Achieving 100% throughput in reconfigurable optical networks

Andrew Brzezinski, and Eytan Modiano, *Senior Member, IEEE*

***Abstract*— We study the maximum throughput properties of dynamically reconfigurable optical network architectures having wavelength and port constraints. Using stability as the throughput performance metric, we outline the single-hop and multi-hop stability regions of the network. Our analysis of the stability regions is a generalization of the BvN decomposition technique that has been so effective at expressing any stabilizable rate matrix for input-queued switches as a convex combination of service configurations. We consider generalized decompositions for physical topologies with wavelength and port constraints. For the case of a single wavelength per optical fiber, we link the decomposition problem to a corresponding Routing and Wavelength Assignment (RWA) problem. We characterize the stability region of the reconfigurable network, employing both single-hop and multi-hop routing, in terms of the RWA problem applied to the same physical topology. We derive expressions for two geometric properties of the stability region: maximum stabilizable uniform arrival rate, and maximum scaled doubly substochastic region. These geometric properties provide a measure of the performance gap between a network having a single wavelength per optical fiber and its wavelength-unconstrained version. They also provide a measure of the performance gap between algorithms employing single-hop versus multi-hop electronic routing in coordination with WDM reconfiguration.**

***Index Terms*— IP-over-WDM, wavelength division multiplexing, matrix decomposition, Birkhoff-von Neumann, WDM reconfiguration, performance evaluation, queueing network, input-queueing**

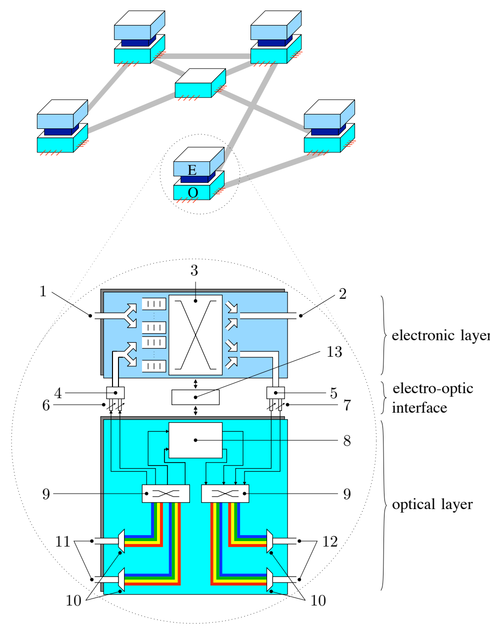

Fig. 1. Network architecture, with each edge node having the following features: 1–electronic inflows; 2–electronic outflows; 3–electronic packet switch; 4–optical to electronic converter; 5–electronic to optical converter; 6–tunable optical receivers; 7–tunable optical transmitters; 8–wavelength converter; 9–optical switch; 10–optical multiplexer/demultiplexer; 11–incoming fiber; 12–outgoing fiber; 13–controller. The network also includes all-optical nodes providing switching/conversion services to incoming fibers.

I. Introduction

We consider an optical networking architecture consisting of nodes having an electronic router overlaying an optical interface, with the nodes interconnected by an optical transport layer. Depicted at the top in Fig. 1 is an example of our architecture with electronic edge nodes interconnected by an optical transport network using optical fiber links. This constitutes the *physical topology* of the network. Optical transceivers, multiplexers/demultiplexers, wavelength converters, and optical switches allow individual wavelength signals to be either *dropped* to the electronic routers at each node or to pass through the node optically. The *logical topology* consists of the lightpath interconnections between the electronic routers and is determined by the configuration of the optical interface at each node [11]. Future optical networks will make use of optical bypass, tunable transceivers, optical switches, and wavelength converters in order to harness the full capacity of the optical transport network. The interaction of these optical components with the electronic interface is depicted at the bottom in Fig. 1. Tunable optical components introduce flexibility to optical networks by enabling logical topology *reconfiguration*. As network traffic changes with time, the optimal logical topology varies as well. In this work, we study the ultimate throughput properties of reconfigurable optical networks. We determine the performance penalty associated with wavelength constraints, and we characterize the performance gap between architectures that employ single-hop versus multi-hop routing at the electronic layer.

The seminal work of Tassiulas and Ephremides underlies much of the existing literature in the area of stability of communication networks [31]. Indeed, the network model considered in this paper easily fits into the framework of Tassiulas and Ephremides, as does much of the switch scheduling literature. To the best of our knowledge, the study of stability properties of optical networks was introduced in [24], [25], [36], where the authors considered optical burst scheduling under dynamic traffic in time-domain wavelength interleaved networks. Subsequent work looking at stability properties of optical networks includes: [4], [5], where scheduling algorithms were introduced for joint electronic routing and WDM

This work was supported in part by the Defense Advanced Research Projects Agency under grant MDA972-02-1-0021 and by NSF under grants ANI-0073730, ANI-0335217, and CNS-0626781.

layer reconfiguration under a variety of practical optical layer constraints; and [34], [35], where the stability properties of optical burst, flow, and packet switched architectures were compared.

A major contribution of our work is a characterization of the stability region for single-wavelength optical networks through a linkage to the Routing and Wavelength Assignment (RWA) problem for WDM networks. This characterization allows us to derive fundamental geometric properties of the stability region for optical networks of arbitrary topologies. In this work, we primarily focus on single-wavelength optical networks. The single wavelength topology is commonly used in traditional metropolitan and access networks operating on one frequency (*e.g.* 1.3nm systems). Moreover, our single-wavelength treatment simplifies the presentation considerably and can be extended, by appropriate scaling of the stability region, to multi-wavelength optical networks.

Our work is conceptually related to Birkhoff-von Neumann (BvN) decompositions, particularly as applied to switching theory [8], [32]. The set of switch configurations (or *service configurations*) available to an $n \times n$ input-queued switch is typically represented by the set of permutation matrices of size n. The result of [31] implies that the convex hull of these service configurations equals the stability region of the input-queued switch. BvN decompositions draw on these concepts to express any stabilizable rate matrix as a convex combination of permutation matrices (service configurations) [8]. An alternative characterization employs a result of Birkhoff [2] to state that the convex hull of the service matrices (permutation matrices) equals the doubly substochastic region [20]. Like BvN decompositions for input-queued switches, our work seeks to express any stabilizable rate matrix as a convex combination of service configurations. Unlike input-queued switches, our optical networking architecture has physical constraints, such as port and wavelength limitations, that affect the set of service configurations. For example, the set of service configurations may not include the full set of permutation matrices, and may include non-permutation matrices. Thus, while the work of [31] allows us to express the stability region as the convex hull of available service configurations, this description can have limited value in providing an understanding of the geometric properties of the stability region. This is in contrast to the case of the input-queued switch, where a result of Birkhoff [2] has been applied to demonstrate that the convex hull of the service matrices (permutation matrices) equals the doubly substochastic region [20]. Recently, the study of [18] has developed order bounds, based on uniform multi-commodity flow, for maximum achievable throughput performance in general network settings. In this paper, we develop a theory of *RWA decompositions* that enables us to *exactly* elicit geometric properties of the stability region of single-wavelength optical networks having general topologies. A preliminary version of this work appeared in [6], [7].

In recent years, tremendous efforts have been made in the research towards so-called "IP-over-WDM" networks. These studies aim to improve network performance through increased electro-optical integration [14], [16], [23], [29], [33], [38]. Several studies consider Optical Burst Switching (OBS) as the mechanism for accessing the optical transport layer [23]–[25], [36], [37], [39]. Most solutions seek to integrate IP and Generalized Multiprotocol Label Switching (GMPLS) functionality. Our work differs from existing studies on electro-optical integration in that we are not tied to a particular protocol suite, but rather employ a "generic" architecture utilizing electronic packet switching along with a reconfigurable optical transport layer. Our approach is to determine the fundamental performance characteristics achievable in general reconfigurable optical networks having varying topology and processing functionalities. We next provide an example of the effects of such physical constraints upon an optical network.

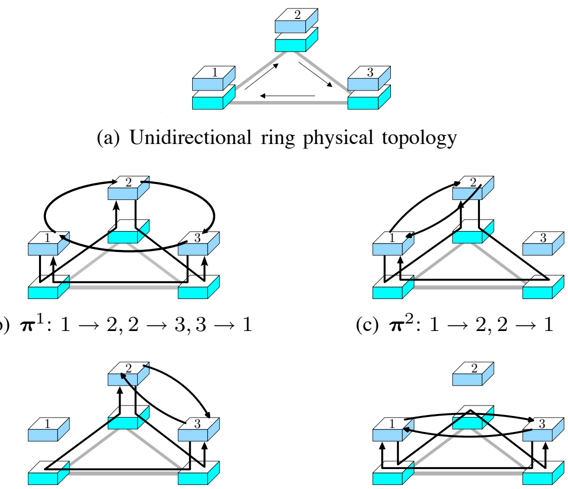

(a) Unidirectional ring physical topology

(b) $\boldsymbol{\pi}^1$: $1 \to 2, 2 \to 3, 3 \to 1$

(c) $\boldsymbol{\pi}^2$: $1 \to 2, 2 \to 1$

(d) $\boldsymbol{\pi}^3$: $2 \to 3, 3 \to 2$

(e) $\boldsymbol{\pi}^4$: $3 \to 1, 1 \to 3$

Fig. 2. There are four maximal logical topology configurations for the unidirectional three-node ring having a single wavelength per optical fiber. The logical configurations are depicted as lightpath routings (straight-edge links with corners) with corresponding logical topology graph overlaid (curved links).

A. Simple motivating example

Consider a unidirectional ring network having 3 nodes. Suppose this network is restricted to a *single wavelength per optical fiber*, with lightpaths routed *only in the clockwise direction*. These constraints restrict the network to four *maximal* logical topologies[1], illustrated in Figure 2.

Consider the traffic matrix $\boldsymbol{\lambda}$, given by

$$\boldsymbol{\lambda} = \begin{bmatrix} \cdot & 0 & \theta \\ \theta & \cdot & 0 \\ 0 & \theta & \cdot \end{bmatrix},$$

where the (i, j)-th entry of $\boldsymbol{\lambda}$ is equal to the average arrival rate of packets to node i destined for node j. We wish to determine the maximum value of θ that the network can support, given that only one packet can be serviced along a logical link per time slot. If we restrict the network to only use single-hop electronic-layer routes, the maximum value of θ is $1/3$. This follows because logical links $1 \to 3$, $2 \to 1$, and $3 \to 2$ each traverse two fibers, which due to the single-wavelength constraint means that only one of these links can be served at

[1]Every valid logical topology is either equal to, or has some subset of links from, one of the maximal topologies.

a time. Sharing time equally between the three links affords a maximum of $1/3$ of the proportion of time to service each link. Thus, $\theta = 1/3$ is the maximum value such that the traffic rate matrix $\boldsymbol{\lambda}$ can be supported.

Suppose instead that we allow the network to make use of multi-hop electronic-layer routes. In this case, a simple policy that maintains logical topology $\boldsymbol{\pi}^1$ (Figure 2(a)) for all time and multi-hops packets along the electronic layer leads to a link load of 2θ on each logical link. Since no more than 1 unit of traffic per time slot can be supported on each wavelength, this policy can support any $\theta \leq 1/2$. This is a clear improvement over the achievable traffic rate matrix supported under single-hop routing. The value $\theta = 1/2$ is also the maximum value achievable, which is easily seen by noting that each physical link has 2θ units of traffic demand that it must service.

For comparison, consider the wavelength-unconstrained case[2] [4], [5], [25], which in the case of the 3-node unidirectional ring topology implies three wavelengths per optical fiber. The maximum value of θ that is supported in this case is $\theta = 1$, which is achievable by maintaining for all time the logical configuration $1 \to 3, 2 \to 1, 3 \to 2$.

This example highlights three important points. First, the wavelength constraint has been shown to reduce the maximum throughput achievable under single-hop and multi-hop routing. This is an example of the intuitively obvious fact that wavelength constraints often lead to throughput penalties. Second, there is a throughput performance gap between electronic layers employing multi-hop versus exclusively single-hop routing. Again, this is intuitively obvious in light of the optical-layer constraints, but this is in contrast to the case of wavelength-unconstrained networks, where single-hop and multi-hop algorithms are identical in terms of throughput performance [5]. Finally, note that both the single-hop and multi-hop cases have made use of service configurations that cannot be equated to permutation matrices, where each input port is always connected to a single output port, each output port is always connected to a single input port, and the connections are exclusively used for single-hop service of packets. This points to the fact that a direct application of BvN decompositions does not apply in constrained network scenarios. These observations suggest three important goals of this paper:

1) to develop a theory of generalized decompositions analogous to BvN decompositions for port and wavelength constrained networks;
2) to explore the throughput penalty of constrained versus wavelength-unconstrained optical networks; and
3) to determine the throughput gap between single-hop and multi-hop electronic-layer routing algorithms.

