Achieving 100% Throughput in Reconfigurable Optical Networks

Andrew Brzezinski and Eytan Modiano
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
brzezin@mit.edu, modiano@mit.edu

***Abstract*— We study the maximum throughput properties of dynamically reconfigurable optical networks 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. We describe throughput-optimal dynamic algorithms employing joint WDM reconfiguration and electronic layer routing decisions. Our approach 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.**

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. Our architecture consists of 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 [1]. 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. 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 Ephrimedes underlies much of the existing literature in the area of stability of communication networks [2]. Indeed, the network model considered in this paper easily fits into the framework of Tassiulas and Ephrimedes, as does much of the switch scheduling literature. A major accomplishment of our work is a characterization of the capacity 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 maximum stability region for optical networks of arbitrary topologies. In this work, we 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 [3]. 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 [2] implies that the convex hull of these service configurations equals the maximum 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) [3]. 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 [2] allows us to express the capacity region as the convex hull of available service configurations, this description can have

This work was supported in part by DARPA under grant MDA972-02-1-0021 and by NSF under grants ANI-0073730 and ANI-0335217.

limited value in providing an understanding of the geometric properties of the capacity region. This is in contrast to the case of the input-queued switch, where a result of Birkhoff has been applied to demonstrate that the convex hull of the service matrices (permutation matrices) equals the doubly substochastic region [4].

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 (see for example [5]–[7]). Several studies consider Optical Burst Switching (OBS) as the mechanism for accessing the optical transport layer [6], [8], [9]. Most solutions seek to integrate IP and Generalized Multiprotocol Label Switching (GMPLS) functionality. Finally, we note the work of [10], where the capacity properties of burst, flow and packet switching architectures were compared. 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.

II. Network properties

We consider a WDM-based packet network with N nodes, labeled $1,\ldots,N$, with node i having P_i transceivers. The physical topology $\mathcal{P}$ consists of L fiber links, labeled $1,\ldots,L$, with each link having a single wavelength available for transmission of data. The network operates in slotted time, with n used to represent the time index. Packets are assumed to have fixed size, with transmission duration of one slot.

Each network node employs virtual output queueing of packets. We designate $\text{VOQ}_{i,j}$ as the virtual output queue containing packets at node i awaiting transmission to node j. Let $X_{i,j}(n)$ be the number of packets enqueued in $\text{VOQ}_{i,j}$ at time n. The arrival process is modeled as a stochastic process, with $A_{i,j}(n)$ representing the cumulative number of arrivals to $\text{VOQ}_{i,j}$ up to and including time n. We assume the arrival processes have a strong law of large numbers property, with long-term *arrival rates* given by

$$\lim_{n\to\infty}\frac{A_{i,j}(n)}{n}=\lambda_{i,j}\ \text{w.p.1}\ \forall i,j.$$

Above, the term w.p.1 represents the statement 'with probability 1'. For convenience, we gather the arrival rates into the $N\times N$ matrix λ. We define $U_{i,j}(n)$ as the cumulative service applied to $\text{VOQ}_{i,j}$ up to and including time slot n. Note that this cumulative variable accounts for the packets that exit the system as well as those that are sent to intermediate nodes of a multi-hop path toward their destination. We gather the arrival and departure processes at each time n into the respective matrices $A(n),U(n)$. Finally, we designate the differential variable $u(n)=(u_{i,j}(n),i,j=1,\ldots,n)$ as the *service configuration* matrix applied at time n. We can then express

$$X(n+1)=X(0)+A(n+1)-U(n).$$

The $N\times N$ non-negative integer matrix $v(n)$ is used to represent the logical topology enabled over the n-th time slot, where $v_{i,j}(n)$ is the number of wavelengths configured for transmission between nodes i and j at time n. This is not to be confused with the service configuration matrix $u(n)$, which captures the actual packet flows over the logical topology $v(n)$ at time n. We define $\mathcal{V}_{\mathcal{P}}$ as the set of all logical configurations that can be enabled over the physical topology $\mathcal{P}$.

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*. To be precise, rate stability is given as follows.

