DynaSPOT: Dynamic Services Provisioned Optical Transport Test-Bed—Achieving Multirate Multiservice Dynamic Provisioning Using Strongly Connected Light-Trail (SLiT) Technology

Ashwin Gumaste, *Member, IEEE*, Nasir Ghani, *Senior Member, IEEE*, Paresh Bafna, Akhil Lodha, Anuj Agrawal, Tamal Das, Jianping Wang, *Member, IEEE*, and Si Qing Zheng, *Senior Member, IEEE*

***Abstract*—In this paper, we report on the dynamic services provisioned optical transport (DynaSPOT) test-bed—a next-generation metro ring architecture that facilitates provisioning of emerging services such as Triple Play, Video-on-Demand (VoD), pseudowire edge-to-edge emulation (PWE3), IPTV, and Data Center Storage traffic. The test-bed is based on the recently proposed strongly connected light-trail (SLiT) technology that enables the triple features of dynamic provisioning, spatial subwavelength grooming and optical multicasting—that are quintessential for provisioning of the aforementioned emerging services. SLiT technology entails the use of a bidirectional optical wavelength bus that is time-shared by nodes through an out-of-band control channel. To do so, the nodes in a SLiT exhibit architectural properties that facilitate bus function. These properties at the network side include ability to support the dual signal flow of drop and continue as well as passive add, while at the client side include the ability to store data in order to support time-shared access. The latter (client side) improvisation is done through a new type of transponder card—called the *trailponder* that provides for (electronic) storage of data and fast transmission (burst-mode) onto the SLiT. Further in order to efficiently provision services over the SLiT, there is a need for an efficient algorithm that facilitates meeting of service requirements. To meet service requirements we propose a dynamic bandwidth allocation algorithm that allocates data time-slots to nodes based on a *valuation* method. The valuation method is principally based on an auctioning scheme whereby nodes send their valuations (bids) and a controller node responds to bids by sending a grant message. The auctioning occurs in the control layer, out-of-band and ahead in time. The novelty of the algorithm is the ability to take into consideration the dual service requirements of bandwidth request, as well as delay sensitivity. At the hardware level, implementation is complex—as our trailponders are layer-2 devices that have limited service differentiation capability. Here, we propose a dual VLAN tag and GFP-based unique approach that is used for providing service differentiation at layer-2. Another innovation in our test-bed is the ability to support multispeed traffic. While some nodes function at 1 Gb/s, and others function at 2.5 Gb/s (using corresponding receivers), a select few nodes can support both 1- and 2.5-Gb/s operation. This novel multispeed support coalesced with the formerly mentioned multiservice support is a much needed boost for services in the metro networks. We showcase the test-bed and associated results, as well as descriptions of hardware subsystems.**

***Index Terms*—Dynamic bandwidth provisioning, light-trail, metro networks.**

I. INTRODUCTION

THE shift of revenues from traditional voice services to a multitude of VoIP, Video-on-Demand (VoD), Pseudowire Edge-to-Edge Emulation (PWE3), Triple play, and Data-Center storage traffics is a strong motivation for new, low-cost, and dynamic optical layer solutions in metro environments. Conventional SONET/SDH hierarchy is now being replaced by more data-centric packet aware technologies like GigE lightpaths and resilient packet rings (RPR) with IP/MPLS overlay. GigE is not efficient (on account of its requirement of end-to-end wavelength granularity) nor is it dynamic, while RPR is expensive (due to OE and EO conversions at every node). A new approach is required that provides efficient grooming of subwavelength traffic preferably at the optical-layer while enabling necessary dynamic bandwidth provisioning thus facilitating emerging services.

We report a solution that enables subwavelength all-optical (spatial) grooming, multiservice support at multiple line-rates using mature and available technology. The proposed solution is deployed through a metro ring WDM test-bed called dynamic services provisioned optical transport (DynaSPOT), and is built on the concept of strongly connected light-trail (SLiT) technology [1]–[3].

A light-trail [4], [5] is a unidirectional optical bus that is provisioned through an out-of-band (OOB) control channel facilitating spatial subwavelength [5] grooming of traffic along multiple nodes.

Manuscript received July 9, 2007; revised November 5, 2007.

A. Gumaste, P. Bafna, A. Lodha, and A. Agrawal are with the Department of Computer Science and Engineering, Indian Institute of Technology, Bombay, Powai, Bombay 400076, India (e-mail: ashwing@ieee.org; pareshpbafna@gmail.com; alodha@cse.iitb.ac.in).

N. Ghani is with the Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM 87106 USA (e-mail: nghani@ece.unm.edu).

T. Das is with the Department of Mathematics, Indian Institute of Technology, Delhi, India (e-mail: dastamal.dinesh@math.iitd.ac.in).

J. Wang is with the Department of Computer Science, City University of Hong Kong, P.R. China (e-mail: jianwang@cityu.edu.hk).

S. Q. Zheng is with the Department of Computer Science, University of Texas at Dallas, Richardson, TX 75080 USA (e-mail: sizheng@utdallas.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2007.913074

A SLiT [1] is a bidirectional implementation of a light-trail with the node architecture first proposed in [1] for implementation in metro rings. SLiT bandwidth is arbitrated by a *controller* and nodes time-share the bandwidth to provision *connections*. Connections are provisioned using burst-mode optics (like in light-trail [5]). In the DynaSPOT test-bed, we show how to dynamically provision these connections with a particular emphasis on services. The time-penalty incurred due to queueing of data while facilitating time-sharing is made proportional to the particular service latency type by a new transponder subsystem—the *trailponder*.

Another salient feature of the DynaSPOT test-bed is the ability to use the same SLiT for connections at multiple line-rates, doing so in a dynamic fashion. This feature enables nodes with diverse transponder inventory to use the same SLiT, thus, reducing capital expenditure and maximizing efficiency. A final focus of the DynaSPOT test-bed is the ability to provision a multitude of services, both delay sensitive and bandwidth intensive.

The rest of the paper is organized as follows: Section II discusses light-trail technology and highlights open issues that are motivations for the SLiT proposal. Section III showcases SLiT technology and associated assumptions. Section IV describes the DynaSPOT test-bed. Section V. showcases results obtained from the test-bed, while Section VI compares these results qualitatively with existing technology solutions. Finally, Section VII summarizes the paper.

II. LIGHT-TRAIL TECHNOLOGY PRIMER—ADVANTAGES AND OPEN ISSUES

In this section, we present a primer on light-trail technology that serves as a preamble to the rest of the paper in particular to the DynaSPOT test-bed. SLiT technology, the subject of investigation in this paper has emerged from our recently proposed concept of light-trails [4], [5].

Light-trails are a generalization of a lightpath (or optical wavelength circuit), such that multiple nodes can take part in communication along the path. A light-trail is analogous to a unidirectional shared wavelength bus. Bandwidth within the bus is shared by arbitration amongst the constituent nodes through an OOB control channel. The control channel is optically dropped, electronically processed and then reinserted back into the network at every node-site [see Fig. 1(a)]. Light-trails have been shown to lead to efficient subwavelength optical grooming (which we define as spatial subwavelength grooming), dynamic provisioning (of bandwidth to nodes), facilitate optical layer multicasting and be built using low-cost and mature technology. To support the above features, light-trail nodes have node architectures that have characteristics to support bus functionality on a per-wavelength basis.

Fig. 1(a) shows a light-trail and its comparison to the well-established lightpaths (point-to-point optical circuits). As is seen, a light-trail is a wavelength bus that is regulated between two extreme nodes, called the convener node and the end node, with the direction of the flow from the convener node to the end node. To set up the light-trail the wavelength is blocked between the two extreme nodes, while the intermediate nodes allow for all-optical pass-through as well as support of the bus functionality. Setting up, tearing down and dimensioning (growing the light-trail) is carried about through optical switch (re)configuration. Since, conventional (typically mechanical) and contemporary optical switches are slow (typically requiring several milliseconds) to change state, we assume that formation (set up), tear down, and dimensioning of the light-trail is infrequent and is semipermanent.

Once a light-trail is set up, nodes can communicate to one another by establishing time-differentiated connections. These connections are short duration transmissions of data over the light-trail. Connections are set up and torn down over the light-trail without any optical switching. The only constraint on the connections is that no two connections can coexist over the light-trail at the same time—which if it would happen, would lead to collisions. To avoid collision and guarantee fairness, nodes communicate to each other and synchronize their transmissions (with respect to each other) through the OOB control channel.

Over an $n-$ node light-trail, a maximum of $n(n-1)/2$ connections are possible. Provisioning of connections, since it does not require any optical switching is called soft-provisioning, while provisioning of the light-trails due to requirement of optical switching is called hard provisioning.

The feature of unidirectional bus functionality in light-trails leads to optical multicasting, while the characteristic of being able to provision connections over a light-trail without optical switching leads to dynamic provisioning. The dynamic provisioning feature results in a property that we term as *spatial subwavelength* support and is now defined: since, multiple nodes time-share the light-trail bandwidth, which in effect means that nodes time-share a wavelength resulting in each node achieving an effective bandwidth that is subwavelength; further since these nodes are spatially separated along the light-trail, this leads to our notion of spatial subwavelength grooming.