II. Network throughput

We consider a reconfigurable WDM-based packet network N, consisting of n nodes (the set of nodes is V). The network symbol N refers to all physical aspects of the optical data network, including the physical topology of the network, the number of wavelengths available in each fiber link, and the number of transceivers (or ports) at each node. We assume that node $v \in V$ has P_v transceivers. The network nodes are interconnected by optical fiber. Let $G_P = (V, E_P)$ be the directed *physical topology graph* of the network N: if there exists a fiber between nodes $v_1, v_2 \in V$ along which data can travel from node v_1 to v_2, then the directed edge (v_1, v_2) belongs to E_P. For example, the network of Fig. 2(a) is represented by a clockwise-oriented 3-cycle (3 node ring) physical topology graph.

A direct optical communication link between two nodes is called a *logical link* or a *lightpath*. Such a link consists of an all-optical path through the network N, connecting the nodes, possibly traversing multiple intermediate nodes, with no intermediate electronic processing (see for example the straight-edge links depicted in Fig. 2(b)-(d)). The edges of the *directed graph* $G_L = (V, E_L)$ represent the set of *logical links* that can be enabled in the network. Denote $m = |E_L|$. In general these logical links may not be able to be activated simultaneously, but resources exist to *at least* allow each link to be active individually. We assume that a lightpath can exist between any two nodes, which implies that G_L is a complete graph. At any time, the network may initiate a logical topology reconfiguration, under which existing lightpaths are torn down and new ones are set up.

Throughout the paper, we will treat data destined for a particular terminal node $j \in V$ as *commodity j data*. For a directed edge $e \in E_L$, let $\sigma(e)$ denote the source (initial) vertex, and $\tau(e)$ denote the terminal (destination) vertex.

Packets are assumed to have fixed size, with transmission duration of one slot. This assumption is for simplicity of exposition and can be relaxed with appropriate envelope algorithms [17]. Each node separately enqueues packets for every destination in the network (*virtual output queueing*), with $Q_{ij}(t)$ equal to the number of enqueued commodity j packets at node i, at the beginning of time slot t. The *differential backlog (backpressure)* of commodity j packets corresponding to link $e \in E_L$ at time t is $Z_{ej}(t) = Q_{\sigma(e)j}(t) - Q_{\tau(e)j}(t)$. For link $e \in E_L$, the *maximum backpressure* at time t is given by $Z^*_e(t) = \max_{j \in V} Z_{ej}(t)$.

Data traffic arrives for service through the network according to a stochastic process, $(A_{ij}(t), t \geq 0)$, where $A_{ij}(t)$ represents the cumulative number of exogenous arrivals of commodity j packets to node i, up to the end of time t. We make the assumption that there is no self-traffic in the network, *i.e.* $A_{ii}(t) = 0$ for all $t \geq 0$ and $i \in V$. The arrival processes are assumed to be general, in that they can be temporally and mutually correlated, with λ_{ij} equal to the long term rate of arrivals for source-destination pair (i, j), where $\lambda_{ij} = \lim_{t\to\infty} A_{ij}(t)/t$ w.p.1 (with probability 1). Denote the $n \times n$ arrival rate matrix $\boldsymbol{\lambda} = (\lambda_{ij}, i, j \in V)$.

Let Π_N denote the set of feasible logical topologies in the network: the $n \times n$ matrix $\boldsymbol{\pi} = (\pi_{ij}, i, j \in V) \in \Pi_N$ is a nonnegative integer matrix, where π_{ij} is the number of active logical links from node i to node j. Clearly, Π_N is constrained by the wavelength/port limitations of the network, as well as

[2] A WDM network is *wavelength-unconstrained* when there are sufficiently many wavelengths available to activate any logical configuration, subject to the number of transceivers available at each node.

the physical topology.

Service is applied to the system at each time slot by activating a logical topology, and routing a packet across each active logical link. We denote the corresponding $m \times n$ *service activation matrix* by $\mathbf{S} = (S_{ej},\, e \in E_L,\, j \in V)$. Here, S_{ej} equals the number of logical links from node $\sigma(e)$ to node $\tau(e)$ used to service commodity j packets under the activation $\mathbf{S}$. Note that an admissible service activation matrix must have a valid underlying logical topology belonging to Π_N. This property characterizes the set of multi-hop service activation matrices, $\mathcal{S}^{\mathrm{mh}}$:

$$\mathcal{S}^{\mathrm{mh}} = \{\mathbf{S} \in \mathbb{Z}_+^{m\times n} : \pi_{\sigma(e)\tau(e)} = \textstyle\sum_{j\in V} S_{ej},\, \boldsymbol{\pi} \in \Pi_N\}.$$

In words, the above equation describes how each integer matrix in $\mathcal{S}^{\mathrm{mh}}$ implies the activation of an underlying logical topology $\boldsymbol{\pi} \in \Pi_N$, where $\pi_{\sigma(e)\tau(e)}$ equals the number of distinct logical links originating at node $\sigma(e)$ and terminating at $\tau(e)$ that are active to service *any* commodity.

The set $\mathcal{S}^{\mathrm{mh}}$ places no restriction on which commodity is allowed to cross an active logical link. This means that service activations belonging to $\mathcal{S}^{\mathrm{mh}}$ can correspond to *multi-hop routing*, where packets are re-enqueued after transmission across a logical link. In this paper, we will also deal with networks in which single-hop routing is exclusively employed. In such a situation, the set of admissible service activations is denoted $\mathcal{S}^{\mathrm{sh}}$, where $\mathbf{S} \in \mathcal{S}^{\mathrm{sh}}$ must satisfy

$$S_{ej} > 0 \text{ implies } j = \tau(e).$$

In words, the above statement means that a link can only be activated to service traffic directly to its destination node. Since the single-hop link activations are available in a multi-hop network, we must have $\mathcal{S}^{\mathrm{sh}} \subseteq \mathcal{S}^{\mathrm{mh}}$.

A service activation matrix $\mathbf{S}$ results in packet transitions through the network. To model the queue evolution implied by invoking $\mathbf{S}$, we introduce for each commodity $j \in V$ the $n \times m$ *routing matrix* $\mathbf{R}^j = (R^j_{ie},\, i \in V,\, e \in E_L)$, where:

$$R^j_{ie} = \begin{cases} 1, & \text{if } \sigma(e) = i \\ -1, & \text{if } \tau(e) = i \text{ and } i \neq j \\ 0, & \text{else} \end{cases}$$

Denote by $d_{ij}(\mathbf{S})$ the *net amount of service*, in number of packets per time slot, experienced by queue Q_{ij} under activation matrix $\mathbf{S}$. Using the above routing matrix we can express $d_{ij}(\mathbf{S}) = \sum_k \mathbf{R}^j_{ik}\mathbf{S}_{kj}$. Since $d_{ij}(\mathbf{S})$ is a net quantity, it can be positive, negative, or zero. We gather these values into the $n \times n$ matrix $\mathbf{d}(\mathbf{S}) = (d_{ij}(\mathbf{S}), i, j \in V)$.

A. Throughput considerations

The performance metric we study here is the network throughput, defined according to the *stability* criterion often referred to as *rate stability*. A system of queues is *rate stable* if [1]

$$\lim_{t\to\infty} Q_{ij}(t)/t = 0 \quad \text{w.p.1} \quad \forall i, j \in V.$$

Our choice of a rate stability criterion is because of its application to the widest possible class of arrival processes, and because it leads to a *closed* stability region of *stabilizable* arrival rate matrices. The results of this paper extend to strong stability (defined in [1]) following additional technical details.

The optimal throughput performance of the reconfigurable WDM network is characterized through the maximum stability region or *stability region* of the network. Since it is of interest in this work to understand the relative throughput performance of algorithms employing multi-hop electronic-layer routing versus algorithms exclusively employing single-hop electronic-layer routes, we distinguish two stability regions: one for achievable rates under multi-hop electronic-layer routing, and one for achievable rates under exclusively single-hop electronic-layer routing.

Definition 2.1: The *single-hop stability region* is denoted $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$. For any arrival process having long-term rate matrix $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\mathrm{sh}}$, there exists a reconfiguration and routing algorithm employing exclusively single-hop electronic routing under which the network is rate stable.

Definition 2.2: The *multi-hop stability region* is denoted $\boldsymbol{\Lambda}^*_{\mathrm{mh}}$. For any arrival process having long-term rate matrix $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\mathrm{mh}}$, there exists a reconfiguration and routing algorithm (possibly employing multi-hop electronic routing) under which the network is rate stable.

Definition 2.3: An algorithm is called *throughput optimal (achieves 100% throughput)* if the set of arrival rates that it can stabilize equals the multi-hop stability region.

In [31], Tassiulas and Ephremides introduced an algorithm that achieves 100% throughput in a general multi-hop-capable network. Their algorithmic description for scheduling in this network setting involves *maxweight decisions*, where at each time $t \geq 0$, the algorithm activates logical topology

$$\boldsymbol{\pi}^* \in \underset{\boldsymbol{\pi}\in\Pi_N}{\arg\max} \sum_{e\in E_L} \pi_{\sigma(e)\tau(e)} Z^*_e(t), \tag{1}$$

and for each active logical link e (having $\pi^*_e > 0$), electronically routes a packet of commodity $j^* \in \arg\max_j Z_{ej}(t)$ across e. In [31], it is shown that this scheduling policy achieves the multi-hop stability region $\boldsymbol{\Lambda}^*_{\mathrm{mh}}$, which can be expressed as the set of non-negative matrices in the convex hull of the available multi-hop service activations. Symbolically,

$$\begin{aligned} \boldsymbol{\Lambda}^*_{\mathrm{mh}} &= \mathbb{R}_+^{n\times n} \cap \mathrm{conv}\left(\{\mathbf{d}(\mathbf{S}) : \mathbf{S} \in \mathcal{S}^{\mathrm{mh}}\}\right) \qquad (2)\\ &= \mathbb{R}_+^{n\times n} \cap \left\{\textstyle\sum_{i=1}^{|\mathcal{S}^{\mathrm{mh}}|} \phi_i \mathbf{d}(\mathbf{S}^i) : \phi_i \geq 0, \sum_i \phi_i = 1\right\}. \end{aligned}$$

Above, we use $|\mathcal{S}^{\mathrm{mh}}|$ to represent the cardinality of the set $\mathcal{S}^{\mathrm{mh}}$, and index the elements of $\mathcal{S}^{\mathrm{mh}}$ with $\mathbf{S}^1, \ldots, \mathbf{S}^{|\mathcal{S}^{\mathrm{mh}}|}$.

When the set of service activations is restricted to single-hop routing, specifically to the set of single-hop service activation matrices $\mathcal{S}^{\mathrm{sh}}$, the algorithm of [31] reduces to activating at each time $t \geq 0$ the logical topology

$$\boldsymbol{\pi}^* \in \underset{\boldsymbol{\pi}\in\Pi_N}{\arg\max} \sum_{i,j\in V} \pi_{ij} Q_{ij}(t), \tag{3}$$

for single-hop electronic routing. In [30], it is shown that this scheduling policy achieves the single-hop stability region $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$, and that $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$ can be expressed as the convex hull of the set of logical topology matrices. Symbolically,

$$\boldsymbol{\Lambda}^*_{\mathrm{sh}} = \mathrm{conv}\left(\Pi_N\right) = \left\{\textstyle\sum_{i=1}^{|\Pi_N|} \phi_i \boldsymbol{\pi}^i : \phi_i \geq 0, \sum_i \phi_i = 1\right\}. \tag{4}$$

To be precise, we emphasize that [31] and [30] demonstrate stability over the regions $\mathbf{\Lambda}^*_{\mathrm{mh}}, \mathbf{\Lambda}^*_{\mathrm{sh}}$, respectively, up to a set of measure zero, under a stricter stability criterion than rate stability, and also under more restrictive assumptions on the arrival processes than what we employ here. In [3], it is proved that the maxweight scheduling policies (1), (3) achieve rate stability for any arrival process whose long-term rates belong to the *closed convex sets* $\mathbf{\Lambda}^*_{\mathrm{mh}}, \mathbf{\Lambda}^*_{\mathrm{sh}}$, respectively.

For context, note that in the case of a $n \times n$ input-queued switch, Π_N contains all permutation matrices of size n, and the stability region equals the *doubly substochastic region*:

$$\Lambda^{\mathrm{IQ}}_{\mathrm{sh}} = \left\{ \boldsymbol{\lambda} : \textstyle\sum_j \lambda_{i,j} \leq 1, \forall i, \ \sum_i \lambda_{i,j} \leq 1, \forall j \right\}. \tag{5}$$

The characterizations (2), (4) of the stability regions are as convex hulls over sets of service matrices. Contrast this with the characterization of the stability region of the input-queued switch in (5) as $2n$ linear inequality constraints. The stability polytopes in general network settings can have exponentially-many non-trivial linear inequality constraints, which leads to a difficulty of tractability in their characterization. In what follows, we develop a theory of *RWA decompositions*. This enables us to *exactly* characterize stability performance properties of wavelength-constrained reconfigurable WDM-based networks having *arbitrary* physical topologies. Thus our analysis establishes quantitative performance metrics that can be used by network designers to evaluate and compare varying network topologies and node functionalities.