Definition 2.1: A system of queues is *rate stable* if

$$\lim_{n\to\infty}\frac{X_{i,j}(n)}{n}=\lim_{n\to\infty}\frac{A_{i,j}(n)-U_{i,j}(n)}{n}=0\ \text{ w.p.1, }\forall i,j.$$

The optimal throughput performance of the reconfigurable optical network is characterized through the maximum stability region or *capacity 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 and algorithms exclusively employing single-hop electronic-layer routes, we distinguish two capacity regions: one for achievable rates under multi-hop electronic-layer routing, and one for achievable rates under exclusively single-hop electronic-layer routing. The *single-hop capacity region*, denoted $\Lambda^{\text{sh}}_{\mathcal{P}}$, is a set of arrival rate matrices that can be stabilized in the network with physical topology $\mathcal{P}$. Any process having an arrival rate belonging to the capacity region can be stabilized by some reconfiguration and routing algorithm employing exclusively single-hop paths. Similarly, the *multi-hop capacity region*, denoted $\Lambda^{\text{mh}}_{\mathcal{P}}$, consists of all those rates such that any arrival process having a rate belonging to the region can be stabilized by some joint reconfiguration and routing algorithm (possibly employing multi-hop routes). An algorithm is called *throughput optimal (achieves 100% throughput)* if the set of arrival rates that it can stabilize equals the multi-hop capacity region.

In [2], Tassiulas and Ephrimedes provided a characterization of the capacity properties of a general multi-hop-capable network. Their algorithmic description for scheduling in this network setting involves *maxweight decisions*, where each network configuration has associated with it a particular weight, and the maximum weighted configuration is chosen at each time. Here, we introduce two versions of this algorithm, specialized to general reconfigurable WDM-based networks.

Algorithm **SHMW** employs WDM reconfiguration and single-hop electronic layer routing. In other words, if a directed logical link exists connecting node i to node j at time n, then that link can only be used at time n to service packets at node i that are destined for node j. For equally-sized matrices ν,X, let their *inner product* be given by

SHMW Single-hop maxweight scheduling algorithm

At time n, select logical topology

$$v(n) = \arg\max_{\nu \in \mathcal{V}_{\mathcal{P}}} \langle \nu, X(n) \rangle.$$

For this WDM configuration, packet scheduling is performed by single-hop routing, which implies

$$u_{i,j}(n) = \min\{X_{i,j}(n), v_{i,j}(n)\}. \quad (1)$$

Equation (1) implies that up to $v_{i,j}(n)$ logical links are available for servicing packets enqueued at node i for destination j, given sufficient backlog in queue $X_{i,j}$.

$\langle \nu, X \rangle = \sum_{i,j} \nu_{i,j} X_{i,j}$. At time n, the algorithm selects a logical topology matrix from $\mathcal{V}_{\mathcal{P}}$ whose inner product with the queue backlog matrix $X(n)$ is maximum. This logical topology is used for single-hop routing of packets to their destinations. Algorithm **SHMW** is detailed above. **SHMW** is a generalized version of the maximum weighted matching (MWM) algorithm for achieving 100% throughput in input-queued switches [4], and in optical networks with no wavelength/RWA constraints [11]. It can be shown that the stability region achieved by **SHMW** is $\Lambda^{\text{sh}}_{\mathcal{P}}$. For context, note that in the case of an input-queued switch, algorithm **SHMW** achieves 100% throughput, with capacity region equal to the *doubly substochastic region*:

$$\Lambda^{\text{sh}}_{\text{IQ}} = \left\{ \lambda : \sum_j \lambda_{i,j} \le 1, \forall i, \ \sum_i \lambda_{i,j} \le 1, \forall j \right\}. \quad (2)$$

The second algorithm of interest, algorithm **MHMW**, employs WDM reconfiguration and multi-hop electronic layer routing. **MHMW** is outlined below. We refer to each packet destined for a particular destination as a unit of a *commodity* that is unique to that destination. The *differential backlog* of commodity k at each link consists of the difference between the number of commodity k packets enqueued at the source node of the link and the number of commodity k packets equeued at the destination node of the link. In words, algorithm **MHMW** can be described for time $n \ge 0$ as follows. For each link $i \to j$, the maximum differential backlog over all commodities is calculated and stored as the (i,j) entry of the $N \times N$ matrix $Z(n)$. The logical topology $v(n)$ is chosen from the set $\mathcal{V}_{\mathcal{P}}$ to maximize the inner product $\langle v(n), Z(n) \rangle$. For each link of the selected logical topology, a commodity that maximizes the differential backlog for that link is selected for electronic routing across the link.