We have shown in [4] and [6] the advantages of light-trails over Gigabit ethernet and resilient packet rings, applied light-trails as an effective candidate technology to storage area networking (SAN) [5] and compared light-trails to optical burst switching [4].

Despite the aforementioned advantages showcased by light-trails, there are a few drawbacks that affect performance as well as increase cost. A primary drawback of light-trail technology is that there is uneven per-span utilization in the unidirectional bus. As can be seen from Fig. 1, the convener node in the light-trail can set up $n-1$ connections (to $n-1$ prospective downstream destination nodes), while the second node can set up $n-2$ connections, and so on. This implies that the span between nodes N_2 and N_3 would result in higher utilization than the span between nodes N_1 and N_2, and so on. This unbalanced utilization implies wastage of bandwidth.

A second disadvantage of light-trails is that they are unidirectional and hence they do not naturally support duplex communication. The unbalanced utilization of spans coupled with the associated delay jitter also has led to unfair allocation of bandwidth to nodes and several schemes [7]–[9] have been proposed for bandwidth allocation and fairness within light-trails.

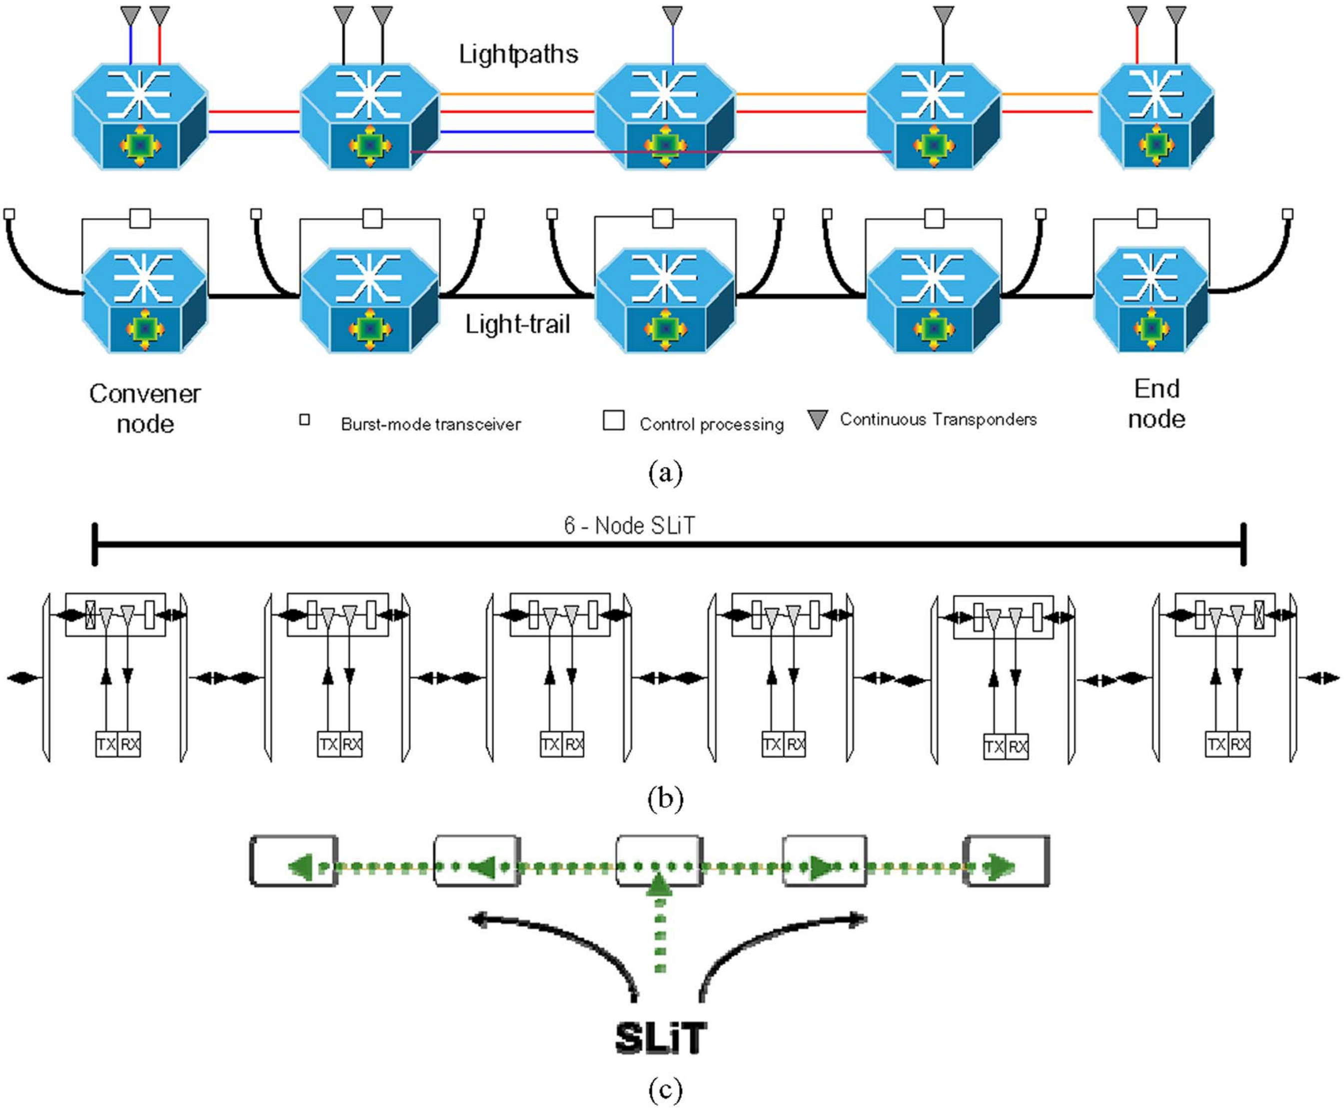

Fig. 1. (a) Comparison of lightpaths and light-trails [10]. (b) A 6-node SLiT. (c) A node transmitting to both East and West directions.

III. SLiT TECHNOLOGY

To alleviate the aforementioned problems in Section II while maintaining the advantages of light-trails, we in [1], extended the basic unidirectional light-trail to the bidirectional SLiT. A SLiT is a bidirectional version of a light-trail that allows communication between nodes over a single wavelength in duplex fashion. Hence, an n-node SLiT is able to support a maximum of n^2 connections (i.e., n^2 source-destination pairs), with the constraint that no two connections can coexist at the same time.

To support duplex communication over a single wavelength, we have proposed in [2] a unique node architecture that facilitates bidirectional support. In addition, in [2] we have proposed a new protocol that guarantees fairness and leads to efficient bandwidth utilization within the SLiT. This protocol, which we have implemented in our DynaSPOT test-bed below and which we will describe in detail in Section IV, takes into consideration service requirements of delay and bandwidth hence facilitating delay sensitive and bandwidth intensive services to be provisioned over the shared wavelength SLiT.

Conceptually, an East-West SLiT is shown in Fig. 1(b), where nodes communicate to other nodes in both directions. Also shown in Fig. 1(c) is an example of a node transmitting in both Eastward, as well as Westward direction. Using the passive optical bus, a node can also receive data from other nodes that are either Eastward or Westward of itself.

A. SLiT Node Architecture

To support the SLiT, a node has an architecture that is shown in Fig. 2(a) and (b). Composite WDM signal from the network enters the node at either of the two Arrayed Waveguides (AWGs) that can act as both multiplexer and demultiplexer depending on whether the signal is entering the node or leaving it. The composite WDM signal is demultiplexed into its constituent wavelengths. Each wavelength is fed to a SLiT optical retrieval section (SORS) that allows the node access to the SLiT signal. The SORS consists of two ON/OFF optical shutters (slow moving optical switches) which are separated by two passive couplers (both in 2×1 configuration). The two couplers are 3-dB (50/50) type, i.e., power at any input port is split in half to the other two output ports.

One of the two couplers is called drop coupler (DC) and the other coupler is called add coupler (AC). Signal that enters

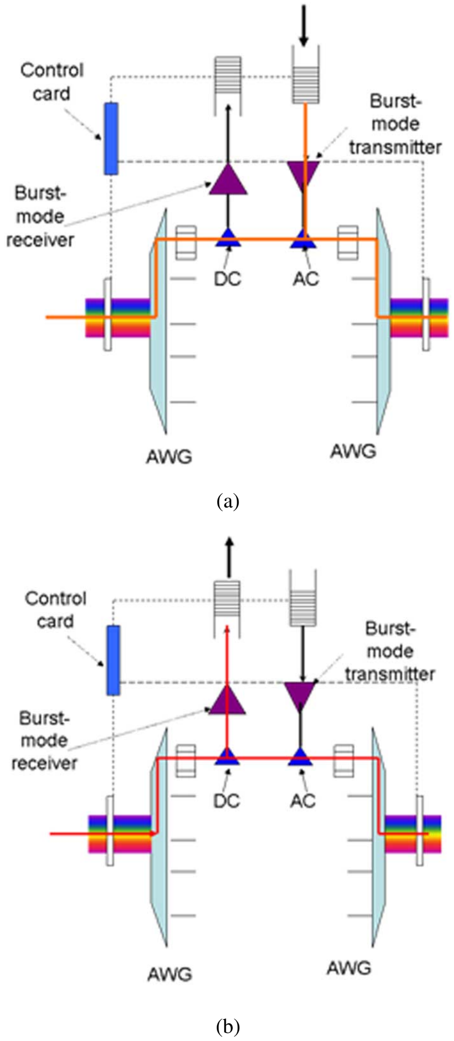

(a)

(b)

Fig. 2. (a) SLiT node architecture supporting passive add function. (b) SLiT node architecture supporting drop and continue function.

the node from either direction (West/East) is dropped for local processing at the drop coupler. The DC also forwards a copy of the signal (using optical splitting property) to nodes further downstream in the SLiT, thus resulting in *drop and continue* operation. The second coupler (AC) allows the node to send in signal into the SLiT. The AC does so using passive properties, i.e., allows addition of signal without any switching operation and this is called as *passive add* operation. Signal sent into the AC is split into two copies, one sent into the Eastward direction and the other sent into the Westward direction, as shown in Fig. 1(c).