B. Implementation considerations

A physical limitation of current reconfigurable optical components is their inherent latency. Thus, assuming current components, it may not be reasonable to assume that reconfiguration of the logical topology is achievable without idle time while transceivers are tuned. Additional implementation issues that can lead to idle time associated with reconfiguration include synchronization of communication along lightpaths and distribution of control information between network nodes. The effect of this *reconfiguration overhead* can be reduced by employing frame-based scheduling, where reconfiguration decisions are only made at frame boundaries as opposed to slot boundaries [5]. By selecting an appropriately large frame size, the network stability region in the presence of this overhead can be made arbitrarily close to the stability region in a network with no such overhead. This has been shown for unconstrained optical networks in [5], [25], and extends easily to our network model. We conclude that the analysis of this paper remains valid when the network is subject to reconfiguration overhead.

We emphasize that this paper focuses on quantifying the throughput properties of reconfigurable WDM networks. These properties are invariant, in that they hold as bounds on the limits of network performance irrespective of the scheduling algorithm employed. Thus, while the maxweight schedulers (1) and (3) are proven to achieve the multi-hop and single-hop stability regions, respectively, these are not the only scheduling algorithms with desirable throughput properties. Nevertheless, it is worthwhile to note that the key computational step of the maxweight schedulers is the selection of a maximum weight logical topology matrix, which, irrespective of single-hop or multi-hop capability, is in general NP-hard. This follows because the *maximum weight independent set problem*, a known NP-hard problem, can be shown to be polynomial-time reducible to the problem of maxweight logical topology selection. There are network configurations however, where polynomial-time solutions to the maxweight scheduling problem exist. For example, in wavelength-unconstrained networks, the maxweight scheduling problem reduces to finding a maximum weight matching in a weighted bipartite graph, an $O(n^3)$ operation (see *e.g.* [4], [5]).

III. RWA DECOMPOSITIONS

In this section, we demonstrate that in any optical network having a single wavelength per physical fiber link, the question of stability for a particular arrival rate matrix can be directly tied to the RWA problem on the same physical topology graph. Note that our work considers stability properties of single-wavelength optical networks. Yet, we use properties of the RWA for multi-wavelength optical networks to characterize the stability region of single-wavelength optical networks. We directly relate the RWA problem with no wavelength conversion to the set of achievable rates using only single-hop electronic routing, and the RWA problem with wavelength conversion to the set of achievable rates using multi-hop electronic layer routes.

A. The RWA problem

The objective of the RWA problem is to minimize the number of wavelengths needed to set up a certain set of lightpaths for a given physical topology. We consider two versions of the RWA problem: RWA with no wavelength conversion capability and RWA with full wavelength conversion capability.

Let $\mathbf{T} = (T_{ij})$ be a non-negative $n \times n$ integer lightpath demand matrix, where T_{ij} is the number of lightpaths, originating at node i and terminating at node j, that must be assigned. In the case of no wavelength conversion capability, the RWA is subject to the *wavelength continuity constraint*, which requires that no lightpath makes use of more than a single color from its source to its destination. In this case, let $W^{\mathrm{nc}}(\mathbf{T})$ be the minimum number of wavelengths required to service the demands of matrix $\mathbf{T}$ with no wavelength conversion (see Appendix A for details). As an example, consider the 3-node *unidirectional* ring physical topology having a single wavelength per optical fiber, and the lightpath demand matrix $\mathbf{T}$ given in Fig. 3(a). A valid RWA with no wavelength conversion is provided in Figure 3(b), It is easy to see for this network that $W^{\mathrm{nc}}(\mathbf{T}) = 4$.

A network node having full wavelength conversion capability can transform any pass-through lightpath, in the optical domain, from its incident wavelength to any other wavelength. In this case, we define $W^{\mathrm{c}}(\mathbf{T})$ to be the minimum number of wavelengths required to service the demands of $\mathbf{T}$ with wavelength conversion (see Appendix A for details). Since using a single color per lightpath is accommodated by the

$$\mathbf{T} = \begin{bmatrix} \cdot & 0 & 2 \\ 1 & \cdot & 0 \\ 1 & 1 & \cdot \end{bmatrix}$$

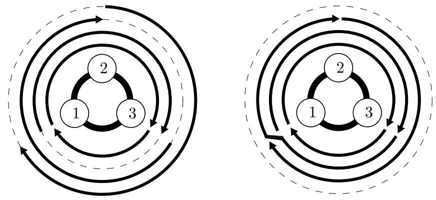

(a) Demand matrix (b) Without (c) With

Fig. 3. RWAs with and without wavelength conversion for traffic $\mathbf{T}$. The physical topology is a unidirectional ring (clockwise oriented). A dashed line indicates an idle wavelength on the corresponding fiber links.

RWA with wavelength conversion, it is clear for any physical topology that $W^{\mathrm{c}}(\mathbf{T}) \leq W^{\mathrm{nc}}(\mathbf{T})$ for all $\mathbf{T}$. For the trivial case of $\mathbf{T} = 0$, we define (for technical reasons) that $W^{\mathrm{nc}}(0) = W^{\mathrm{c}}(0) = 1$. For the traffic demand $\mathbf{T}$ of Fig. 3(a), Fig. 3(c) depicts the RWA employing wavelength conversion. In this case, $W^{\mathrm{c}}(\mathbf{T}) < W^{\mathrm{nc}}(\mathbf{T})$ (the inequality is strict).

B. Examples of RWA decompositions

In the RWA problem, multiple *single-wavelength* logical configurations are *multiplexed* through the use of frequency division (WDM). In our reconfigurable network setting, restricted to a single wavelength per optical fiber, multiple single-wavelength logical configurations are multiplexed through the use of time division (by enabling logical reconfiguration and adjustable electronic-layer routing over time). Through careful interchange of time and frequency, we can conceptually link the RWA problem to the stability issue in our reconfigurable network. Consequently, we will demonstrate how to transform a RWA for a particular wavelength traffic demand into a sequence of arrival rate matrices belonging to the network stability region, when the network N has a single wavelength per optical fiber. We next demonstrate this relationship with examples for both the single-hop and multi-hop scenarios.

1) Single-hop RWA decompositions: Consider the RWA with no wavelength conversion for traffic $\mathbf{T}$ in Fig. 3(a). The RWA of Figure 3(b) multiplexes the traffic demand $\mathbf{T}$ over 4 wavelengths. This RWA can be expressed as a *decomposition* of $\mathbf{T}$ into a superposition of single-wavelength logical topology configurations (expressed in matrix form) as follows,

$$\mathbf{T} = \begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 1 & 0 & \cdot \end{bmatrix} + \begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix} + \begin{bmatrix} \cdot & 0 & 0 \\ 0 & \cdot & 0 \\ 0 & 1 & \cdot \end{bmatrix} + \begin{bmatrix} \cdot & 0 & 0 \\ 1 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix}, \tag{6}$$

where the matrices from left to right represent the rings depicted in Figure 3(a) in order of increasing radius. Note that each of the matrices in the decomposition of (6) is a valid single-wavelength logical configuration.

Assuming there is a constant number $W \geq 4$ wavelengths available in each optical fiber, then we can say that each single-wavelength logical configuration in the RWA utilizes a fraction of $1/W$ of the total available *multiplexing resources* in the network. The *utilization* of the multiplexing resources is then given by $4/W \leq 1$.

We also consider *time as a multiplexing resource*; however, since we consider the evolution of our system over an infinite horizon, the time resource is normalized to unity. Consequently, when a particular single-wavelength logical configuration utilizes a fraction of the time resource, this is a measure of the long-term fraction of time spent servicing that logical configuration.

Consider equation (6), the valid RWA for traffic matrix $\mathbf{T}$ on the physical topology G_P, and re-interpret each wavelength configuration as utilizing $1/W$ of the available time resources in a single-wavelength network N. We have established that each wavelength configuration from the RWA is a valid single-wavelength logical topology and that the total utilization of multiplexing resources can be no more than 1. Consequently, we have validly multiplexed time in the single-wavelength network N. The resulting rate matrix corresponding to time sharing of service configurations is given for $W \geq 4$ by

$$\boldsymbol{\lambda}_W = \frac{1}{W}\mathbf{T} = \frac{1}{W}\begin{bmatrix} \cdot & 0 & 2 \\ 1 & \cdot & 0 \\ 1 & 1 & \cdot \end{bmatrix}.$$

Using (6), we have an explicit decomposition of $\boldsymbol{\lambda}_W$ into a convex combination of valid single-hop service matrices, subject to a single-wavelength per optical fiber,

$$\begin{aligned} \boldsymbol{\lambda}_W = \frac{1}{W}\begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 1 & 0 & \cdot \end{bmatrix} + \frac{1}{W}\begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix} + \frac{1}{W}\begin{bmatrix} \cdot & 0 & 0 \\ 0 & \cdot & 0 \\ 0 & 1 & \cdot \end{bmatrix} \\ + \frac{1}{W}\begin{bmatrix} \cdot & 0 & 0 \\ 1 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix} + \frac{W-4}{W}\begin{bmatrix} \cdot & 0 & 0 \\ 0 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix}. \end{aligned} \tag{7}$$

From the decomposition of (7), we can immediately conclude that $\boldsymbol{\lambda}_W \in \boldsymbol{\Lambda}^*_{\mathrm{sh}}$ for $W \geq 4$ (this follows directly from the definition of $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$). In words, the arrival rate matrix $\boldsymbol{\lambda}_W$ belongs to the single-hop stability region for any $W \geq 4$. We call this decomposition a *single-hop RWA decomposition* of $\boldsymbol{\lambda}_W$. In summary, by interchanging frequency and time, we have used a RWA for a particular wavelength traffic demand to produce a sequence of arrival rate matrices belonging to the single-hop stability region of N, when N has a single wavelength per optical fiber.

2) Multi-hop RWA decompositions: For the RWA with wavelength conversion, each wavelength routing can be considered a valid single-wavelength logical configuration. The difference from the RWA with no wavelength conversion is that lightpaths on a particular wavelength can have endpoints on that wavelength, corresponding to the use of a wavelength converter. We can re-interpret the RWA problem in our reconfigurable setting by noting that while the RWA problem uses wavelength converters to take advantage of available resources at *different frequencies* (equivalently, wavelengths), our reconfigurable network uses electronic-layer queues to take advantage of available resources at *different times*. Thus, wherever RWA invokes a wavelength converter, the reconfigurable network can be understood to terminate a lightpath at that node and electronically enqueue the carried data for multi-hop transmission to its destination at a different time.

We demonstrate multi-hop RWA decompositions in the following example. Consider the RWA problem for the wavelength traffic demand $\mathbf{T}$ in Fig. 3(a). We have established that wavelength conversion can be used to service $\mathbf{T}$ with only 3 wavelengths, as depicted in Figure 3(c). This RWA can be expressed as the following decomposition, with the matrices successively representing the rings depicted in Figure 3(b) in order of increasing radius,

$$\mathbf{T} = \begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 1 & 0 & \cdot \end{bmatrix} + \begin{bmatrix} \cdot & -1 & 1 \\ 0 & \cdot & 0 \\ 0 & 1 & \cdot \end{bmatrix} + \begin{bmatrix} \cdot & 1 & 0 \\ 1 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix}. \quad (8)$$

The above decomposition can be interpreted as follows. The first wavelength fully services demands $\{1 \to 3, 3 \to 1\}$. The second wavelength services demand $1 \to 3$ and services demand $3 \to 2$ *only up to node* 1. Consequently, the '-1' in the second matrix of (8) represents the $3 \to 2$ traffic that is enqueued for multi-hop transmission at node 1. The third wavelength services the remainder of demand $3 \to 2$ from node 1 to node 2 as well as fully servicing demand $2 \to 1$.

Thus, for $W \geq 3$, the arrival rate matrix $\boldsymbol{\lambda}_W = (1/W)\mathbf{T}$ can be expressed using (8) as a convex combination of valid single-wavelength multi-hop service matrices,

$$\boldsymbol{\lambda}_W = \frac{1}{W}\begin{bmatrix} \cdot & 0 & 1 \\ 0 & \cdot & 0 \\ 1 & 0 & \cdot \end{bmatrix} + \frac{1}{W}\begin{bmatrix} \cdot & -1 & 1 \\ 0 & \cdot & 0 \\ 0 & 1 & \cdot \end{bmatrix} + \frac{1}{W}\begin{bmatrix} \cdot & 1 & 0 \\ 1 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix} + \frac{W-3}{W}\begin{bmatrix} \cdot & 0 & 0 \\ 0 & \cdot & 0 \\ 0 & 0 & \cdot \end{bmatrix}.$$

We conclude that $\boldsymbol{\lambda}_W \in \boldsymbol{\Lambda}^*_{\mathrm{mh}}$ for $W \geq 3$.