The stability region of **MHMW** *under the rate stability criterion* can be expressed as the *closed* convex hull of the available multi-hop service configurations. A result of [2] is that **MHMW** achieves the capacity region $\Lambda^{\text{mh}}_{\mathcal{P}}$.

B. The RWA problem

The routing and wavelength assignment (RWA) problem takes as input a physical topology $\mathcal{P}$ and the integer traffic

MHMW Multi-hop maxweight scheduling algorithm

At time n, calculate for each source-destination pair $l = i \to j$ and commodity k the ***differential backlog***, $d_{l,k}(n)$:

$$d_{l,k}(n) = X_{i,k}(n) - X_{j,k}(n).$$

Define the $N \times N$ matrix of maximum differential backlogs, $Z(n)$, by selecting for each link l a commodity k that maximizes the differential backlog across that link:

$$Z_{i,j}(n) \triangleq \max_k \{d_{l,k}(n)\}.$$

Select the logical topology $v(n) \in \mathcal{V}_{\mathcal{P}}$ to maximize the inner product $\langle v(n), Z(n) \rangle$. Electronic routing on $v(n)$ is performed by transmitting along each active logical link a packet whose destination maximizes the differential backlog along that link. Thus, if logical link $l = i \to j$ is active, we transmit along l a packet destined for node $k^*_l \in \arg\max_k d_{l,k}(n)$.

matrix T, corresponding to wavelength demands that are to be fully satisfied by a static lightpath configuration on $\mathcal{P}$. The output of the RWA problem is an integer W, which is the minimum number of wavelengths required to service T on physical topology $\mathcal{P}$. We consider two versions of the RWA problem: RWA with no wavelength conversion capability and RWA with full wavelength conversion capability.

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, for the particular physical topology $\mathcal{P}$, let $W^{\text{nc}}_{\mathcal{P}}(T)$ be the minimum number of wavelengths required to service traffic T with no wavelength conversion.

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 $W^{\text{c}}_{\mathcal{P}}(T)$ to be the minimum number of wavelengths required to service traffic T with wavelength conversion on physical topology $\mathcal{P}$. Since using a single color per lightpath is accommodated by the RWA with wavelength conversion, it is clear that for any physical topology $\mathcal{P}$, $W^{\text{c}}_{\mathcal{P}}(T) \le W^{\text{nc}}_{\mathcal{P}}(T)$ for all T. For the trivial case of $T = 0$ (the zero matrix), we define (for technical reasons) that $W^{\text{nc}}_{\mathcal{P}}(0) = W^{\text{c}}_{\mathcal{P}}(0) = 1$.

III. Capacity regions from RWA decompositions

In this section, we will demonstrate that the single-hop and multi-hop capacity regions for single-wavelength optical networks can be fully described by the RWA functions $W^{\text{nc}}_{\mathcal{P}}$ and $W^{\text{c}}_{\mathcal{P}}$, respectively. 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 link the RWA problem to the stability issue in our reconfigurable network. Consequently, this enables us to draw on the rich RWA literature to understand the throughput properties of particular physical topologies.

We begin by considering the single-hop capacity region. For integer traffic matrix T, the RWA function provides the minimum number of wavelengths required to satisfy the wavelength demands of T: $W_{\mathcal{P}}^{\mathrm{nc}}(T)$ in the case of no conversion. An *RWA decomposition* of T consists of $W_{\mathcal{P}}^{\mathrm{nc}}(T)$ $N \times N$ matrices (in the case of no conversion), with each matrix associated with a different wavelength of the RWA for T. Thus the i-th matrix contains the number of lightpaths established between each source-destination pair in the RWA for T on the i-th wavelength. The concept of a RWA decomposition provides key intuition for the results of this section. Additional details can be found in [12].