Since nodes time-share the SLiT bandwidth, it implies that when a node establishes a connection (over the SLiT), other nodes have to queue their data and wait for transmission rights over the SLiT. A new subsystem is required that enables nodes to queue data and then efficiently transmit (and likewise receive) whenever the protocol allows the node rights to form a connection. We propose a subsystem called a *trailponder* that stores data in a format that supports layer-2 data without affecting layer-2 functionality. The trailponder is similar to a transponder in the function that it transmits and receives signals between the network side (SLiT) and the client-side. However, the trailponder has the additional onus of storing data in a format that allows effective communication (explained in Section IV). Current layer-2 communication protocols are in principle not delay tolerant, meaning that if a destination node does not receive data from a source node for some duration, the destination node would trigger a loss of signal alarm.

A second function of the trailponder is to efficiently utilize the SLiT bandwidth. To do so, the trailponder must be able to send data (with minimal delay), once the node is given rights to form a connection. This means that the transmitter (laser) must be able to switch "ON" in a very short time interval. Conversely, the receiver in the trailponder must also be able to "latch-on" to an incoming signal with minimal preprocessing or training bits. To provide this kind of fast ON/OFF and rapid reception capability, the trailponder is designed using burst-mode optics [11] as shown in the architecture of Fig. 2.

B. *Control Protocol and Timing Issues*

The control protocol for bandwidth arbitration in the SLiT is designed based on an auctioning algorithm. The SLiT bandwidth is assumed to be the object that is being bid for by multiple nodes in the SLiT (that act as bidders). The data channel (SLiT) is assumed to be time-slotted with slots of large duration (typically 0.3-5 ms). It is not necessary to synchronize all the nodes with respect to each other. Slot boundaries are assumed to be only loosely synchronized. Loose synchronization is achieved by the OOB control channel that is optically dropped and electronically processed at every node before being inserted back into the network. Due to this OEO function on the control channel, it is fair to assume (and hence implement) the control channel to be tightly synchronized with respect to all the nodes (sharing common clock). Since the control card at every node is responsible for transmission of data (and hence formation of a connection) over the SLiT, it is assumed that the trailponder, which actually sends and receives data is also pseudosynchronized. The data time-slots are selected large enough so that the auction algorithm can result in a convergent bidding process. In each SLiT, a particular node is selected as the *controller* node. This node acts as an arbiter for bids in the SLiT. The bid is sent as a *valuation* that is a single numerical quantity representing both the bandwidth intensity and delay sensitivity of the data that is stored at the node (in its trailponder). In every data time-slot, each node sends a valuation to the controller node through the control channel. The controller receives valuations from every node in the SLiT and then selects the node that sent the highest valuation. It then sends a *grant* message to this node (that sent the highest valuation). This node then sets up a connection (over the SLiT) in the *next* data time-slot. In this way bandwidth is dynamically allocated to nodes in a SLiT. The process of computing valuations is explained in Section IV-F.

The entire procedure of sending valuations by the nodes to the controller, computing the node with the highest valuation and signaling back to the successful bidding node, is done *ahead-in-time* and through the OOB control channel. We have shown in [16] that the procedure of computing valuations leads to proportional fairness [12].

TABLE I
NODE CONFIGURATION AND SERVICE PROVISIONING IN THE DYNASPOT TEST-BED

Controller

	N_1	N_2	N_3	N_4
Voice	x	x	x	
Video	x		x	x
Data	x	x	x	x
Storage/DCN	x		x	

	N_1	N_2	N_3	N_4
1GbTX		x		x
1GbRX	x	x		x
2.5GbTX	x		x	
2.5GbRX	x		x	

C. SLiT Assumptions

This subsection lists SLiT assumptions. A node cannot be part of two SLiTs on the same wavelength even if the two SLiTs are graphically nonoverlapping. This feature is another major differentiator from light-trails [13]. Two SLiTs on the same wavelength can coexist if and only if they do not have any common nodes between them, while in light-trails, two light-trails can coexist if they have a single common node between them. Further, the control channel in a SLiT is bidirectional—implying two separate wavelengths, one Eastward and the other Westward.

IV. DESCRIPTION OF THE DYNASPOT TEST-BED

The DynaSPOT test-bed is a 4-node open-optical metro-ring that can support single wavelength communication at both 1 and 2.5 Gb/s. It is in the future expandable to support WDM communication at 100-GHz channel spacing and bit-rates up to 10 Gb/s. The test-bed is built to support metro applications and makes use of WDM optics at 100 GHz spacing and a channel count of 40. SLiT technology forms the cornerstone of the test-bed.

In the present version of the test-bed, the SLiT is statically set up while connections over the SLiT are dynamically set up based on the valuation protocol that we describe in Section IV-F.

Our objectives of the test-bed are: i) to demonstrate dynamic bandwidth allocation within a SLiT; ii) the allocation of subwavelength flows to each node in the SLiT; iii) achieving multirate communication within the same SLiT (using different *type* of trailponder cards, i.e., at different bit-rates; and iv) provisioning metro services such as VoIP, Triple Play, VoD, PWE3, and data centers.

A. Design Choices of the DynaSPOT Test-Bed

The following design choices are involved in the DynaSPOT test-bed.

SLiT: A 4-node SLiT is created to facilitate subwavelength dynamic service provisioning. A node can communicate to any of the other three nodes using SLiT principles of communication (all-optical, time-sharing of bandwidth) mentioned in Section III. An out-of-band control channel is used for arbitration. At each node we have designed and implemented a trailponder with support of one or more plausible bit-rates. Each trailponder is provided with a PowerPC embedded in an FPGA (*Xilinx Virtex 2Pro*) for arbitration and control purposes. Nodes can establish connections with other nodes that have compatible receivers—i.e., a source node can communicate with destination node(s) under the condition that the destination node has a receiver at the same line-rate as the source-node. In this way multiple nodes can time-share the SLiT at different line-rates.

For example, in the 4-node SLiT, nodes N_1 and N_3 communicate at 2.5 Gb/s while nodes N_2 and N_4 communicate at 1 Gb/s. Further, node N_1 can receive information at both 2.5 Gb/s, as well as 1 Gb/s.

The transmitter, receiver and traffic profile that is provisioned at the nodes is shown in Table I. The SLiT is assumed to be time-slotted with data time-slots of 400 μs duration separated by guard-bands of 10 μs . It is possible to change the duration of the time-slots in the range of 300 to 5 μs depending on specific user requirement.

A general guideline for slot duration selection is that larger the time-slots greater the average delay and an overall betterment of efficiency (at higher loads). Our choice of smaller data time-slots (of 400 μs) is based on the logic that we desire to support delay sensitive services and also the fact that at low to medium loads the efficiency of the system (defined as average utilization of occupied data time-slots) is comparable to the efficiency obtained when the time-slots are of larger duration. This is because, with large time-slots, at low to medium loads the nodes do not have enough data in the buffer (within the trailponders) to occupy entire slot width.

B. Node Architecture

The node architecture used in DynaSPOT is shown in Fig. 3 and the critical subsystem used for provisioning traffic—the *trailponder* is shown in Fig. 4. The conceptual layout of the test-bed is shown in Fig. 5 while the photographs of the test-bed are shown in Figs. 6 and 7. Specifically, Fig. 6 shows a single-node trial version of the DynaSPOT test-bed.

