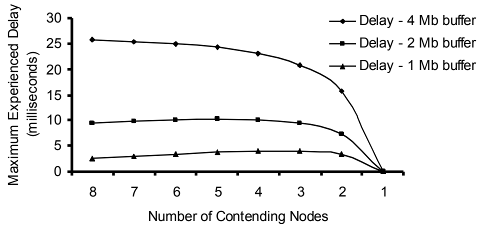

Fig. 4c

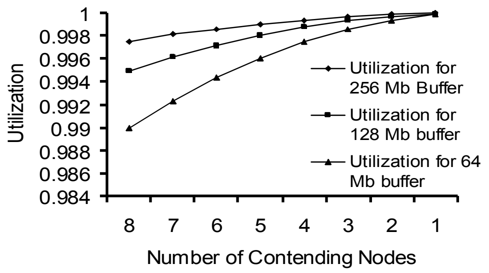

Fig. 5a

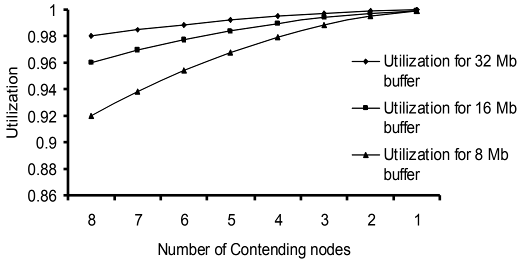

Fig. 5b

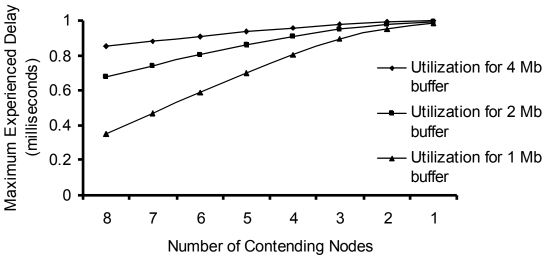

Fig. 5c

VII. CONCLUSION

In this paper we have proposed a model for optical RPR using our previously proposed concept of light-trails. The bus-like properties of light-trails coupled with a control channel produce a low-cost and highly dynamic system that enables bandwidth on demand provisioning much needed for RPR type networks. However, the optical nature of light-trails creates a paradigm shift from electronic processing and hence lowers the infrastructure cost. In addition the light-trail system allows for dynamic provisioning as well as optical multicasting – much needed features of next generation dynamic optical networks. The light-trail system eventually poses as an alternative to RPR in terms of both technological as well as economical enhancement.

REFERENCES

[1] A. Gumaste and I. Chlamtac, “Light-trails: A Novel Solution for IP Centric Communication in the Optical Domain,” *IEEE Proc. of the High Performance Switching and Routing (HPSR)* Torino 2003, Italy

[2] A. Gumaste and I. Chlamtac, “Light-trails: A Solution for Dynamic Optical Communications,” *OSA Journal of Optical Networking* (JON) May 2004, *invited paper*.

[3] A. Gumaste, G. Kuper and I. Chlamtac, “Light-trail Assignment to Optical Mesh Networks,” *IEEE 13th LANMAN* San Francisco April 2004

[4] A. Gumaste and I. Chlamtac, “Mesh Implementation of Light-trails: A Solution to IP Centric Communication,” 12th *Proc of IEEE International Conference on Computers, Communication and Networks (ICCCN)* Dallas TX, 2003

[5] I. Chlamtac, A. Ganz, and G. Karmi, “Lightpath Communcication: A Novel Way to Increasing Capacity of Optical WANs,” *IEEE Trans. On Commun*. Mar 1992 pp1152

[6] A. Gumaste and S. Q Zheng, “Optimizing Light-trail Assignment for WDM Ring Networks,” submitted

[7] C. Qiao et al, “Labelled Optical Burst Switching LOBS: IP over WDM Networks, ” *IEEE Communication Magazine.* Sept 2001

[8] A. Gumaste and I. Chlamtac, “Adaptations to a GMPLS Framework: the Light-trail Approach,” *Proc of National Fiber Optic Engineers Conference, NFOEC*, Orlando 2003

[9] A. Gumaste and S. Q. Zheng, “Network Level Dissemination of Light-Frame Communication,” UT-Dallas Tech Report UTDCS-14-04, April 2004.

[10]. Bertsekas, *Auction Algorithm*, MIT, 1978

[11]. RPR Alliance

[12] A. Gumaste and S. Q Zheng, “Protection Scheme for Light-trail Networks,” Submitted under review.