For each non-negative integer traffic matrix T, consider the sequence of arrival rates created by dividing T by all integers $W \geq W_{\mathcal{P}}^{\mathrm{nc}}(T)$. In this section we consider all such arrival rates, gathered over all possible integer traffics T in the RWA problem. Let $\mathcal{R}_{\mathcal{P}}^{\mathrm{nc}}$ be the set of all such arrival rates,

$$\mathcal{R}_{\mathcal{P}}^{\mathrm{nc}} = \left\{ \lambda = \frac{1}{W} T : T \in \mathbb{Z}_+^M, W \in \mathbb{Z}_+, W \geq W_{\mathcal{P}}^{\mathrm{nc}}(T) \right\},$$

where $\mathbb{Z}_+$ is the set of non-negative integers, $\mathbb{Z}_+^M$ is the M-fold Cartesian product of $\mathbb{Z}_+$, and $M = N(N-1)$. Recall also 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, $\Lambda_{\mathcal{P}}^{\mathrm{sh}}$ is the single-hop capacity region on $\mathcal{P}$, subjected to the single-wavelength constraint. For the set $\mathcal{R}$, let $\mathrm{cl}(\mathcal{R})$ represent the closure of $\mathcal{R}$. The following theorem is proved in [12].

Theorem 3.1: For physical topology $\mathcal{P}$, $\Lambda_{\mathcal{P}}^{\mathrm{sh}} = \mathrm{cl}(\mathcal{R}_{\mathcal{P}}^{\mathrm{nc}})$.

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

$$\mathcal{R}_{\mathcal{P}}^{\mathrm{c}} = \left\{ \lambda = \frac{1}{W} T : T \in \mathbb{Z}_+^M, W \in \mathbb{Z}_+, W \geq W_{\mathcal{P}}^{\mathrm{c}}(T) \right\}.$$

Through similar steps as in the single-hop case, we can establish the following theorem (proved in [12]).

Theorem 3.2: For physical topology $\mathcal{P}$, $\Lambda_{\mathcal{P}}^{\mathrm{mh}} = \mathrm{cl}(\mathcal{R}_{\mathcal{P}}^{\mathrm{c}})$.

IV. Geometric properties of the stability region

While the stability properties of our dynamically reconfigurable electronic-over-optical network are well characterized in the single-hop and multi-hop cases through Theorems 3.1 and 3.2, 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 (2), which provides the wavelength-unconstrained capacity region as the set of doubly substochastic matrices.

The remainder of this paper is dedicated to extracting geometric properties of the single-hop and multi-hop capacity regions in our wavelength-constrained reconfigurable network setting. Unlike the wavelength-unconstrained case, the physical topology $\mathcal{P}$ has an effect on the stability properties.

A. Maximum uniform (all-to-all) arrival rate matrices

In this section, we make use of RWA decompositions to establish geometric properties of the single-hop and multi-hop capacity regions. Define J as the $N \times N$ matrix having (i,j) entry equal to 1 if $i \neq j$ and equal to 0 otherwise. We then seek to determine the maximum values $\theta_{\mathcal{P}}^{\mathrm{sh}}, \theta_{\mathcal{P}}^{\mathrm{mh}}$ such that $\theta_{\mathcal{P}}^{\mathrm{sh}} J$ belongs to the single-hop capacity region, and $\theta_{\mathcal{P}}^{\mathrm{mh}} J$ belongs to the multi-hop capacity region.

Theorem 4.1: For physical topology $\mathcal{P}$, the maximum values $\theta_{\mathcal{P}}^{\mathrm{sh}}, \theta_{\mathcal{P}}^{\mathrm{mh}}$ such that $\theta_{\mathcal{P}}^{\mathrm{sh}} J \in \Lambda_{\mathcal{P}}^{\mathrm{sh}}$ and $\theta_{\mathcal{P}}^{\mathrm{mh}} J \in \Lambda_{\mathcal{P}}^{\mathrm{mh}}$, respectively, are given by

$$\theta_{\mathcal{P}}^{\mathrm{sh}} = \sup_{l \in \mathbb{Z}_+} \frac{l}{W_{\mathcal{P}}^{\mathrm{nc}}(lJ)}, \qquad \theta_{\mathcal{P}}^{\mathrm{mh}} = \sup_{l \in \mathbb{Z}_+} \frac{l}{W_{\mathcal{P}}^{\mathrm{c}}(lJ)}. \tag{3}$$

The equations in (3) 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 topology $\mathcal{P}$, with or without multi-hop capability. Theorem 4.1 is proved in [12].