The node architecture has evolved from a reconfigurable optical add-drop multiplexer (ROADM), with incoming WDM signal demultiplexed by an arrayed wave guide (AWG). Our AWGs have a special slot to demultiplex and multiplex signal

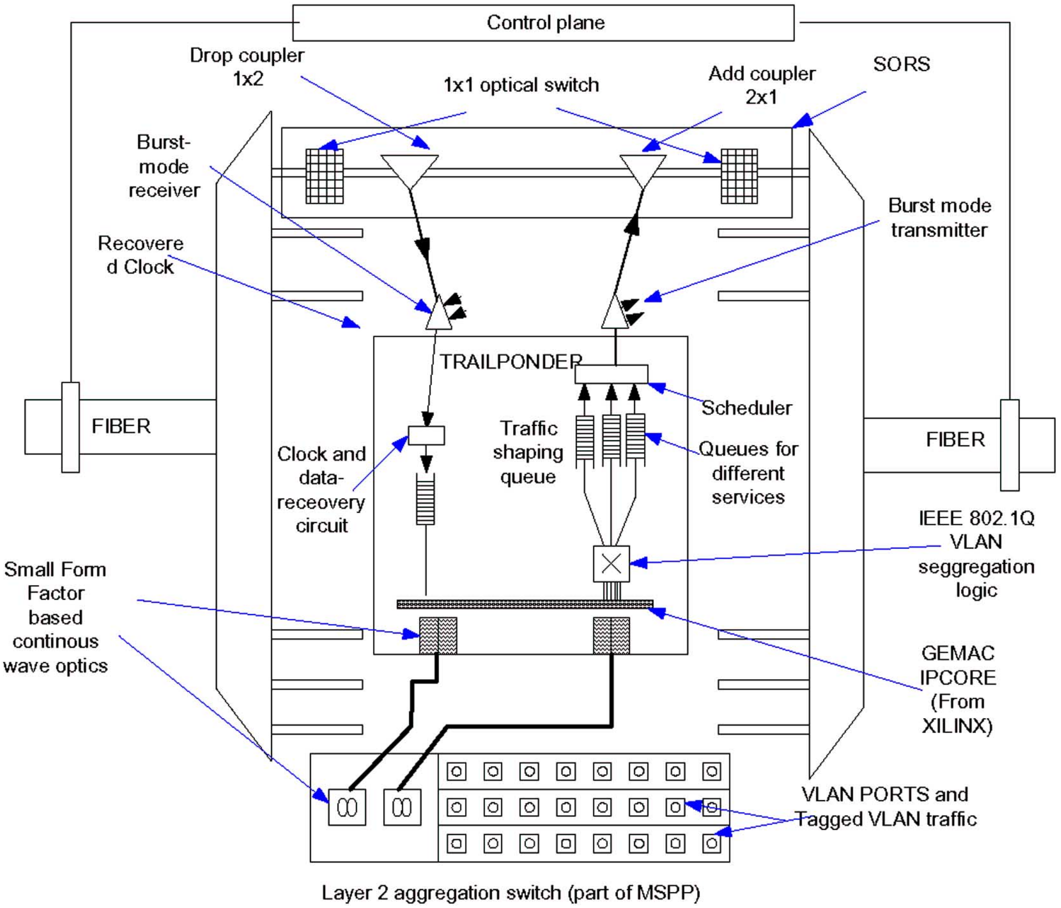

Fig. 3. Node architecture in the DynaSPOT test-bed.

at 1550 nm (non-ITU wavelength) to support off-the-shelf PON optics which are presently used in the test-bed. Each constituent wavelength (in the C-band) is fed to a SLiT Optical Retrieval Section that consists of a series of two ON/OFF (optical) switches (with 1-dB insertion loss and 5-ms switching time) and two passive optical couplers (3 dB) in 2 × 1 configuration.

Incoming signal (in either direction) is dropped and continued at the drop-coupler while local signal can be added into the SLiT passively through the add-coupler. The two couplers have 50/50 splitting/combining ratio. The two optical switches are in the ON state at all intermediate nodes in the SLiT. At the extreme nodes the switches on the outer end are in the OFF state, i.e., for the East-most node the East-most switch is in the OFF state, while for the West-most node the West-most switch is in the OFF state.

In addition, a single variable optical attenuator (VOA) is used to stabilize optical power-level in the SORS. The VOA enables features of flattening gain tilt as a result of skewed bidirectional amplification. The VOA is connected to the arbiter. The arbiter runs a simple gain-control algorithm to stabilize signal power.

To add and drop signal, to and from the time-shared SLiT, as well as to facilitate service provisioning, we use the trailponder. The trailponder (as shown in Fig. 4) consists of burst-mode optics (laser and receiver) [1], [2] and associated electronics (memory, processor) and is triggered through a control card. The trailponder is analogous to a transponder—it facilitates client signals to be transmitted (and received) over the SLiT. The OE (and EO) trailponder card has the added function of storing data as well as scheduling stored data in an optimal manner to facilitate service provisioning. The trailponder stores data in a way that the client side layer-2 equipment (typically an aggregation switch) is oblivious to the storage and scheduling of data over the shared medium SLiT. To do so, it has to manipulate the layer-2 protocol which will be described in Section IV-D.

Apart from providing access to the SLiT, the trailponder is a key device that enables services to be inculcated into our valuation based bandwidth provisioning algorithm. Flows corresponding to different services are fed into the trailponder as shown in Fig. 4. The trailponder then differentiates services based on the following technique. It uses either a VLAN based differentiator that segregates incoming packets (from a client) based on VLAN tags (explained in Section IV-E), or uses the *type* field for differentiation if the transmission format is based on *Generic Framing Procedure* (GFP), ITU G.7041. Packets based on service types are stored in corresponding *service buffers*. Whenever a node is granted a data time-slot for transmission, the data is then sent into the SLiT. To do so, the trailponder uses a TX_EN signal that enables the burst-mode laser. Once the laser is enabled, the trailponder maps the stored data into Ethernet frames or GFP payload. To do so, it selects a corresponding MAC address using the GEMAC IP CORE from Xilinx. The data from the buffers is

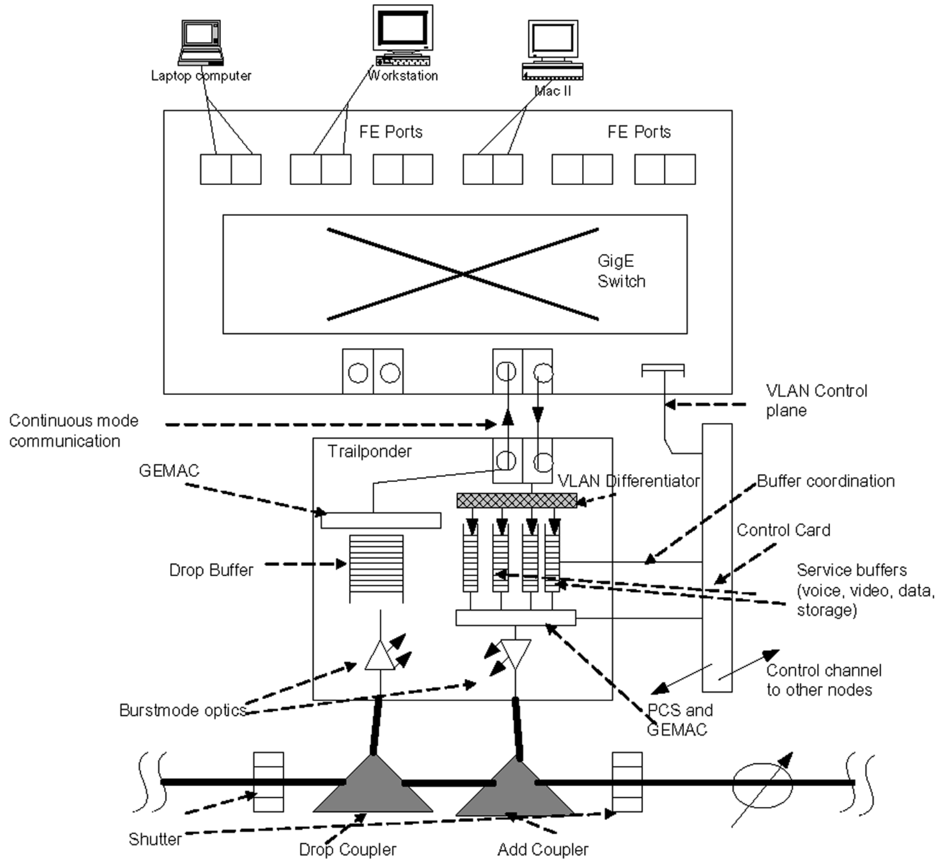

Fig. 4. Trailponder (for VLAN support) and control card.