IV. Capacity regions from RWA decompositions in single-wavelength networks

The examples of the previous section have shown how the RWA with and without wavelength conversion *for a single traffic demand* $\mathbf{T}$ can be translated to a sequence of arrival rate matrices belonging to the single-hop and multi-hop stability regions, respectively. In this section, we will demonstrate that the single-hop and multi-hop stability regions for single-wavelength optical networks can be fully described by the RWA functions W^{nc} and W^{c}, respectively.

A. Single-hop stability region

We begin by considering the single-hop stability region. In networks with no wavelength constraints, this region is characterized in [4], [5]. In [34], [35], this region is studied as the stability region of general optical flow switched networks. Our characterization, which is exclusive to single-wavelength networks, is useful in our subsequent development of geometric properties of the stability region. In particular, it allows us to express the entire capacity region as a collection of limit points based on the solution to the RWA problem.

The example of Section III-B.1 provided a sequence of arrival rates belonging to $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$ for a single integer traffic demand matrix $\mathbf{T}$ in the RWA problem with no conversion. In this section we consider all such arrival rates, gathered over all possible integer traffics $\mathbf{T}$ in the RWA problem. Let $\mathcal{R}^{\mathrm{nc}}$ be the set of all such arrival rates,

$$\mathcal{R}^{\mathrm{nc}} = \left\{ \boldsymbol{\lambda} = \frac{1}{W}\mathbf{T} : \mathbf{T} \in \mathbb{Z}_+^{n\times n}, W \in \mathbb{Z}_+, W \geq W^{\mathrm{nc}}(\mathbf{T}) \right\}. \quad (9)$$

Recall that we are restricting attention to joint optical reconfiguration and electronic layer routing algorithms where the optical layer has only a single wavelength available in each optical fiber. Consequently, $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$ is the single-hop stability region of the single-wavelength network N.

For the set $\mathcal{R}$, let $\mathrm{cl}(\mathcal{R})$ represent the closure[3] of $\mathcal{R}$. We next establish that every matrix in $\mathrm{cl}(\mathcal{R}^{\mathrm{nc}})$ belongs to $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$, and conversely, that every matrix in $\boldsymbol{\Lambda}^*_{\mathrm{sh}}$ belongs to $\mathrm{cl}(\mathcal{R}^{\mathrm{nc}})$.

Theorem 4.1: $\boldsymbol{\Lambda}^*_{\mathrm{sh}} = \mathrm{cl}(\mathcal{R}^{\mathrm{nc}})$

Proof: See Appendix B. ■

B. Multi-hop stability region

The multi-hop stability region is characterized in a similar manner. In [4], [5], this region is characterized for networks having no wavelength constraints, where it is shown that the single-hop and multi-hop stability regions are equal. In [34], [35], a queueing model is enlisted to study the throughput properties of optical packet switched networks (OPS). The OPS stability region of [34], [35] is related to the multi-hop stability region of our reconfigurable optical network, with differences arising depending on the set of available optical layer network configurations Π_N. Our characterization in the single-wavelength setting is tailored to our subsequent analysis of geometric properties of the multi-hop stability region.

Here, we gather all possible arrival rates generated by multi-hop RWA decompositions over all possible traffic demand matrices $\mathbf{T}$ into the set $\mathcal{R}^{\mathrm{c}}$,

$$\mathcal{R}^{\mathrm{c}} = \left\{ \boldsymbol{\lambda} = \frac{1}{W}\mathbf{T} : \mathbf{T} \in \mathbb{Z}_+^{n\times n}, W \in \mathbb{Z}_+, W \geq W^{\mathrm{c}}(\mathbf{T}) \right\}. \quad (10)$$

Through similar steps as in the single-hop case, we can establish the following theorem.

Theorem 4.2: $\boldsymbol{\Lambda}^*_{\mathrm{mh}} = \mathrm{cl}(\mathcal{R}^{\mathrm{c}})$

Proof: The proof is similar to the proof of Theorem 4.1, and is omitted for brevity. ■

V. Geometric properties of the stability region

While the stability properties of our dynamically reconfigurable electronic-over-optical network are well characterized in the multi-hop and single-hop cases through equations (2) and (4), respectively, these expressions do not easily yield simple geometric properties of the stability regions. This is in contrast to the characterization of the input-queued switch stability region of equation (5).

The remainder of this work is dedicated to extracting geometric properties of the single-hop and multi-hop stability regions in the wavelength-constrained WDM network setting.

[3] An accumulation point of $\mathcal{R}$ is such that there exist other points of $\mathcal{R}$ arbitrarily close by. The closure of $\mathcal{R}$ is then given by the union of $\mathcal{R}$ and all its accumulation points [19].

In what follows, we will occasionally refer to the wavelength-unconstrained network setting. From our assumption that node $v \in V$ has P_v transceivers available, when the network has no wavelength constraint, the stability region (single-hop and multi-hop) equals [3]–[5]

$$\mathbf{\Lambda}_{\text{port}} = \left\{ \boldsymbol{\lambda} : \textstyle\sum_j \lambda_{ij} \leq P_i \, \forall i, \, \sum_i \lambda_{ij} \leq P_j \, \forall j \right\}. \tag{11}$$

A. Maximum uniform arrival rate matrices

In this section, we make use of RWA decompositions to establish geometric properties of the single-hop and multi-hop stability regions. Define $\mathbf{J}$ as the $n \times n$ matrix having (i, j) entry equal to 1 if $i \neq j$:

$$\mathbf{J} = \begin{bmatrix} \cdot & 1 & \cdots & 1 \\ 1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & 1 \\ 1 & \cdots & 1 & \cdot \end{bmatrix}.$$

We then seek to determine the maximum values $\theta^{\text{sh}}, \theta^{\text{mh}}$ such that $\theta^{\text{sh}}\mathbf{J}$ belongs to the single-hop stability region, and $\theta^{\text{mh}}\mathbf{J}$ belongs to the multi-hop stability region.

Theorem 5.1: For network N having a single wavelength per optical fiber, let $\theta^{\text{sh}} = \sup\{\theta : \theta\mathbf{J} \in \mathbf{\Lambda}^*_{\text{sh}}\}$. Then,

$$\theta^{\text{sh}} = \limsup_{k \to \infty} k / W^{\text{nc}}(k\mathbf{J}). \tag{12}$$

For the multi-hop scenario, let $\theta^{\text{mh}} = \sup\{\theta : \theta\mathbf{J} \in \mathbf{\Lambda}^*_{\text{mh}}\}$. Then,

$$\theta^{\text{mh}} = \limsup_{k \to \infty} k / W^{\text{c}}(k\mathbf{J}). \tag{13}$$

Proof: See Appendix C. ∎

Equations (12) and (13) essentially capture the maximum ratio of the uniform traffic load l to the number of wavelengths needed to support that traffic demand. These values are a measure of the most efficient way that the uniform traffic demand l can be packed over network N, with or without wavelength conversion.

Theorem 5.1 allows us to draw on the literature regarding RWA algorithms for various physical topologies to obtain geometric properties of the single-hop and multi-hop stability regions. As an example, consider the unidirectional ring having a single transceiver per node ($P_i = 1$). In this case, it can be shown that the minimum numbers of wavelengths required to service traffic $l\mathbf{J}$ with or without wavelength conversion are equal: $W^{\text{nc}}(l\mathbf{J}) = W^{\text{c}}(l\mathbf{J}) = nl(n-1)/2$. Applying (12) and (13) we obtain a maximum uniform arrival rate of $\theta^{\text{sh}} = \theta^{\text{mh}} = 2/(n^2 - n)$. Thus, there is no single-hop versus multi-hop performance gap for uniform arrival rates under the unidirectional ring. However, noting in the wavelength-unconstrained case (see (11)), the maximum uniform arrival rate is given by $\theta^{\max} = 1/(n-1)$, we find a constrained versus unconstrained performance gap of $2/n = O(1/n)$.[4]

[4]We employ O-notation to represent an *asymptotically tight bound* [13] on the performance gap.

We draw the RWA values $W^{\text{nc}}(l\mathbf{J}), W^{\text{c}}(l\mathbf{J})$ from [10], [26]–[28], and summarize the single-hop and multi-hop maximum uniform arrival rates for several physical topologies in Table I. The table lists the maximum uniform arrival rates achievable in the single-wavelength setting, as well as the corresponding maximum uniform arrival rate achievable in the wavelength-unconstrained case, $\theta^{\max}$, and the implied unconstrained versus constrained performance gap. For the tree topology $\mathcal{T}$, denote $\mathcal{N}_{e,1}, \mathcal{N}_{e,2}$ as the node sets in the cut corresponding to edge $e \in \mathcal{T}$.

A remarkable property evident from Table I is that for all physical topologies considered, there is no single-hop versus multi-hop performance gap with respect to uniform arrival rates. This follows for all physical topologies considered in Table I, because under uniform traffic demand, RWA with and without wavelength conversion can achieve the same minimum number of wavelengths. It is conjectured in [26] that this result holds generally over all physical topologies.

Note that *the geometric properties listed in the table are exact*. For physical topologies besides rings, trees, tori, hypercubes, and others where the solution to the RWA problem is known, the exact characterizations of (12) and (13) can be approximated through evaluation of the RWA functions over multiple all-to-all integer traffic demands. Techniques for solving the integer RWA problem are well-studied in the literature. In [12], various RWA methodologies are classified, based on their optimization criteria, and their approach to solving the problem. Additional comments regarding the solution to the RWA problem can be found in Appendix A.

B. Maximum scaled doubly substochastic set

In this section, we take advantage of RWA decompositions to derive bounds on the maximum scaling that can be applied to the set of doubly substochastic matrices, such that every matrix in the scaled set is contained within the stability region. For a mathematical description of this property we require the following definitions.

Definition 5.1: For matrix $\mathbf{A}$, let the maximum row/column sum of $\mathbf{A}$ be given by $\|\mathbf{A}\|_{\max}$:

$$\|\mathbf{A}\|_{\max} = \max\left\{ \max_i \textstyle\sum_j A_{ij}, \max_j \sum_i A_{ij} \right\}.$$

Definition 5.2: Let the set $\mathcal{D}_s$ denote the doubly substochastic region, scaled by factor s,

$$\mathcal{D}_s = \left\{ \boldsymbol{\lambda} \in \mathbb{R}_+^{n \times n} : \|\boldsymbol{\lambda}\|_{\max} \leq s \right\}.$$

We seek the maximum values $\alpha^{\text{sh}}, \alpha^{\text{mh}}$ such that the sets $\mathcal{D}_{\alpha^{\text{sh}}}, \mathcal{D}_{\alpha^{\text{mh}}}$ are respectively subsets of the single-hop and multi-hop stability regions. We will demonstrate that there are cases in which the multi-hop stability region provides improved performance over the single-hop stability region, in terms of this geometric property. Consequently, we can conclude that there are indeed cases in which multi-hop routing can provide a strict throughput performance improvement over algorithms that exclusively employ single-hop routes. This is in contrast to the case of a crossbar switch, where single-hop algorithms can achieve the stability region.

Definition 5.3: The integer matrix $\mathbf{T} = (T_{ij}) \in \mathbb{Z}_+^{n \times n}$ is called k-allowable if it satisfies $\|\mathbf{T}\|_{\max} \leq k$. Let $\mathcal{K}_k$ be the set of all k-allowable matrices.

TABLE I

MAXIMUM VALUES $\theta^{\rm sh}, \theta^{\rm mh}$ FOR VARIOUS PHYSICAL TOPOLOGIES HAVING A SINGLE WAVELENGTH PER OPTICAL FIBER. THE CORRESPONDING WAVELENGTH-UNCONSTRAINED VALUES ARE LISTED UNDER $\theta^{\rm max}$, ALONG WITH THE RESULTING THROUGHPUT PERFORMANCE GAP.

Physical topology G_P	P_i, $\forall i$	$W^{\rm nc}(l\mathbf{J}) = W^{\rm c}(l\mathbf{J})$	$\theta^{\rm sh} = \theta^{\rm mh}$	$\theta^{\rm max}$	Performance gap $\frac{\theta^{\rm mh}}{\theta^{\rm max}}$
Tree $\mathcal{T}$	1	$l \max_{e\in\mathcal{T}} \lvert\mathcal{N}_{e,1}\rvert\lvert\mathcal{N}_{e,2}\rvert$	$1/(\max_{e\in\mathcal{T}} \lvert\mathcal{N}_{e,1}\rvert\lvert\mathcal{N}_{e,2}\rvert)$	$1/(n-1)$	$\dfrac{n-1}{\max_{e\in\mathcal{T}} \lvert\mathcal{N}_{e,1}\rvert\lvert\mathcal{N}_{e,2}\rvert}$
Unidirectional ring	1	$nl(n-1)/2$	$2/(n^2-n)$	$1/(n-1)$	$2/n = O(1/n)$
Bidirectional ring					
n odd	2	$l(n^2-1)/8$	$8/(n^2-1)$	$2/(n-1)$	$4/(n+1) = O(1/n)$
n even	2	$\lceil ln^2/8\rceil$	$8/n^2$	$2/(n-1)$	$4(n-1)/n^2 = O(1/n)$
2D Torus R rows, C cols (R, C divisible by 4)	4	$\lceil lRC(R+C)/16\rceil$	$16/(RC(R+C))$	$4/(n-1)$	$\frac{4(n-1)}{n(R+C)} = O(1/(R+C))$
Binary hypercube	$\log_2 n$	$nl/2$	$2/n$	$(\log_2 n)/(n-1)$	$\frac{2(n-1)}{n\log_2 n} = O(1/\log_2 n)$

Let $\mathcal{W}^{\rm nc}(k)$ be the minimum number of wavelengths required to service *any* k-allowable traffic matrix in the RWA with no conversion: $\mathcal{W}^{\rm nc}(k) = \max_{\mathbf{T}\in\mathcal{K}_k} W^{\rm nc}(\mathbf{T})$. Similarly, let the corresponding value with wavelength conversion be $\mathcal{W}^{\rm c}(k)$. The RWA problem for k-allowable matrices was introduced in [15] and subsequently studied in [10], [26]–[28]. These papers seek to understand the values of the quantities $\mathcal{W}^{\rm nc}(k), \mathcal{W}^{\rm c}(k)$ for various physical topologies. The bidirectional ring with no wavelength conversion is considered in [15], [28], tree topologies with no wavelength conversion were considered in [15], [27], and ring and torus topologies with wavelength conversion were considered in [10]. Additional results for k-allowable traffics can be found in [26].