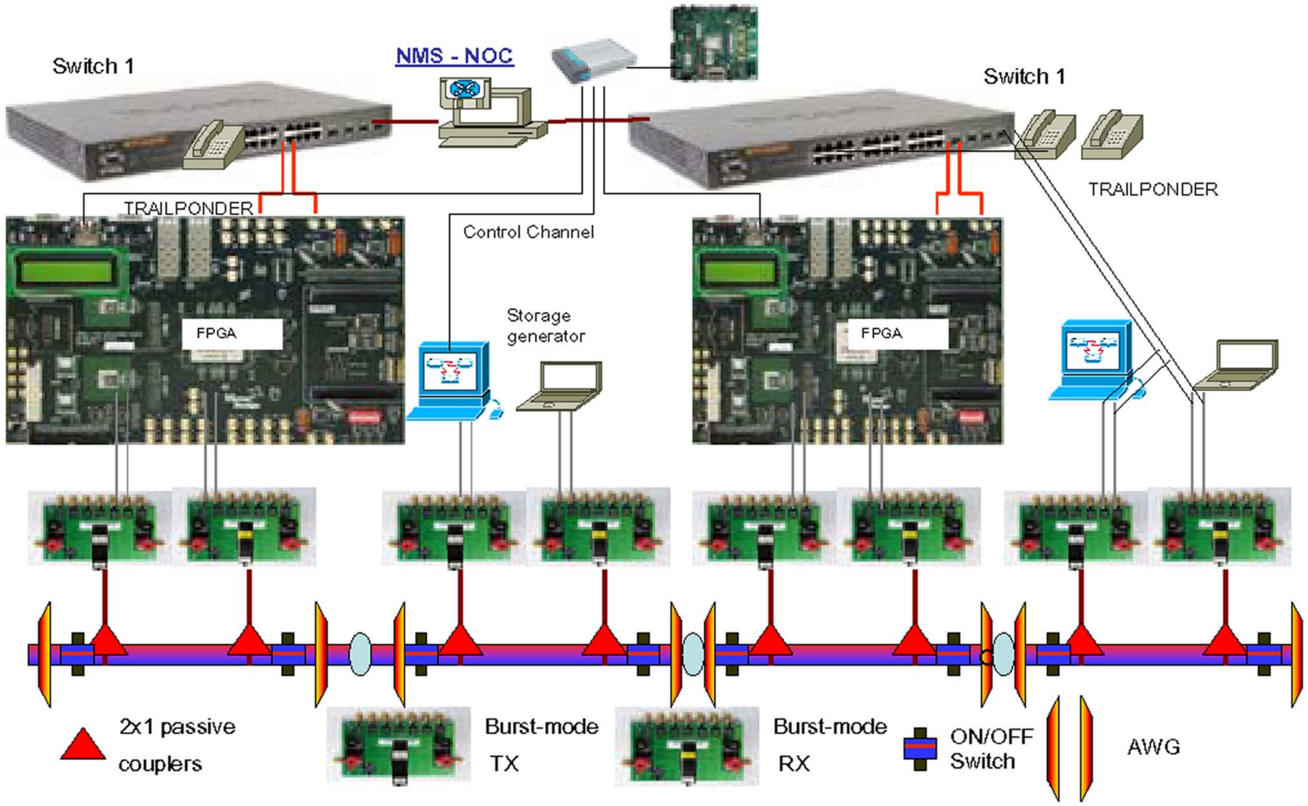

Fig. 5. Conceptual layout of the test-bed.

now perfectly aligned into an Ethernet frame. The multiple service buffers are emptied in the following priority order: storage → voice → video → data. The total memory allocated for buffering in a trailponder is 2 Mb (1 Mb for TX, 1 Mb

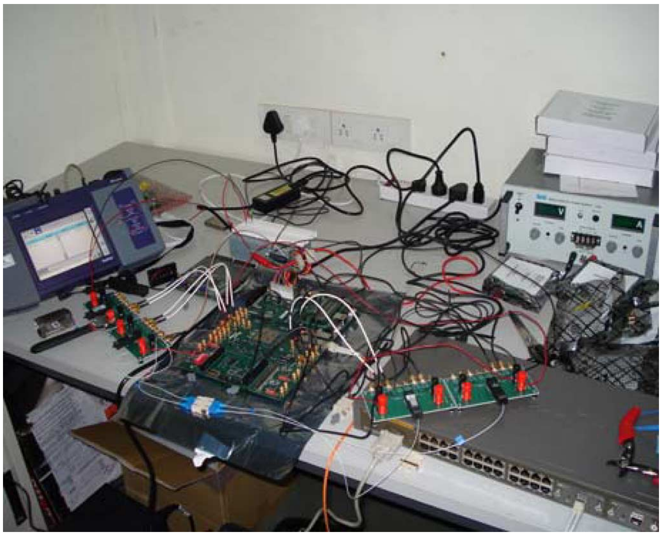

Fig. 6. Single-node configuration for testing.

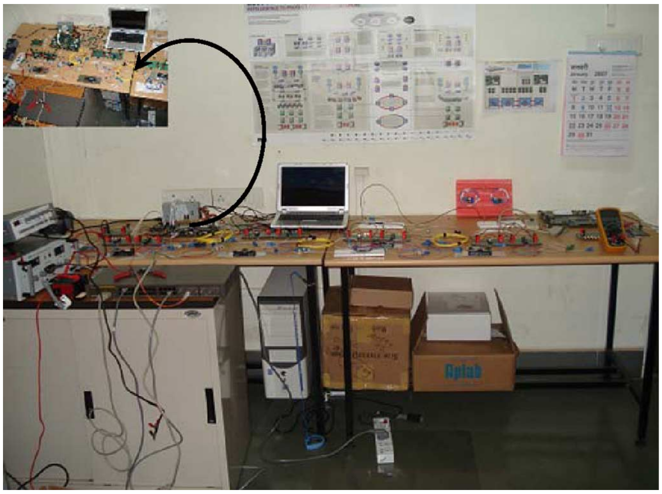

Fig. 7. 4-node DynaSPOT test-bed using SLiT technology.

for RX). Further, there are four service buffers of size 200, 200, 200, and 400 kb. The 400-kb buffer is used for data (due to it being bandwidth intensive) while the remaining three buffers are used for voice, video and storage/data center/pseudowire traffic. The 1-Mb buffer at receiver side is used for traffic engineering (a future function that is not yet implemented in the DynaSPOT test-bed).

Note on buffer requirement in the trailponder: Data time-slot duration, number of nodes in the SLiT and line-rate are all dependent on buffer size available in the trailponder. A detailed discussion on buffer sizing for light-trails (also valid for SLiT) is available in [17]. Larger the available electronic buffer in the trailponder, the higher the efficiency of the SLiT due to larger data time-slot size. However, with large slot-size, the average delay suffers (as each node transmits for a longer time). Further, keeping the slot-size constant, a larger buffer implies more nodes can be added into the SLiT, but here again, this results in a negative penalty on average delay. Through simulation experiments we have observed that any slot-size increase beyond 3 ms is detrimental to delay for SLiTs of size 5–7 nodes (typical metro environments). For SLiTs that have more than 8 nodes, the recommended slot-size is 1–2 ms and the buffer desired is of the order of 1–2 Mb. A second problem with larger buffer is the time penalty experienced in reassembly of Ethernet frames. Hence it is overall desired to have a smaller buffer even if it implies a small drop in efficiency, resulting in low delay and ease of operation.

C. Control Card

This is built using a Virtex2 FPGA board (XUP V2P) and is connected to all the nodes through a 10/100 Ethernet switch. For sake of simplicity, it is assumed that the control channel is nonblocking, implying that when two nodes send requesting signals on the control channel the two signals (in form of Ethernet packets) do not collide. In practice, the control channel would be a dedicated wavelength (at 1510 nm) which would be time-slotted with miniature control slots (compared to the data time-slots). However, to maintain simplicity and due to collocation of the nodes (in the test-bed), we simply use fast Ethernet-based control channel that is connected to a single control card. This single control card acts as an arbiter and is provisioned at node N_3.

D. Connection Provisioning

This subsection discusses how connections are provisioned within the SLiT. At the beginning of each data time-slot, every node sends a request to the arbiter node (node N_3 in our case) to form a connection. The request is sent as a utility *valuation*. The valuation is computed by the trailponder based on a method that is described in Section IV-F. The trailponder computes valuation (in every data time-slot) and sends this to the arbiter through a trailponder fast Ethernet interface. To do so, the computed valuation (a numerical quantity in [0,1]) is mapped to an Ethernet frame. This frame is sent to the arbiter using a Windowed automatic repeat request (ARQ) protocol: as part of this protocol, the trailponder sends its valuation to the arbiter. If the arbiter receives the valuation correctly, (without any transmission-line errors that are detected based on single bit parity); then the arbiter sends the same valuation back to the trailponder in another Ethernet frame. If the returned valuation is within a certain time-window since the trailponder first sent data, then the trailponder knows that the arbiter has correctly received its request for bandwidth in the next time-slot (connection formation). If however, the initiating trailponder does not receive the acknowledgement frame within the specified window, it then resends the valuation. Typically, the window period is set to half data time-slot length. Upon receiving the valuations from all the trailponders the arbiter decides which node would transmit in the next slot based on the highest utility valuation. This node is *granted* permission to set up a connection (of max duration 400 μs). Between two data time-slots there is a guard-band of 10 μs. This is necessary to *reset* the burst-mode receiver logic bias.

E. VLAN/GFP Type Based Service Differentiation

As shown in Fig. 4, traffic from the client-side arrives into the trailponder in a continuous-mode fashion (typically through 10/100/1000 Mbps interfaces). This traffic is aggregated by a layer-2 switch and then sent to the trailponder. The trailponder receives layer-2 frames (Ethernet) with VLAN (IEEE 802.1Q) tags and either sends them into the network as Ethernet frames (at 1 Gb/s), or maps these into GFP frames (at 2.5 Gb/s) depending on the line-rate of the burst-mode transmitter. While transmitting data, the trailponder preserves the mapping between ingress VLAN tag and egress GFP-type field. These frames (Ethernet/GFP) are queued in service buffers. For the former case of end-to-end Ethernet traffic, the frames enter the trailponder FPGA through the Xilinx GEMAC that is engineered to capture VLAN tags as well as to perform physical coding sublayer (PCS) functions. As part of PCS function, the header of the incoming Ethernet frame is stripped off. The header is separately stored on an on-board FPGA cache memory, while the data part is stored into a FIFO that is created in an SDRAM that is available on the Xilinx Virtex 2Pro 1152 board. Interconnection between the SDRAM and the FPGA is managed by one of the two on-board (FPGA) PowerPCs, and the actual interface is operated by a memory access module embedded in the FPGA.

Since VLAN tags have 3-bits allocated for service type, the FPGA is able to store the incoming data (within the Ethernet frame) into one of the four appropriate buffers (as shown in Figs. 3 and 4). The information is stored into a buffer depending on the match of the buffer id ($000 = \mathrm{voice}, 001 = \mathrm{video}, 011 = \mathrm{storage/datacenter/pseudo\text{-}wire}$ and all $\mathrm{other} = \mathrm{data}$) to the corresponding VLAN tag or GFP-type.

F. Utility Valuation and Provisioning

In this subsection, we discuss how valuations are computed. Valuation is a quantitative measure reflecting how much a node requires the next data time-slot in order to meet both its bandwidth and delay needs. The challenge in computing valuation is that bandwidth intensity and delay sensitivity are difficult to normalize with respect to each other. What we are interested in is the net utility that the network would get if a node is allowed to provision bandwidth while meeting the node's requirement. From the work of Shenker [14] we know that from a utility standpoint, a sigmoidal like utility function best describes the time-variant need of real-time services. This implies that a function that has a sigmoidal-like curvature as a function of time is ideal for real-time services like voice, video, etc. Likewise, Lee, Shroff, and Mazumdar [15] have shown that concave utility functions serve as optimal allocation strategy for bursty data traffic.

Hence, the valuation that we compute is based on both utilities—bandwidth as well as delay and is now shown: For every buffer, the trailponder computes a value called *time-to-service*, defined as the time remaining before which a buffer must be scheduled (into the SLiT) or else the longest waiting packet (of that service) would be timed-out (service latency not met). If $a_i^j(t)$ is the time-to-service for the jth buffer at node N_i at time t, then $b_i(t) = \min_j \left(a_i^j(t)\right)$ is called the delay criticality [10] and represents the minimum time before which the node must be serviced, or else packets from at least one buffer would be timed out.

The trailponder also computes the buffer activity period as follows: Let $c_i^j(t)$ be the time elapsed since the first packet entered buffer j at node N_i at time t; then, $d_i(t) = \max_j \left(c_i^j(t)\right)$ represents the activity period of the trailponder at node N_i. The trailponder can now compute service valuation (resulting from the delay sensitive services) as

$$s_i(t) = \frac{b_t(t)}{1 + d_i(t)}. \tag{1}$$

In [16] we have shown that probability distribution function of the above can be reduced to a sigmoidal like function.

Similarly the trailponder also computes a value of *buffer utilization*, defined as the ratio of the number of bits in the buffer to the total buffer capacity. Hence, if $y_i(t)$ is the total number of bits in all the buffers at node N_i at time t, and Y is the total size (max capacity) of the buffers combined, then buffer utilization (valuation) is computed as

$$z_i(t) = \frac{y_i(t)}{Y}. \tag{2}$$

Again, as shown in [16], the above function $(z_i(t))$ for Poisson and Pareto arrivals has a concave distribution.

Since both service valuation and buffer utilizations are entities yielding ratios in the range 0 [1], the trailponder passes on the maximum of the two ratios to the control card. Hence the valuation that a node sends is given by

$$\mathrm{val}_i(t) = \max[z_i(t), s_i(t)]. \tag{3}$$

Each control card then sends this valuation to the arbiter. The valuation is analogous to the utility that the node has for the bandwidth (in the next data time-slot) and hence we also call the valuation as a utility valuation. The control card also sends a *destination* list to the arbiter—that consists of all possible destination MAC addresses of packets stored in the buffer. The multiplicity in destinations facilitates optical multicasting feature in the SLiT bus and also supports multirate/speed communication.

G. Multirate/Speed Support

When a connection is provisioned at a certain line-rate, the source and destination node are assumed to have the requisite laser/receiver at that line-rate. It may however happen that a node has 2 receivers at different line rates (1 and 2.5 Gb/s) connected passively (in drop and continue fashion). In such a case, the node must switch OFF the receiver that is not in sync with the line-rate of the connection. This is done ahead in time through the OOB control packet (*grant*) that is sent by the arbiter node and triggers the corresponding receiver bias OFF. This works as follows: if N_1 is the source node and N_2 is the destination node, and if N_2 has receivers for both 1 and 2.5 Gb/s while N_1 has a transmitter at 1 Gb/s, then N_2 has to switch OFF its receiver at 2.5 Gb/s. When N_1 gets a grant message from the arbiter to form a connection, the arbiter also tells N_2 (through an Ethernet frame in the control channel) that N_1 would be transmitting at 1 Gb/s line-rate. N_2 then switches OFF its receiver at 2.5 Gb/s. The assumption here is that the arbiter has global knowledge about which node has what line-rate capabilities. The assumption is valid since the arbiter acts as a central point of intelligence connected to all the nodes through the control channel.

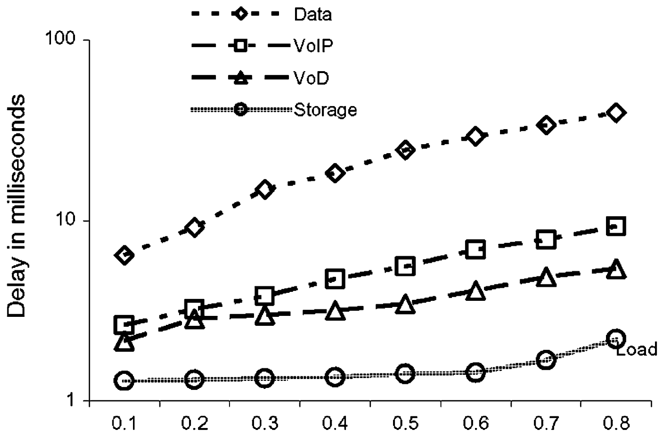

Fig. 8. Delay as a function of load.

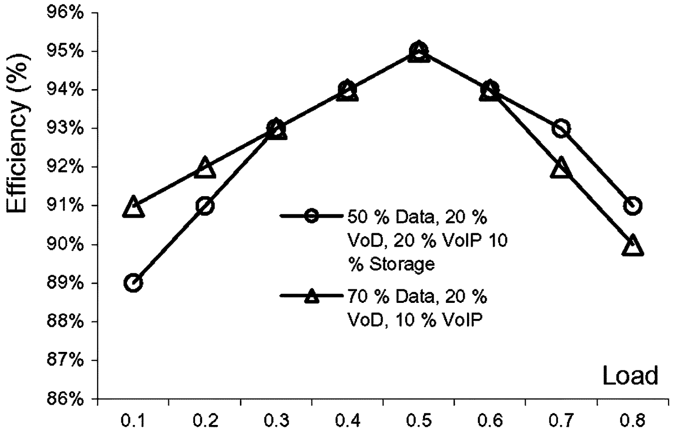

Fig. 9. Efficiency of the system.

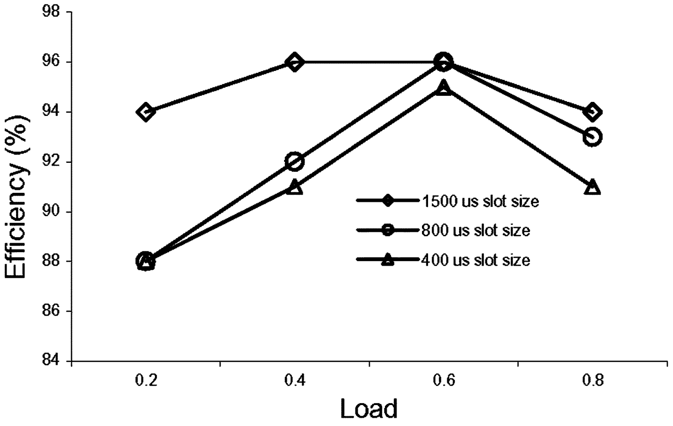

Fig. 10. Efficiency within the SLiT for different data slot-sizes.

V. EXPERIMENT AND RESULTS

We performed an experiment with a 4-node SLiT at 1550 nm in the DynaSPOT test-bed. The SLiT supported line-rates of 1 Gb/s (based on Gigabit Ethernet) as well as 2.5 Gb/s (based on GFP). The setup assumed four traffic types—VoIP, VoD, data and storage/data-center/pseudo wire; feeding into an aggregation switch at all the nodes. All incoming (client-side) Ethernet traffic was appended with VLAN tags while GFP traffic was differentiated based on the type field. Traffic was varied by a generator through a layer-2 aggregation switch as well as a Gigabit Ethernet VLAN tester that generated tagged VLAN based packets. 2 nodes were connected to the aggregation switch while two other nodes were connected to the Gigabit Ethernet VLAN tester. Traffic intensity (load) could be varied at both generators. Measurements were performed by increasing traffic from 100 to 800 Mbps for GigE and 200 Mbps to 2 Gb/s for GFP. Load was computed as a ratio of the total number of bits that enter the SLiT (computed over all the nodes) in one second, to the SLiT line-rate.

Shown in Fig. 8 is the average end-to-end delay profile for data, VoIP, VoD and storage traffic as a function of load. The profile of traffic is shown in Table I. VoIP traffic was generated as part of the VLAN tester as well as by emulating a point-to-point Skype connection. The VoD model is representative of a *Video Hub Office* (VHO) and several *Video Serving Offices* (VSOs). Video on demand traffic was emulated by connecting one node to a video server (through the Ethernet switch) while the other nodes acted as recipients of video traffic. In Fig. 8 we observe that the average end-to-end delay is well within the acceptable service latency requirements even when the SLiT is heavily loaded with duplex VoIP traffic and delay sensitive storage traffic.