The following theorem establishes the quantity $\alpha^{\rm sh}$ as the maximum scale factor on the substochastic region, such that the scaled region is a subset of the single-hop stability region. The analogous result for the multi-hop case is also provided.

Theorem 5.2: Let $\alpha^{\rm sh} = \sup\{\alpha : \mathcal{D}_\alpha \subseteq \mathbf{\Lambda}^*_{\rm sh}\}$. Then,

$$\alpha^{\rm sh} = \limsup_{k\to\infty} k/\mathcal{W}^{\rm nc}(k). \tag{14}$$

Similarly, let $\alpha^{\rm mh} = \sup\{\alpha : \mathcal{D}_\alpha \subseteq \mathbf{\Lambda}^*_{\rm mh}\}$. Then,

$$\alpha^{\rm mh} = \limsup_{k\to\infty} k/\mathcal{W}^{\rm c}(k). \tag{15}$$

Proof: See Appendix D. ∎

Equations (14) and (15) provide the limiting ratios of k to the worst-case number of wavelengths required to support any k-allowable traffic, in their respective RWA problems. This is a measure of the most efficient way that the worst-case k-allowable traffic can be packed over network N, in the limit of large k.

Applying Theorem 5.2, we can use results from the RWA literature [9], [10], [22], [26]–[28] to characterize the values $\alpha^{\rm sh}, \alpha^{\rm mh}$ for various physical topology configurations. Consider for example the bidirectional ring having an even number $n \geq 8$ nodes. For the RWA with no wavelength conversion, the worst-case k-allowable traffic requires $\lceil kn/3\rceil$ wavelengths, resulting in a maximum scaling of $\alpha^{\rm sh} = 3/n$. The RWA with wavelength conversion requires at most $\lceil kn/4\rceil$ wavelengths for any k-allowable traffic, yielding $\alpha^{\rm mh} = 4/n$. Consequently, we have a single-hop versus multi-hop performance gap of $3/4$, irrespective of the number of nodes in the network. Designating the maximum scale value achievable in the wavelength-unconstrained case by $\alpha^{\rm max}$, we note that the bidirectional ring has $\alpha^{\rm max} = 2$, since the architecture employs two transceivers per node (one for each incident fiber). This yields a constrained versus unconstrained performance gap in the unidirectional ring of $2/n$. Our results for various physical topologies are summarized in Table II. Note that the value of $\mathcal{W}^{\rm c}(k)$ for a bidirectional ring when n is odd remains an open problem. Consequently, Table II provides the tightest known interval in which this value resides [9], and the interval in which $\alpha^{\rm mh}$ resides. The lower limit of this interval is derived based on the next theorem (see Theorem 5.3 and the subsequent discussion). Also note that for the tree network, throughput performance depends on the tree topology employed, and particularly on the worst-case cut that maximizes the number of nodes on the smaller side of the cut. We call this number $c_{\mathcal{T}}$. Recalling our definition of $\mathcal{N}_{e,1}, \mathcal{N}_{e,2}$ as the node sets in the cut corresponding to edge e, we have $c_{\mathcal{T}} \triangleq \max_{e\in\mathcal{T}} \min\{|\mathcal{N}_{e,1}|, |\mathcal{N}_{e,2}|\}$.

Theorem 5.2 provides an *exact* characterization of the maximum scaled doubly substochastic region fully contained within $\mathbf{\Lambda}^*_{\rm mh}$. If an order bound is sufficient, then we can use [18, Lem. 1] to provide the following connection between the geometric properties studied in this section.

Theorem 5.3: $n\theta^{\rm mh}/2 \leq \alpha^{\rm mh} \leq (n-1)\theta^{\rm mh}$

Proof: Lemma 1 of [18] can be understood in our reconfigurable WDM network setting as follows: if $\theta\mathbf{J} \in \mathbf{\Lambda}^*_{\rm mh}$, then $\mathcal{D}_\alpha \subseteq \mathbf{\Lambda}^*_{\rm mh}$ when $\alpha \leq n\theta/2$. The lower bound follows. The upper bound follows since $(\alpha^{\rm mh}/(n-1))\mathbf{J} \in \mathcal{D}_{\alpha^{\rm mh}} \subseteq \mathbf{\Lambda}^*_{\rm mh}$, which implies $\theta^{\rm mh} \geq \alpha^{\rm mh}/(n-1)$. ∎

Theorem 5.3 allows us to obtain a refined bound on $\alpha^{\rm mh}$ for the bidirectional ring when n is odd. For this physical topology, Theorem 5.3 provides that $\alpha^{\rm mh} \geq 4n/(n^2-1)$. Based only on the fact (from Table II) that $\lceil k(n-1)/4\rceil \leq \mathcal{W}^{\rm c}(k) \leq \lceil kn/4\rceil$, we find that $4/n \leq \alpha^{\rm mh} \leq 4/(n-1)$. However, since $4n/(n^2-1) > 4/n$ for $n \geq 2$, we can obtain the refined bound, $4n/(n^2-1) \leq \alpha^{\rm mh} \leq 4/(n-1)$.

TABLE II

MAXIMUM VALUES $\alpha^{\rm sh}, \alpha^{\rm mh}$ FOR VARIOUS PHYSICAL TOPOLOGIES HAVING A SINGLE WAVELENGTH PER OPTICAL FIBER. ALSO LISTED FOR EACH TOPOLOGY IS THE SINGLE-HOP VERSUS MULTI-HOP PERFORMANCE GAP, AS WELL AS THE CONSTRAINED VERSUS UNCONSTRAINED PERFORMANCE GAP.

Physical topology G_P	$\mathcal{W}^{\rm nc}(k)$	$\alpha^{\rm sh}$	$\mathcal{W}^{\rm c}(k)$	$\alpha^{\rm mh}$	$\alpha^{\rm sh}/\alpha^{\rm mh}$	$\alpha^{\rm mh}/\alpha^{\rm max}$
Star	k	1	k	1	1	1
Tree $\mathcal{T}$	$kc_{\mathcal{T}}$	$1/c_{\mathcal{T}}$	$kc_{\mathcal{T}}$	$1/c_{\mathcal{T}}$	1	$1/c_{\mathcal{T}}$
Unidirectional ring	kn	$1/n$	$k(n-1)$	$1/(n-1)$	$1-1/n$	$1/(n-1)$
Bidirectional ring ($n \geq 7$)						
n odd	$\lceil kn/3 \rceil$	$3/n$	$\left\lceil \frac{k(n-1)}{4} \right\rceil \leq \mathcal{W}^{\rm c}(k) \leq \left\lceil \frac{kn}{4} \right\rceil$	$\frac{4n}{n^2-1} \leq \alpha^{\rm mh} \leq \frac{4}{n-1}$	$\leq \frac{3}{4} - \frac{3}{4n^2}$	$\leq \frac{2}{n-1}$
n even	$\lceil kn/3 \rceil$	$3/n$	$\lceil kn/4 \rceil$	$4/n$	$3/4$	$2/n$

A similar statement to Theorem 5.3 cannot be made for the quantity $\alpha^{\rm sh}$, because the argument of [18, Lem. 1] is inherently a multi-hop result.

VI. EXTENSIONS

A. Additional geometric properties

Theorem 5.2 can be easily extended to provide for any polytope the maximum scale factor such that the scaled polytope remains within the stability region. The proof of the following theorem is similar to that of Theorem 5.2, and can be found in [3].

Theorem 6.1: Let $\mathcal{P}$ be a convex, compact, full-dimensional subset of $\mathbb{R}_+^{n\times n}$, and $\alpha_{\mathcal{P}}^{\rm sh} = \sup\{\alpha : \alpha\boldsymbol{\lambda} \in \boldsymbol{\Lambda}_{\rm sh}^*, \forall \boldsymbol{\lambda} \in \mathcal{P}\}$. Then

$$\alpha_{\mathcal{P}}^{\rm sh} = \limsup_{k\to\infty} k/\mathcal{W}_{\mathcal{P}}^{\rm nc}(k),$$

where $\mathcal{W}_{\mathcal{P}}^{\rm nc}(k) = \max_{\mathbf{T}\in\mathcal{K}_k^{\mathcal{P}}} W^{\rm nc}(\mathbf{T})$, and $\mathcal{K}_k^{\mathcal{P}} = \mathbb{Z}_+^{n\times n} \cap k\mathcal{P}$. An identical result applies for the multi-hop case.

Theorem 6.1 can be used to recover the result of Theorem 5.2, by employing the set $\mathcal{P} = \left\{\boldsymbol{\lambda} \in \mathbb{R}_+^{n\times n} : \|\boldsymbol{\lambda}\|_{\max} \leq 1\right\}$.

B. Multi-wavelength WDM networks

Our network model in Section II is sufficiently general that it applies much more broadly than in WDM networks having a single wavelength per optical fiber. In particular, the model can accommodate any number of wavelengths available in each fiber, and other architectural assumptions that affect the logical topologies and electronic routing allowed in the network. For such a network, designate by $\mathcal{S}$ the set of allowable service activation matrices. Recall from Section II that every matrix belonging to $\mathcal{S}$ jointly represents a valid logical topology and electronic routing. As in definitions 2.1 and 2.2, we designate by $\boldsymbol{\Lambda}^*$ the stability region of arrival rates that can be rate stabilized when the service activation set in network N is $\mathcal{S}$.

The following definition generalizes the RWA functions $W^{\rm nc}, W^{\rm c}$ to this more general network setting.

Definition 6.1: For the non-negative integer matrix $\mathbf{T} = (T_{ij})$, let $\chi(\mathbf{T})$ equal the minimum number of service activation matrices belonging to $\mathcal{S}$ required to dominate $\mathbf{T}$:

$$\chi(\mathbf{T}) = \min\left\{k : \exists \mathbf{S}^1, \ldots, \mathbf{S}^k \in \mathcal{S},\, \mathbf{T} \leq \textstyle\sum_{l=1}^k \mathbf{d}(\mathbf{S}^l)\,\forall i,j\right\}.$$

The following theorems generalize Theorems 4.1, 4.2, 5.1, and 5.2 to multi-wavelength networks. Their proofs follow identically to the single-wavelength proofs, only replacing the RWA function $W^{\rm nc}(\cdot)$ with $\chi(\cdot)$.

Theorem 6.2: Define the set $\mathcal{R}$ as the set of integer traffic matrices scaled by their respective χ values,

$$\mathcal{R} = \left\{\boldsymbol{\lambda} = \frac{1}{W}\mathbf{T} : \mathbf{T} \in \mathbb{Z}_+^{n\times n},\, W \in \mathbb{Z}_+,\, W \geq \chi(\mathbf{T})\right\}.$$

Then $\boldsymbol{\Lambda}^* = \mathrm{cl}(\mathcal{R})$.

Theorem 6.3: Define $\theta^* = \sup\{\theta : \theta\mathbf{J} \in \boldsymbol{\Lambda}^*\}$. Then

$$\theta^* = \limsup_{k\to\infty} k/\chi(k\mathbf{J})$$

Theorem 6.4: Define $\alpha^* = \sup\{\alpha : \alpha\mathcal{D} \subseteq \boldsymbol{\Lambda}^*\}$. Then

$$\alpha^* = \limsup_{k\to\infty} k/\max_{\mathbf{T}\in\mathcal{K}_k} \chi(\mathbf{T})$$

While these theorems broaden the class of networks to which generalized RWA decompositions can be applied to characterize network stability properties, their key drawback is that they rely on the function χ, which does not in general tie to a well-studied optimization problem. This is in contrast to the special case of single-wavelength networks, where we characterized the stability region and geometric properties in terms of the well-known RWA problem.

Because the function χ may be difficult in general to characterize, we can use the single-wavelength geometric properties to bound their multi-wavelength counterparts. It can be shown (see [3] for details) that the stability region when the network has w wavelengths available in each optical fiber, denoted $\boldsymbol{\Lambda}_{{\rm mh},w}^*$, satisfies

$$\boldsymbol{\Lambda}_{{\rm mh},w}^* \subseteq \mathrm{conv}\left(w\boldsymbol{\Lambda}_{\rm mh}^* \cap \boldsymbol{\Lambda}_{\rm port} \cap \mathbb{Z}_+^{n\times n}\right).$$

This outer bound in combination with Theorems 5.1 and 5.2 provides simple outer bounds for the throughput properties of w-wavelength networks, based on the single-wavelength characterizations of this paper.

VII. SUMMARY

We have studied the optimal throughput performance properties of reconfigurable WDM-based packet networks. We developed a theory of RWA decompositions that establishes the stability regions of WDM networks having single-hop and multi-hop routing capability in terms of the RWA problem.

This characterization enabled us to *exactly* determine certain geometric properties of the stability region under any physical topology, restricted to a single-wavelength per optical fiber:

the maximum all-to-all arrival rate and maximum doubly substochastic region that can be supported by the network. We presented closed-form solutions for certain network topologies such as rings, trees, and tori.

These geometric properties provide a measure of the optimal achievable throughput under any physical topology. Consequently, a network designer could use such a metric in comparing and evaluating network topologies and/or varying node functionality. For example, we have *exactly* demonstrated the throughput performance gap between wavelength-limited and wavelength-unconstrained networks having particular physical topologies. Additionally, we have exactly characterized the throughput performance gap between networks employing exclusively single-hop routing and those employing multi-hop routing. In the case of the bidirectional ring, we have observed a performance improvement of 33% of multi-hop over single-hop enabled networks.

Appendix A
The RWA optimization

The RWA problem with full wavelength conversion is an integer multicommodity flow problem, which can be formulated as follows [21]. Let $\mathbf{T} = (T_{ij}) \in \mathbb{Z}_+^{n\times n}$ represent the set of lightpath demands, and let f_{ij}^e be a flow variable that represents the number of lightpaths from node i to node j that cross the fiber link e. For physical topology graph G_P, let E_v^σ be the set of edges originating at node v: $E_v^\sigma = \{e \in E_P : \sigma(e) = v\}$. Similarly, let E_v^τ denote the set of edges terminating at node v: $E_v^\tau = \{e \in E_P : \tau(e) = v\}$.

$$\min \quad W \tag{16}$$

$$\text{s.t.} \quad W \geq \textstyle\sum_{i,j\in V} f_{ij}^e, \quad \forall e \in E_P \tag{17}$$

$$\sum_{e\in E_v^\sigma} f_{ij}^e - \sum_{e\in E_v^\tau} f_{ij}^e = \begin{cases} T_{ij} & v = i \\ -T_{ij} & v = j \\ 0 & \text{else} \end{cases} \quad \forall v, i, j \in V \tag{18}$$

$$f_{ij}^e \in \mathbb{Z}_+, \quad \forall i, j \in V, e \in E_P \tag{19}$$

The minimum value W reached by the optimization is $W^{\text{c}}(\mathbf{T})$.

The RWA problem with no wavelength conversion can be formulated through the addition of the following constraints in the optimization (16)-(19), which impose the *wavelength-continuity constraint* on the RWA problem.

$$f_{ij}^e = \textstyle\sum_{w=1}^W c_{ij}^{e,w} \quad \forall i, j \in V, e \in E_P$$

$$\sum_{e\in E_v^\sigma} c_{ij}^{e,w} - \sum_{e\in E_v^\tau} c_{ij}^{e,w} \begin{cases} \geq 0 & v = i \\ \leq 0 & v = j \\ = 0 & \text{else} \end{cases} \quad \forall v, i, j \in V$$

$$c_{ij}^{e,w} \in \{0, 1\} \quad \forall i, j \in V, e \in E_P, w \in \{1, \ldots, W\}$$

The minimum value W reached by this optimization is $W^{\text{nc}}(\mathbf{T})$.

Commonly, the RWA problem is solved in two stages, first by solving the lightpath routing problem, followed by obtaining a wavelength assignment for the routing determined in the first step [12]. The routing problem can be solved sequentially using shortest-path algorithms, or through standard integer programming solution methods such as randomized rounding. The wavelength assignment algorithm is typically studied as a graph coloring algorithm, with common approaches to the problem including sequential assignment, genetic algorithms, simulated annealing, and randomized rounding. See [12] and the references contained therein for details.

Appendix B
Proof of Theorem 4.1

Here we divide the proof as follows. First we demonstrate that $\text{cl}(\mathcal{R}^{\text{nc}}) \subseteq \mathbf{\Lambda}_{\text{sh}}^*$, and second we prove $\mathbf{\Lambda}_{\text{sh}}^* \subseteq \text{cl}(\mathcal{R}^{\text{nc}})$.

Proof that $\text{cl}(\mathcal{R}^{\text{nc}}) \subseteq \mathbf{\Lambda}_{\text{sh}}^*$*:* Suppose $\boldsymbol{\lambda} \in \mathcal{R}^{\text{nc}}$. Then from (9) there must exist $\mathbf{T}, W$ such that $\boldsymbol{\lambda} = (1/W)\mathbf{T}$, with $\mathbf{T} \in \mathbb{Z}_+^{n\times n}$ and $W \geq W^{\text{nc}}(\mathbf{T})$. We establish a RWA decomposition for $\boldsymbol{\lambda}$ as a subconvex combination of $W^{\text{nc}}(\mathbf{T})$ matrices as follows. For each index $i = 1, \ldots, W^{\text{nc}}(\mathbf{T})$, we construct the matrix $\tilde{\boldsymbol{\pi}}^i$, corresponding to a valid single-wavelength logical topology configuration: let $\tilde{\pi}_{kl}^i = 1$ if logical link $k \to l$ is enabled on the i-th color of the RWA of $\mathbf{T}$ employing $W^{\text{nc}}(\mathbf{T})$ wavelengths, and $\tilde{\pi}_{kl}^i = 0$ otherwise. Clearly $\tilde{\boldsymbol{\pi}}^i$ is a valid logical topology subject to the single-wavelength constraint, since the same configuration had a valid routing on the i-th color under the RWA of $\mathbf{T}$. Let the elements of Π_N be indexed by $\boldsymbol{\pi}^1, \ldots, \boldsymbol{\pi}^{|\Pi_N|}$, where $|\Pi_N|$ is the cardinality of Π_N. Thus, it must be true that

$$\begin{aligned} \boldsymbol{\lambda} &= \tfrac{1}{W} \textstyle\sum_{j=1}^{W^{\text{nc}}(\mathbf{T})} \tilde{\boldsymbol{\pi}}^j, \\ &= \textstyle\sum_{i=1}^{|\Pi_N|} \frac{\sum_{j=1}^{W^{\text{nc}}(\mathbf{T})} \mathbb{1}_{\{\tilde{\boldsymbol{\pi}}^j = \boldsymbol{\pi}^i\}}}{W} \boldsymbol{\pi}^i, \\ &= \textstyle\sum_{i=1}^{|\Pi_N|} \alpha^i \boldsymbol{\pi}^i, \end{aligned} \tag{20}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function and for all i,

$$\alpha^i \triangleq \left(\textstyle\sum_{j=1}^{W^{\text{nc}}(\mathbf{T})} \mathbb{1}_{\{\tilde{\boldsymbol{\pi}}^j = \boldsymbol{\pi}^i\}} \right) / W.$$

By definition we have that $\alpha^i \geq 0, \forall i$, and since $W \geq W^{\text{nc}}(\mathbf{T})$, $\sum_i \alpha^i \leq 1$. We conclude that $\boldsymbol{\lambda} \in \mathbf{\Lambda}_{\text{sh}}^*$.

Next, suppose $\boldsymbol{\lambda} \in \text{cl}(\mathcal{R}_{\mathcal{P}}^{\text{nc}}) \setminus \mathcal{R}^{\text{nc}}$. By the definition of the closure of a set, there must exist a sequence $\{\boldsymbol{\lambda}^k\}$, with $\boldsymbol{\lambda}^k \in \mathcal{R}^{\text{nc}}$ for all k, such that $\boldsymbol{\lambda}^k \to \boldsymbol{\lambda}$ as $k \to \infty$. From (20), each $\boldsymbol{\lambda}^k$ has a RWA decomposition given by

$$\boldsymbol{\lambda}^k = \textstyle\sum_{i=1}^{|\Pi_N|} \alpha_k^i \boldsymbol{\pi}^i.$$

For each k, the vector $(\alpha_k^1, \ldots, \alpha_k^{|\Pi_N|})$ belongs to the compact set of non-negative real vectors having L^1 norm no greater than one. Using this compactness property, the Bolzano-Weierstrass Theorem [19] guarantees the existence of a vector $(\alpha^1, \ldots, \alpha^{|\Pi_N|})$ and a subsequence $\{k_j\}_{j=1}^\infty$ with

$$\alpha_{k_j}^i \to \alpha^i \text{ as } j \to \infty, \quad \text{for } i = 1, \ldots, |\Pi_N|. \tag{21}$$

To demonstrate that $\boldsymbol{\lambda} = \sum_i \alpha^i \boldsymbol{\pi}^i$, we make use of the following chain of relations. Let $\varepsilon > 0$, and let $\|\cdot\|$ be the L^1 norm operator.

$$\begin{aligned} \left\|\boldsymbol{\lambda} - \textstyle\sum_i \alpha^i \boldsymbol{\pi}^i\right\| &\leq \left\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^{k_j}\right\| + \left\|\boldsymbol{\lambda}^{k_j} - \textstyle\sum_i \alpha^i \boldsymbol{\pi}^i\right\|, \\ &= \left\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^{k_j}\right\| + \left\|\textstyle\sum_{i=1}^{|\Pi_N|} (\alpha_{k_j}^i - \alpha^i) \boldsymbol{\pi}^i\right\|, \\ &< \varepsilon. \end{aligned} \tag{22}$$

Equation (22) follows for j sufficiently large from the convergence property of the sequence $\{\boldsymbol{\lambda}^k\}$ and by (21). Finally, we have that $\alpha_k^i \geq 0, \forall k, i$, and that $\sum_i \alpha_k^i \leq 1, \forall k$, from which it must be true that the limiting quantities $\alpha^1, \ldots, \alpha^{|\Pi_N|}$ satisfy $\alpha^i \geq 0, \forall i$, and $\sum_i \alpha^i \leq 1$. This implies $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\text{sh}}$.

Proof that $\boldsymbol{\Lambda}^*_{\text{sh}} \subseteq \text{cl}(\mathcal{R}^{\text{nc}})$: Suppose $\boldsymbol{\lambda} \notin \text{cl}(\mathcal{R}^{\text{nc}})$. Then we must show that $\boldsymbol{\lambda} \notin \boldsymbol{\Lambda}^*_{\text{sh}}$. Suppose $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\text{sh}}$. Then there exist $\alpha^1, \ldots, \alpha^{|\Pi_N|}$ such that $\boldsymbol{\lambda} = \sum_i \alpha^i \boldsymbol{\pi}^i$. Define α_k^i to be the value α^i truncated to k decimal places. This truncation ensures that $\alpha_k^i \geq 0, \forall i, k$, $\sum_i \alpha_k^i \leq 1, \forall k$, and that $\alpha_k^i \to \alpha^i$ as $k \to \infty$ for $i = 1, \ldots, |\Pi_N|$. For each k, define $\boldsymbol{\lambda}^k = \sum_i \alpha_k^i \boldsymbol{\pi}^i$, $\mathbf{T}^k = 10^k \boldsymbol{\lambda}^k$, and $W_k = \sum_i 10^k \alpha_k^i$. Clearly $\mathbf{T}^k$ is an integer matrix for every k. The decomposition property of $\boldsymbol{\lambda}^k$ implies

$$\mathbf{T}^k = \textstyle\sum_{i=1}^{|\Pi_N|} 10^k \alpha_k^i \boldsymbol{\pi}^i. \tag{23}$$

Since $10^k \alpha_k^i$ is an integer for all i, k, we may interpret (23) as a valid RWA for traffic $\mathbf{T}^k$ using $W_k \leq 10^k$ wavelengths. This follows because each $\boldsymbol{\pi}^i$ can be routed on a single wavelength. By definition, it must be true that $W_k \geq W^{\text{nc}}(\mathbf{T}^k)$. Thus, $\boldsymbol{\lambda}^k = \mathbf{T}^k / W_k \in \mathcal{R}^{\text{nc}}$ for each k. Since $\boldsymbol{\lambda}^k \to \boldsymbol{\lambda}$, then $\boldsymbol{\lambda} \in \text{cl}(\mathcal{R}^{\text{nc}})$, which is a contradiction.

Appendix C
Proof of Theorem 5.1

We consider the single-hop case only, since the multi-hop case follows similarly. Denote $\theta^* = \limsup_{k \to \infty} k / W^{\text{nc}}(k\mathbf{J})$.