For storage/data-center traffic we assume a largely dynamic and extremely delay sensitive traffic characteristic (delay tolerance <10 ms). To emulate storage traffic we create end-to-end pseudo wires that connect two hard-drives. The pseudowire traffic is then mapped into Ethernet frames based on RFC 3985 PWE3 (pseudowire edge to edge emulation). Dynamism is brought about in the network by a C# applet that controls each hard-disk and that requests for data transfer from one hard-disk to another (over the SLiT). The rate of requests and amount of data transfer can all be varied to observe performance.

Fig. 9 shows the efficiency of the system with two different traffic mixes, one with 10% storage traffic and the other with no storage traffic. In Fig. 9, we observe that with more data traffic (and less dynamic—storage traffic), efficiency is better at lower and medium loads but degrades at heavy loads due to data-burstiness. Not shown in the figure but also observed is that the fall in efficiency is primarily because data time-slots are scantily utilized for highly delay sensitive storage traffic. This behavior is not observed in simulation results in [5]. Another important characteristic of provisioning storage traffic through our test-bed is that the efficiency for storage traffic alone is lesser (about 10%–20%) in the test-bed than as observed in simulation. This behavior can be explained due to the way in which data is stored and fetched in the trailponder: the payload of a layer-2 frame is stored separately and in order to provision fast access (as in case of storage traffic) the time-slot is somewhat under utilized while the frame is being reassembled for transport through the SLiT.

Fig. 10 shows the effect of slot size on efficiency. It is seen that larger the slot-size better the efficiency. The complete picture, i.e., the effect on delay is shown in Fig. 11. Shown in Fig. 11 is the effect of slot-size increase on efficiency and the corresponding increase in delay. We here define a parameter p

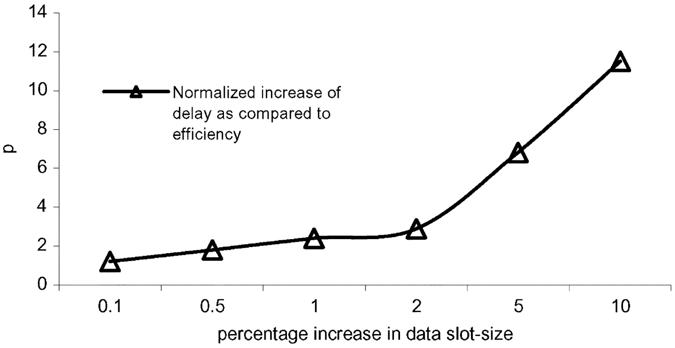

Fig. 11. Normalized increase of delay as a function of efficiency for increase in data time-slot size.

that is the ratio of increase in efficiency to the corresponding increase in delay, normalized over the corresponding increase in slot-size. Hence, p would be unity if a 10% increase in slot-size results in a 10% increase in efficiency and a 10% increase in delay. As can be seen, p is not linear for increase in slot-size (starting from 400 μs). This shows that by using larger slot-sizes we can achieve better efficiency but we have severe penalty in terms of delay.

The result (in Fig. 11) has so far not been seen in simulation experiments and is dependent on trailponder behavior, control channel provisioning, and slot-size.

VI. COMPARISONS WITH EXISTING TECHNOLOGIES

In this section, we compare SLiT technology to existing approaches especially from the test-bed perspective.

A. *Comparison With SONET/SDH*

The primary advantage of our approach as opposed to the SONET/SDH scheme is that of cost. Since SONET/SDH solutions require OEO devices in-line at every node and further each OEO cross-connect is synchronized with every other similar device in the network, the equipment cost is high. In our case, all the electronics is relegated to the periphery of the network, and the optical line is based on bus principles, thus, reducing electronics and hence cost.

A second advantage of our scheme over SONET/SDH is that of being able to support dynamic bandwidth provisioning using the valuation protocol we discussed. Our scheme is based on statistical TDM, while SONET/SDH pure TDM. In our scheme, the assignment of a data time-slot to a node is dependent on the node's delay and bandwidth requirement. This means that when the node's delay or bandwidth requirement (service) changes, so does the slot allocation pattern, resulting in on-demand bandwidth provisioning.

A third advantage of our proposal is in providing optical layer multicasting that is essential for video services. The end-to-end delay observed in SONET systems is very low, while we do have a significant delay penalty especially at high loads. However, the advantages we experience through dynamic provisioning, low-cost and optical multicasting while maintaining delay below the tolerance level (see Fig. 8) makes our proposal more attractive than SONET/SDH. The simulation results in [4] and [10] are quite close to our observed results in the test-bed. The minor difference is because: (a) the simulation results are primarily comparing SONET/SDH with light-trails that are lesser "busy" than SLiTs, and (b) because of hardware issues like finite guard band and packet assembly, etc., seen in the test-bed that were not considered in the simulation.

B. *Comparison With Optical Burst Switching and Optical Burst Transport*

Light-trails have been compared to optical burst switching (OBS) in [4] and to optical burst transport (OBT) in [18]. Those results were based on simulation models and have been ratified through our experiment. We observe that due to switch-less nature of SLiTs there is a 80% efficiency improvement over OBS for high-loads. This is because of the large time required in OBS to configure the optical switches (assuming use of mechanical switches similar to the ones used in our test-bed). As compared to optical burst-transport we have a 40% efficiency benefit at low loads that increases to about 55% at higher loads. This is due to the time lost in OBT for burst aggregation. If OBT assumed a similar aggregation technique like ours, i.e., fixed sized time-slots the situation would better by about 20% (in terms of efficiency). However, the assumption there would be that N^2 wavelengths are available. Since, SLiTs are a two-layered provisioning model (an $n-$ node SLiT once provisioned can support n^2 connections), hence, the requirement of a large number of wavelengths is removed. Another differentiator between SLiT and OBS/OBT is the ability to support delay sensitive and bandwidth intensive services. Our valuation protocol forms the cornerstone of being able support the four services through our test-bed. To the best of our knowledge no data exists about how OBS/OBT would performed for all the four services that we have provisioned through our test-bed.

C. *Comparison to RPR*

The IEEE 802.17 standard, i.e., RPR is a forerunner for metro ring networks. We have theoretically compared this in [6]. The main advantages our approach over RPR is that of being able to provide on-demand bandwidth, better efficiency, lower-cost (no OEO), and providing fairness. The DynaSPOT test-bed results conform to these claims. Of primary interest is the ability of the test-bed to showcase fairness. The protocol proposed in [16] for light-trails, which we have extended for SLiTs in this test-bed has been theoretically proved to provide proportional fairness [12]. By demonstrating this protocol over the test-bed we claim that our proposal has better fairness than RPR while being lower cost and resulting in similar efficiency [6].

VII. CONCLUSION

We demonstrate the DynaSPOT test-bed to support subwavelength grooming, dynamic service provisioning (VoIP, data, VoD and storage), and multirate communication (1 and 2.5 Gb/s) over a single wavelength using SLiT technology. Services such as VoIP, data, video on demand, storage (data-center) are provisioned over our test-bed using a dynamic bandwidth allocation protocol. An architectural overview and experimental results are presented. Dynamic bandwidth allocation is based on a novel auctioning algorithm that computes bids as valuations reflecting bandwidth and delay needs of the services is shown and implemented in the test-bed. Compliance to high efficiency and delay requirements of the provisioned services are shown. In summary, we demonstrate different features of SLiT technology showcasing it as an enabler for next generation metro applications, by providing the required features for service provisioning and using a low-cost and evolutionary (from ROADM) set up.

ACKNOWLEDGMENT

The authors would like to thank R. Nachane (JDS Uniphase), R. Subramaniam, and N. Varma (Xilinx/Avnet) for their equipment support, and Prof. K. Ramamritham (IIT Bombay) and Prof. B. Menezes for their encouragement and Prof. B. Mukherjee (UC-Davis) for his input on their test-bed.

REFERENCES

[1] A. Gumaste, S. Jain, and S. Q. Zheng, "SLiT: Strongly connected light-trail solution for cost efficient and dynamic optical networking," in *Proc. 22nd IEEE/OSA Opt. Fiber Commun. Conf. OFC 2006*, Anahiem, CA, 2006, JThB61.

[2] A. Gumaste, N. Ghani, P. Bafna, A. Lodha, S. Srivastava, T. Das, and S. Zheng, "Achieving multi-rate dynamic subwavelength service provisioning in strongly connected light-trails (SLiTs)," in *Proc. IEEE/OSA OFC Opt. Fiber Conf.*, Anaheim, CA, Mar. 2007, Postdeadline paper.

[3] P. Bafna, A. Gumaste, and N. Ghani, "Delay sensitive smoothed round robin scheduler (DS2R2) for light-trail and SLiT networks," in *Proc. IEEE/OSA OFC 2007*, Anaheim, CA, TuG.

[4] A. Gumaste and I. Chlamtac, "Light-trails: An optical solution for IP transport," *OSA J. Opt. Netw.*, pp. 864–891, May 2004.