From the definition of θ^*, there must exist a sequence $\{k_l\}$ such that $k_l \to \infty$ as $l \to \infty$, and

$$k_l / W^{\text{nc}}(k_l \mathbf{J}) \to \theta^*. \tag{24}$$

Define the uniform arrival rate matrix $\boldsymbol{\lambda}^l = k_l \mathbf{J} / W^{\text{nc}}(k_l \mathbf{J})$. From the definition of the set $\mathcal{R}^{\text{nc}}$, we have that $\boldsymbol{\lambda}^l \in \mathcal{R}^{\text{nc}}$ for all l. Due to the convergence property (24), it must be true that $\theta^* \mathbf{J} \in \text{cl}(\mathcal{R}^{\text{nc}})$. By Theorem 4.1 we then have that $\theta^* \mathbf{J} \in \boldsymbol{\Lambda}^*_{\text{sh}}$.

Conversely, suppose that $\boldsymbol{\lambda}$ is a uniform arrival rate matrix, with uniform arrival rate $r > \theta^*$, for which $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\text{sh}}$. Theorem 4.1 provides that $\boldsymbol{\lambda} \in \text{cl}(\mathcal{R}^{\text{nc}})$. Thus, there must exist a sequence of matrices $\{\boldsymbol{\lambda}^k\}$ such that $\boldsymbol{\lambda}^k \to \boldsymbol{\lambda}$ as $k \to \infty$, and $\boldsymbol{\lambda}^k \in \mathcal{R}^{\text{nc}}$ for all k. Consequently, by the definition of the set $\mathcal{R}^{\text{nc}}$, there must exist a sequence of traffics $\{\mathbf{T}^k\}$ and integers $\{W_k\}$ such that $\boldsymbol{\lambda}^k = \mathbf{T}^k / W_k$ with $W_k \geq W^{\text{nc}}(\mathbf{T}^k)$ for all k. Define the sequence of traffics $\{\tilde{\mathbf{T}}^k\}$, with $\tilde{\mathbf{T}}^k = (\min_{i \neq j} T_{ij}^k)\mathbf{J}$. Since $\lambda_{ij}^k \to r$ for all $i \neq j$, it must be true that $(\min_{i \neq j} \lambda_{ij}^k) \to r$. This implies that $\tilde{\mathbf{T}}^k / W_k = (\min_{i \neq j} T_{ij}^k / W_k)\mathbf{J} \to r\mathbf{J}$. Clearly, since the traffic $\tilde{\mathbf{T}}^k$ is integer and fully dominated (entry-by-entry) by $\mathbf{T}^k$, it must be true that $\tilde{\mathbf{T}}^k$ can be satisfied using W_k wavelengths. This follows by using the RWA for $\mathbf{T}^k$ using W_k wavelengths in order to build a RWA for $\tilde{\mathbf{T}}^k$ using W_k wavelengths. Since $r > \theta^*$, there must exist k^* such that when $k > k^*$, for $i \neq j$, $\tilde{T}_{ij}^k / W_k > \theta^*$. Since W_k wavelengths are sufficient for a RWA with no conversion of traffic $\tilde{\mathbf{T}}^k$, we must have that $W_k \geq W^{\text{nc}}(\tilde{\mathbf{T}}^k)$. Thus for all $i \neq j$, $\tilde{T}_{ij}^k / W^{\text{nc}}(\tilde{\mathbf{T}}^k) > \theta^*$, which implies by the definition of $\tilde{\mathbf{T}}^k$ that for $k > k^*$,

$$(\min_{i \neq j} T_{ij}^k) / W^{\text{nc}}((\min_{i \neq j} T_{ij}^k)\mathbf{J}) > \theta^*. \tag{25}$$

For integer $c > 0$, the traffic $c\tilde{\mathbf{T}}^k$ can be satisfied using cW_k wavelengths, by simply repeating the RWA for traffic $\tilde{\mathbf{T}}^k$ a total of c times. Consequently, we must have $W^{\text{nc}}(c\tilde{\mathbf{T}}^k) \leq cW_k$. Combining this fact with (25), we have for any $k > k^*$, and any $c \geq 1$,

$$(c \min_{i \neq j} T_{ij}^k) / W^{\text{nc}}(c(\min_{i \neq j} T_{ij}^k)\mathbf{J}) > \theta^*.$$

This violates the definition of θ^* and provides a contradiction.

Appendix D
Proof of Theorem 5.2

In this appendix, we focus on the single-hop quantity, α^{sh}. The proof for the multi-hop quantity α^{mh} follows identically. Denote $\alpha^* = \limsup_{k \to \infty} k / \mathcal{W}^{\text{nc}}(k)$.

Definition D.1: Let the set $\partial\mathcal{D}_s$ denote the set of doubly substochastic matrices having at least one row or column sum equal to s: $\partial\mathcal{D}_s = \{\boldsymbol{\lambda} \in \mathcal{D}_s : \|\boldsymbol{\lambda}\|_{\max} = s\}$.

Proof that $\alpha^{\text{sh}} \geq \limsup_{k \to \infty} k / \mathcal{W}^{\text{nc}}(k)$: Suppose $\boldsymbol{\lambda} \in \mathcal{D}_{\alpha^*}$, with $\boldsymbol{\lambda} \neq 0$ (since $\boldsymbol{\lambda} = 0$ has a trivial RWA decomposition). Define the sequence of integer traffic matrices $\{\mathbf{T}^k\}$, such that for $i \neq j$, $T_{ij}^k = (\lfloor \lambda_{ij} \mathcal{W}^{\text{nc}}(k) - \eta_k \rfloor)^+$. Here, the operator $(\cdot)^+$ sets to zero any negative elements of its matrix operand, and $\lfloor \cdot \rfloor$ is the floor operator. We seek to ensure that $\mathbf{T}^k \in \mathcal{K}_k$, $\forall k$. To this end, consider the following series of relations. For sequence $\{\eta_k\}$, which we define subsequently, and k sufficiently large,

$$\begin{aligned}
\|\mathbf{T}^k\|_{\max} &= \left\| (\lfloor \boldsymbol{\lambda} \mathcal{W}^{\text{nc}}(k) - \eta_k \mathbf{J} \rfloor)^+ \right\|_{\max} \\
&\leq \left\| (\boldsymbol{\lambda} \mathcal{W}^{\text{nc}}(k) - \eta_k \mathbf{J})^+ \right\|_{\max} && (26) \\
&\leq \|\boldsymbol{\lambda} \mathcal{W}^{\text{nc}}(k)\|_{\max} - \eta_k && (27) \\
&\leq \alpha^* \mathcal{W}^{\text{nc}}(k) - \eta_k && (28) \\
&\leq k + \varepsilon_k \mathcal{W}^{\text{nc}}(k) - \eta_k, && (29)
\end{aligned}$$

where for $k \in \mathbb{Z}_+$,

$$\varepsilon_k = \sup_{\tilde{k} \geq k} \left| (\tilde{k} / \mathcal{W}^{\text{nc}}(\tilde{k})) - \alpha^* \right|.$$

In (26), if we assume that $\eta_k / \mathcal{W}^{\text{nc}}(k) \to 0$ as $k \to \infty$, then (27) follows for k sufficiently large, since there is at least one non-zero element on the row/column having maximum sum in $\boldsymbol{\lambda}$. Note that $\mathcal{W}^{\text{nc}}(k)$ increases at least linearly in k. Since $\boldsymbol{\lambda} \in \mathcal{D}_{\alpha^*}$, (28) must follow. By (14) we then have (29).

To ensure $\mathbf{T}^k \in \mathcal{K}_k$, we simply choose $\eta_k = \varepsilon_k \mathcal{W}^{\text{nc}}(k)$. Clearly, $\eta_k / \mathcal{W}^{\text{nc}}(k) \to 0$ as $k \to \infty$, since (14) implies that $\varepsilon_k \to 0$ as $k \to \infty$. Next, define $\boldsymbol{\lambda}^k = (1/\mathcal{W}^{\text{nc}}(k))\mathbf{T}^k$. Since $\mathbf{T}^k \in \mathcal{K}_k$, it must be true that $\boldsymbol{\lambda}^k \in \mathcal{R}^{\text{nc}}$. To demonstrate that $\boldsymbol{\lambda}$ has a RWA decomposition, we need to show that $\boldsymbol{\lambda}^k \to \boldsymbol{\lambda}$ as $k \to \infty$. Since $\eta_k / (\mathcal{W}^{\text{nc}}(k)) \to 0$ as $k \to \infty$, this is clearly true. Thus, $\boldsymbol{\lambda} \in \text{cl}(\mathcal{R}^{\text{nc}})$, which implies by Theorem 4.1 that $\boldsymbol{\lambda} \in \boldsymbol{\Lambda}^*_{\text{sh}}$. Since this holds for all $\boldsymbol{\lambda} \in \mathcal{D}_{\alpha^*}$, it must be true that $\alpha^{\text{sh}} \geq \alpha^*$.

Proof that $\alpha^{\text{sh}} \leq \limsup_{k \to \infty} k / \mathcal{W}^{\text{nc}}(k)$: Suppose there exists $\alpha > \alpha^*$ such that $\mathcal{D}_\alpha \subseteq \boldsymbol{\Lambda}^*_{\text{sh}}$. Consider any positive integer u. Let $\boldsymbol{\lambda}^{u,1}, \ldots, \boldsymbol{\lambda}^{u,K_u}$ be a finite set of matrices belonging to $\partial\mathcal{D}_\alpha$, such that

$$\partial\mathcal{D}_\alpha \subseteq \bigcup_{l=1}^{K_u} \left\{ \boldsymbol{\lambda} : |\lambda_{ij} - \lambda_{ij}^{u,l}| \leq 1/u, \forall i, j \in V \right\}.$$

In words, the set of points $\{\boldsymbol{\lambda}^{u,1},\ldots,\boldsymbol{\lambda}^{u,K_u}\}$ are the center locations of a set of $(1/u)$-balls that cover the outer boundary $\partial\mathcal{D}_\alpha$. The compactness of $\mathcal{D}_\alpha$ is sufficient to ensure the existence of a covering such that K_u is finite-valued [19].

Since $\boldsymbol{\lambda}^{u,l}\in\mathcal{D}_\alpha$, and by our assumption that $\mathcal{D}_\alpha\subseteq\boldsymbol{\Lambda}^*_{\mathrm{sh}}$, Theorem 4.1 provides that there must exist a set of integer traffics $\{\mathbf{T}^{u,1},\ldots,\mathbf{T}^{u,K_u}\}$, and a set of positive integers $\{W^{u,1},\ldots,W^{u,K_u}\}$ such that for $l=1,\ldots,K_u$,

$$(1/W^{u,l})\mathbf{T}^{u,l}\in\left\{\boldsymbol{\lambda}:|\lambda_{ij}-\lambda_{ij}^{u,l}|\le 1/u,\ \forall i,j\in V\right\}, \quad (30)$$

where $W^{u,l}\ge W^{\mathrm{nc}}(\mathbf{T}^{u,l})$ for all l. Since K_u is finite and $\mathbf{T}^{u,l}$ is an integer matrix for all l, there must exist integers $\kappa_1^u,\ldots,\kappa_{K_u}^u$ and k_u^*, such that for $l=1,\ldots,K_u$,

$$\kappa_l^u\|\mathbf{T}^{u,l}\|_{\max}=k_u^*.$$

The integer traffic $\kappa_l^u\mathbf{T}^{u,l}$ must have a RWA using $\kappa_l^uW^{u,l}$ wavelengths. This RWA is constructed by repeating the RWA for traffic $\mathbf{T}^{u,l}$, that makes use of $W^{u,l}$ wavelengths, a total of κ_l^u times over $\kappa_l^uW^{u,l}$ wavelengths. While the maximum row/column sum of $\boldsymbol{\lambda}^{u,l}$ is α, that of $(\kappa_l^u/W^{u,l})\mathbf{T}^{u,l}$ is $k_u^*/W^{u,l}$ for each l. Applying (30), we then have for $l=1,\ldots,K_u$,

$$\left|\alpha-k_u^*/(\kappa_l^uW^{u,l})\right|\le(n-1)/u. \quad (31)$$

Consider any traffic $\mathbf{T}\in\mathcal{K}_{k_u^*}$, with maximum row/column sum equal to k_u^*. Then $(\alpha/k_u^*)\mathbf{T}\in\partial\mathcal{D}_\alpha$, which implies there exists l^* such that for all $i,j\in V$,

$$\left|(\alpha/k_u^*)T_{ij}-\lambda_{ij}^{u,l^*}\right|\le 1/u. \quad (32)$$

Combining (30) with (32), we have

$$\left|(\alpha/k_u^*)T_{ij}-T_{ij}^{u,l^*}/W^{u,l^*}\right|\le 2/u. \quad (33)$$

Multiplying (33) through by $\kappa_{l^*}^uW^{u,l^*}$ provides

$$\left|(\alpha\kappa_{l^*}^uW^{u,l^*}/k_u^*)T_{ij}-\kappa_{l^*}^uT_{ij}^{u,l^*}\right|\le(2/u)\kappa_{l^*}^uW^{u,l^*}. \quad (34)$$

Note that if $a>0$, and $|ax-y|\le c$, then if $ax-y>0$, we have $x-y\le c/a+((1-a)/a)y$, and if $ax-y<0$, we have $x-y\le((1-a)/a)y$. Consequently, equation (34) implies

$$\begin{aligned}\left|T_{ij}-\kappa_{l^*}^uT_{ij}^{u,l^*}\right|&\le\frac{2}{u}\kappa_{l^*}^uW^{u,l^*}\frac{k_u^*}{\alpha\kappa_{l^*}^uW^{u,l^*}}\\&\quad+\kappa_{l^*}^uT_{ij}^{u,l^*}\left(k_u^*/(\alpha\kappa_{l^*}^uW^{u,l^*})-1\right).\end{aligned}$$

The difference between the integer traffic demand matrix $\mathbf{T}$ and the matrix $\kappa_{l^*}^u\mathbf{T}^{u,l^*}$ can then be bounded as

$$\begin{aligned}\textstyle\sum_{i,j}\left|T_{ij}-\kappa_{l^*}^uT_{ij}^{u,l^*}\right|&\le n(n-1)\frac{2}{u}\kappa_{l^*}^uW^{u,l^*}\frac{k_u^*}{\alpha\kappa_{l^*}^uW^{u,l^*}}\\&\quad+nk_u^*\left(k_u^*/(\alpha\kappa_{l^*}^uW^{u,l^*})-1\right)\\&\triangleq\omega_{u,l^*}.\end{aligned}$$

Then,

$$\frac{\omega_{u,l^*}}{k_u^*}=n(n-1)\frac{2}{\alpha u}+n\left(\frac{k_u^*}{\alpha\kappa_{l^*}^uW^{u,l^*}}-1\right).$$

Applying (31), it is clear that $\omega_{u,l^*}/k_u^*\to 0$ as $u\to\infty$.