[5] A. Gumaste and S. Zheng, "Next generation optical storage area networks: The light-trails approach," *IEEE Commun. Mag.*, vol. 21, no. 3, pp. 72–79, Mar. 2005, and references therein.

[6] A. Gumaste and S. Q. Zheng, "Optical implementation of resilient packet rings using light-trails," in *21st Proc. Opt. Fiber Conf./Nat. Fiber Opt. Eng. Conf. NFOEC/OFC*, CA, 2005.

[7] A. Gumaste and S. Zheng, "Dual auction (and recourse) opportunistic protocol for light-trail network design," in *Proc. IEEE Wireless Opt. Commun. Conf. (WOCN)*, Bangalore, India, 2006.

[8] A. Gumaste and P. Palacharla, "Heuristic and optimal assignment techniques for light-trail ring WDM networks," *Elsevier's Computer Commun. J. (CCJ)*, pp. 21–32, Mar. 2007.

[9] W. Zhang, G. Xue, and K. Thulasiraman, "Dynamic light trail routing and protection issues in WDM optical networks," in *Proc. IEEE Globecom 2005*, Dec. 2005.

[10] A. Gumaste *et al.*, "On control channel for service provisioning in light-trail WDM networks," in *Proc. 43rd IEEE Int. Conf. Commun. ICC 2007*, Glasgow, UK, Jun. 2007.

[11] Y. Ota and R. Swartz, "Burst-mode compatible optical receiver with a large dynamic range," *IEEE J. Lightw. Technol.*, vol. 8, no. 12, pp. 1897–1903, Dec. 1990.

[12] F. Kelly, Models for a Self Managed Internet Univ. Cambridge, U.K., online version.

[13] A. Gumaste and I. Chlamtac, "Mesh implementation of light-trails: A solution to IP centric communication in the optical domain," in *Proc. 13th IEEE Int. Conf. Computer Commun. Netw. ICCCN*, Dallas, TX, Oct. 2003.

[14] S. Shenker, "Fundamental design issues for future internet," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, pp. 1176–1185, Dec. 1995.

[15] J. Lee, R. Mazumdar, and N. B. Shroff, "Non-convex optimization and rate control for multi-class services in the internet," *IEEE/ACM Trans. Netw.*, vol. 13, no. 4, pp. 827–840, Aug. 2005.

[16] A. Gumaste *et al.*, "Two stage auction algorithm for fair bandwidth allocation and topology growth in light-trail based WDM networks," *IEEE J. Sel. Areas Commun.*, 2007, submitted for publication.

[17] A. Gumaste, P. Palacharla, and T. Naito, "Performance evaluation and demonstration of light-trails in shared wavelength optical networks (SWON)," in *Proc. 31st Europ. Conf. Opt. Commun.*, 2005, (ECOC) Glasgow.

[18] A. Gumaste, A. Jukan, A. Lodha, X. Cin, and N. Ghani, "A novel node architecture for light-trail provisioning in mesh WDM metro networks," in *Proc. OFC 2008*, San Diego, CA, 2008.

Ashwin Gumaste (M'04) received the Ph.D. from the University of Texas, Dallas, in 2003. As part of his degree, he proposed the light-trail and light-frame concept that has received wide acceptance both in industry and academia.

He is currently a faculty member with the Department of Computer Science and Engineering, Indian Institute of Technology (IIT), Bombay (2005–2007). In 2008, he will be on sabbatical leave as a Visiting Scholar with the Massachusetts Institute of Technology (MIT), Cambridge. From 2001 to 2005, he was with Fujitsu Laboratories (USA) Inc. as a Member of Research Staff in the Photonics Networking Laboratory. Prior to this, he was with Fujitsu Network Communications R&D and prior to that, with Cisco Systems in the Optical Networking Group (ONG). He has more than 40 pending U.S. and EU patents and has published approximately 70 papers in referred conferences and journals. He has authored three books in broadband networks, namely *DWDM Network Designs and Engineering Solutions* (Indianapolis, IN: Pearson Education/Cisco, 2002), a networking bestseller, *First-Mile Access Networks and Enabling Technologies* (Indianapolis, IN: Pearson Education/Cisco, 2004), and *Broadband Services: User Needs, Business Models and Technologies* (New York: Wiley, 2005). He is also an active consultant to industry and has worked with both service providers and vendors.

Dr. Gumaste has served as Program Chair, Co-Chair, Publicity Chair, and Workshop Chair for various IEEE conferences. He is also a Guest Editor for the IEEE COMMUNICATIONS MAGAZINE and serves as the General Chair of the 1st International Symposium on Advanced Networks and Telecommunication Systems (ANTS 2007) held in Bombay, India.

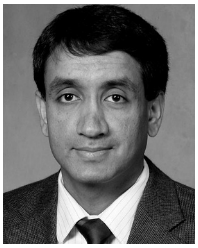

Nasir Ghani (SM'03) received the B.S. degree in computer engineering from the University of Waterloo, Canada, in 1991, the M.S. degree in electrical engineering from McMaster University, Canada, in 1992, and the Ph.D. degree in electrical and computer engineering from the University of Waterloo, in 1997.

He has gained a wide range of industrial and academic experience in the telecommunications area and in the past has held senior positions at Nokia, IBM, Motorola, Sorrento Networks, and Tennessee Tech University. Currently, he is an Associate Professor with the Department of Electrical and Computer Engineering, University of New Mexico, where he is actively involved in a wide range of funded research projects in the area of optical networks and cyber-infrastructures. He has published more than 80 journal and conference papers, several book chapters, various standardization proposals, and has two patents granted.

Dr. Ghani recently served as a co-chair for the optical networking symposia for *IEEE ICC 2006* and *IEEE GLOBECOM 2006* and is a Program Committee Member for *OFC 2007* and *OFC 2008*. Furthermore, he has been a Program Committee Member for numerous IEEE, SPIE, ACM, and IEC conferences and has served regularly on NSF, DOE, and other international panels. He is an Associate Editor of the IEEE COMMUNICATIONS LETTERS journal and has served as Guest Editor for special issues of IEEE NETWORK, IEEE COMMUNICATIONS MAGAZINE, and *Cluster Computing*. He is a recipient of the prestigious NSF CAREER Award.

Paresh Bafna received M. Tech. degree from the Indian Institute of Technology (IIT), Bombay.

His research interests include optical networks and core networking. His research interests include delay sensitive, short-term fair connection scheduler for which he proposed an algorithm called Delay Sensitive Smoothed Round Robin (DS2R2). He is also interested in growth algorithms for light-trail networks especially in the metro and access area.

Akhil Lodha received the B.Tech. degree from the Indian Institute of Technology (IIT), Bombay, in 2007.

He is currently pursuing the M.Tech. degree from IIT Bombay. He has worked on anomaly detection that has been widely cited. As part of his Master's degree work, he has published four papers and has several more under submission. His research interests include networks, security, and optimization

Mr. Lodha has presented his work on stochastic optimization of light-trails at the IEEE Conference on High Performance Routing and Switching (HPSR 2007).

Anuj Agrawal received the B.Tech. degree in electronics and communication engineering from Uttar Pradesh Technical University, Lunknow, India, in 2005.

He is currently pursuing the Masters degree in the Department of Computer Science and Engineering, Indian Institute of Technology (IIT), Bombay. His research interests include broadband access networks.

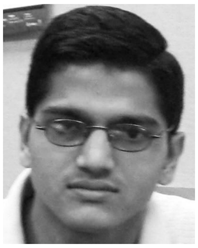

Tamal Das is currently with the Department of Mathematics, Indian Institute of Technology (IIT), Delhi, pursuing a 5-yr Master of Technology degree in mathematics and computing. He has performed research internships at IIT Bombay and the National Center for Supercomputing Applications (UIUC). His research interests are in network security and optimization, algorithms, and software development. He has coauthored a chapter on Game Theory and Networking in *Uncertainties in Next Generation Networks: Using Bayesian Models and Stochastic Analysis*, (Hoboken, NJ: Wiley, 2007). He has also coauthored a paper on the DYNASPOT test-bed at the IEEE Optical Fiber Communication Conference (OFC 2007).

Jianping Wang (M'03) received the B.Sc. and M.Sc. degrees in computer science from Nankai University, Tianjin, China, in 1996 and 1999, respectively, and the Ph.D. degree in computer science from the University of Texas at Dallas in 2003.

She is currently an Assistant Professor with the Department of Computer Science, City University of Hong Kong. Her research interests include optical networks and dependable networking.

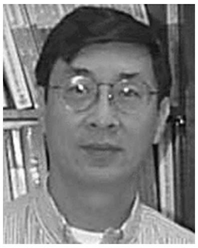

Si Qing Zheng (SM'98) received the Ph.D. degree from the University of California, Santa Barbara, in 1987.

Currently, he is a Professor of computer science, computer engineering, and telecommunications engineering at the University of Texas at Dallas. He is Director of the Advance Computer Network Architecture (ACNA) Laboratory. He was previously with Louisiana State University for 11 years. His research interests include algorithms, computer architectures, networks, parallel and distributed processing, telecommunications, and VLSI. He has published more than 200 papers in these areas.

Dr. Zheng serves on the Program Committees and as Chairman for numerous conferences, as well as Editor for reputed journals.