If each additional integer demand in traffic $\mathbf{T}$ over that in traffic $\kappa_{l^*}^u\mathbf{T}^{u,l^*}$ is serviced using a unique wavelength, the value of ω_{u,l^*} can be used to infer an upper bound on the minimum number of wavelengths required to service $\mathbf{T}$. This holds, given the appropriate choice for the index l^*, for any $\mathbf{T}\in\mathcal{K}_{k_u^*}$ having maximum row/column sum of k_u^*, from which we obtain $\mathcal{W}^{\mathrm{nc}}(k_u^*)\le\max_l(\kappa_l^uW^{u,l}+\omega_{u,l})$. Thus,

$$k_u^*/\mathcal{W}^{\mathrm{nc}}(k_u^*)\ge k_u^*/(\max_l\kappa_l^uW^{u,l}+\max_l\omega_{u,l}). \quad (35)$$

Applying (31), the right side of (35) must converge to α as $u\to\infty$. Thus, there must exist $\bar{u}$ such that for all $u\ge\bar{u}$,

$$k_u^*/\mathcal{W}^{\mathrm{nc}}(k_u^*)\ge(\alpha+\alpha^*)/2>\alpha^*.$$

Clearly, if $k_u^*\to\infty$, this is in violation of (14), which provides a contradiction. Thus, it remains to show that $k_u^*\to\infty$ as $u\to\infty$. Suppose this is not true, and there exists integer $\tilde{k}^*$ such that $k_u^*\le\tilde{k}^*$ for all u. We can then bound the cardinality of $\mathcal{K}_{\tilde{k}^*}$ as $|\mathcal{K}_{\tilde{k}^*}|\le(\tilde{k}^*)^{n(n-1)}$. The number of distinct $(1/u)$-balls required to cover $\partial\mathcal{D}_\alpha$ must increase with u. This can be seen as follows: consider any two neighboring (sharing the same face) non-zero vertices of $\mathcal{D}_\alpha$. The line segment joining these two vertices is completely contained in $\partial\mathcal{D}_\alpha$. This line segment is isomorphic to an interval of equal length on the real line, for which a covering by $(1/u)$-balls clearly requires an increasing number of balls as u increases. Furthermore, since the line segment is not collinear with the origin (this would violate that one of the end points is a vertex of $\mathcal{D}_\alpha$), the number of covering $(1/u)$-balls that exist such that no two balls contain *any* matrices that are scaled versions of one another, is also increasing with u. Consequently, for sufficiently large u, there must be more than $(\tilde{k}^*)^{n(n-1)}$ traffics in the set $\{\mathbf{T}^{u,1},\ldots,\mathbf{T}^{u,K_u}\}$ that are not scaled versions of one another. Since the line joining each of these traffics to the origin has a unique direction, the common boundary that these traffics will be scaled to (using the integers from the set $\{\kappa_l^u\}$) must contain more than $(\tilde{k}^*)^{n(n-1)}$ integer matrices. This however, is in violation of our assumption that $|\mathcal{K}_{k_u^*}|\le(\tilde{k}^*)^{n(n-1)}$ for all u. Thus, $k_u^*\to\infty$ as $u\to\infty$.

References

[1] M. Ajmone Marsan, E. Leonardi, M. Mellia, and F. Neri, "On the stability of isolated and interconnected input-queueing switches under multiclass traffic," *IEEE Trans. Inform. Theory*, vol. 45, no. 3, pp. 1167–1174, Mar. 2005.

[2] G. Birkhoff, "Tres observaciones sobre el algebra lineal," *Rev. Universidad Nacional de Tucuman*, pp. 147–150, 1946.

[3] A. Brzezinski, "Scheduling algorithms for throughput maximization in data networks," Ph.D. dissertation, M.I.T., 2007.

[4] A. Brzezinski and E. Modiano, "Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic," in *IEEE Proc. INFOCOM'05*, Miami, FL, Mar. 2005.

[5] ——, "Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic," *IEEE/OSA J. Lightw. Technol.*, vol. 23, no. 10, pp. 3188–3205, Oct. 2005.

[6] ——, "Achieving 100% throughput in reconfigurable optical networks," in *Proc. IEEE INFOCOM High-Speed Networking Workshop*, Apr. 2006.

[7] ——, "RWA decompositions for optimal throughput in reconfigurable optical networks," in *Proc. INFORMS Telecom*, Mar. 2006.

[8] C. S. Chang, W. J. Chen, and H. Y. Huang, "Birkhoff-von Neumann input buffered crossbar switches," in *IEEE Proc. Information Communications (INFOCOM)*, 2000, pp. 1614–1623.

[9] L. Chen and E. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks with wavelength converters," in *Proc. IEEE INFOCOM'03*, Apr. 2003.

[10] ——, "Dynamic routing and wavelength assignment with optical bypass using ring embeddings," *Opt. Switch. Netw.*, vol. 1, no. 1, pp. 35–49, Jan. 2005.

[11] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high bandwidth optical WAN's," *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–1182, 1992.

[12] J. S. Choi, N. Golmie, F. Lapeyrere, F. Mouveaux, and D. Su, "A functional classification of routing and wavelength assignment schemes in DWDM networks: static case," in *Proc. OPNET 2000*, Jan. 2000, pp. 1109–1115.

[13] T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms*. The MIT Press, 2001.

[14] A. Fumagalli and L. Valcarenghi, "IP restoration vs. WDM protection: is there an optimal choice?" *IEEE Network*, vol. 14, no. 6, pp. 34–41, Nov-Dec 2000.

[15] O. Gerstel, G. Sasaki, S. Kutten, and R. Ramaswami, "Worst-case analysis of dynamic wavelength allocation in optical networks," *IEEE/ACM Trans. Netw.*, vol. 7, no. 6, pp. 833–845, Dec. 1999.

[16] N. Ghani, S. Dixit, and T. Wang, "On IP-over-WDM integration," *IEEE Commun. Mag.*, pp. 72–84, Mar. 2000.

[17] K. Kar, D. Stiliadis, T. Lakshman, and L. Tassiulas, "Scheduling algorithms for optical packet fabrics," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 7, pp. 1143–1155, Sep. 2003.

[18] R. Madan and D. Shah, "Capacity-delay scaling in arbitrary wireless networks," in *Proceedings Allerton Conf. Communication, Control, and Computing*, Sep. 2005.

[19] J. E. Marsden and M. J. Hoffman, *Elementary Classical Analysis*, 2nd ed. W. H. Freeman and Company, 2000.

[20] N. McKeown, A. Mekkittikul, V. Anantharam, and J. Walrand, "Achieving 100% throughput in an input-queued switch," *IEEE Trans. Commun.*, vol. 47, no. 8, pp. 1260–1267, Aug. 1999.

[21] B. Mukherjee, *Optical Communication Networks*. McGraw-Hill, 1997.

[22] A. Narula-Tam, P. J. Lin, and E. Modiano, "Efficient routing and wavelength assignment for reconfigurable WDM networks," *IEEE J. Select. Areas Commun.*, vol. 20, no. 1, pp. 75–88, Jan. 2002.

[23] C. Qiao, "Labeled optical burst switching for IP-over-WDM integration," *IEEE Commun. Mag.*, pp. 104–114, Sep. 2000.

[24] K. Ross, N. Bambos, K. Kumaran, I. Saniee, and I. Widjaja, "Dynamic scheduling of optical data bursts in time-domain wavelength interleaved networks," in *High Performance Interconnects*, Aug. 2003.

[25] ——, "Scheduling bursts in time-domain wavelength interleaved networks," *IEEE J. Select. Areas Commun.*, vol. 21, no. 9, pp. 1441–1451, Nov. 2003.

[26] P. Saengudomlert, "Architectural study of high-speed networks with optical bypassing," Ph.D. dissertation, Massachusetts Institute of Technology, Sep. 2002.

[27] P. Saengudomlert, E. Modiano, and R. Gallager, "Dynamic wavelength assignment for WDM all-optical tree networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 4, pp. 895–905, Aug. 2005.

[28] ——, "On-line routing and wavelength assignment for dynamic traffic in WDM ring and torus networks," *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 330–340, Apr. 2006.

[29] L. Sahasrabuddhe, S. Ramamurthy, and B. Mukherjee, "Fault management in IP-over-WDM networks: WDM protection versus IP restoration," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 21–33, Jan. 2002.

[30] A. Stolyar, "Maxweight scheduling in a generalized switch: state space collapse and workload minimization in heavy traffic," *Ann. Appl. Probab.*, vol. 14, no. 1, pp. 1–53, Jan. 2004.

[31] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Automat. Contr.*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.

[32] J. von Neumann, "A certain zero-sum two-person game equivalent to the optimal assignment problem," *Contributions to the Theory of Games, vol. 2*, pp. 5–12, 1953.

[33] J. Y. Wei, "Advances in the management and control of optical Internet," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 4, pp. 768–785, May 2002.

[34] G. Weichenberg, V. W. S. Chan, and M. Medard, "On the capacity of optical networks: A framework for comparing different transport architectures," Jul. 2005, Technical Report LIDS-P-2655, MIT LIDS.

[35] ——, "On the capacity of optical networks: A framework for comparing different transport architectures," in *IEEE Proc. Information Communications (INFOCOM)*, Apr. 2006.

[36] I. Widjaja, I. Saniee, R. Giles, and D. Mitra, "Light core and intelligent edge for a flexible, thin-layered and cost-effective optical transport network," *IEEE Commun. Mag.*, vol. 41, pp. S30–S36, May 2003.

[37] C. Xin and C. Qiao, "A comparative study of OBS and OFS," in *IEEE/OSA Optical Fiber Conference (OFC)*, 2001, pp. ThG7–1–ThG7–3.

[38] Y. Yinghua, C. Assi, S. Dixit, and M. A. Ali, "A simple dynamic integrated provisioning/protection scheme in IP over WDM networks," *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 174–182, Nov. 2001.

[39] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-over-WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 2062–2071, Oct. 2000.

Andrew Brzezinski received the B.A.Sc. degree in electrical engineering from the University of Toronto, Canada, in 2000, the M.S. degree in electrical engineering from Stanford University in 2002, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT). He is currently with Fidelity Investments, Boston, MA.

Eytan Modiano received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and PhD degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992 respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992-1993 while he was conducting research on security and performance issues in distributed network protocols.

Between 1993 and 1999 he was with the Communications Division at MIT Lincoln Laboratory where he designed communication protocols for satellite, wireless, and optical networks and was the project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. He joined the MIT faculty in 1999, where he is presently an Associate Professor. His research is on communication networks and protocols with application to satellite, wireless, and optical networks.

He is currently an Associate Editor for Communication Networks for *IEEE Transactions on Information Theory* and for the *International Journal of Satellite Communications*. He had served as a guest editor for the *IEEE Journal on Selected Areas in Communications* special issue on WDM network architectures; the *Computer Networks Journal* special issue on Broadband Internet Access; the *Journal of Communications and Networks* special issue on Wireless Ad-Hoc Networks; and for *IEEE Journal of Lightwave Technology* special issue on Optical Networks. He was the Technical Program co-chair for WiOpt 2006, IEEE Infocom 2007, and ACM MobiHoc 2007.