Theorem 4.1 allows us to draw on the literature regarding all-to-all RWA properties for various physical topologies to obtain geometric properties of the single-hop and multi-hop stability regions [13]–[16].

B. Maximum scaled doubly substochastic region

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 capacity region. For a mathematical description of this property we require the following definitions.

Definition 4.1: For matrix A, let the maximum row/column sum of A be given by $\|A\|_{\max}$:

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

Definition 4.2: Let the set $\mathcal{D}_s$ denote the doubly substochastic region, scaled by factor s,

$$\mathcal{D}_s = \left\{ \lambda \in \mathbb{R}^M : \lambda_{i,j} \geq 0 \,\forall i,j,\ \|\lambda\|_{\max} \leq s \right\}.$$

We seek the maximum values $\alpha_{\mathcal{P}}^{\mathrm{sh}}, \alpha_{\mathcal{P}}^{\mathrm{mh}}$ such that the regions $\mathcal{D}_{\alpha_{\mathcal{P}}^{\mathrm{sh}}}, \mathcal{D}_{\alpha_{\mathcal{P}}^{\mathrm{mh}}}$ are respectively subsets of the single-hop and multi-hop capacity regions.

Definition 4.3: The integer matrix $T \in \mathbb{Z}^M$ is called k-allowable if it satisfies

$$T_{i,j} \in \mathbb{Z}_+,\ \forall i,j,\ \|T\|_{\max} \leq k.$$

We denote by $\mathcal{K}_k$ the set of all k-allowable matrices.

Let the minimum number of wavelengths required to service *any* k-allowable traffic in the RWA with no conversion be $\mathcal{W}_{\mathcal{P}}^{\mathrm{nc}}(k) = \max_{T \in \mathcal{K}_k} W_{\mathcal{P}}^{\mathrm{nc}}(T)$. Similarly, let the corresponding value with wavelength conversion be $\mathcal{W}_{\mathcal{P}}^{\mathrm{c}}(k)$.

The following theorem (proved in [12]) establishes for each physical topology $\mathcal{P}$ the quantity $\alpha_{\mathcal{P}}^{\text{sh}}$ as a scale factor on the substochastic region, such that the scaled region is contained within the single-hop capacity region. The analogous result for the multi-hop case is also provided.

Theorem 4.2: $\mathcal{D}_{\alpha_{\mathcal{P}}^{\text{sh}}} \subseteq \Lambda_{\mathcal{P}}^{\text{sh}}$ and $\mathcal{D}_{\alpha_{\mathcal{P}}^{\text{mh}}} \subseteq \Lambda_{\mathcal{P}}^{\text{mh}}$, where

$$\alpha_{\mathcal{P}}^{\text{sh}} = \limsup_{k\to\infty} \frac{k}{\mathcal{W}_{\mathcal{P}}^{\text{nc}}(k)}, \qquad \alpha_{\mathcal{P}}^{\text{mh}} = \limsup_{k\to\infty} \frac{k}{\mathcal{W}_{\mathcal{P}}^{\text{c}}(k)}. \quad (4)$$

The equations in (4) 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 topology $\mathcal{P}$, in the limit of large k.

The following theorem (proved in [12]) is the converse to Theorem 4.2, by demonstrating for each physical topology $\mathcal{P}$, that the quantity $\alpha_{\mathcal{P}}^{\text{sh}}$ is an upper bound on the maximum scale factor on the substochastic region, such that the scaled region is contained within the single-hop capacity region. For the multi-hop case, we also provide the analogous result.

Theorem 4.3: For $\alpha > \alpha_{\mathcal{P}}^{\text{sh}}$, there exists $\lambda \in \mathcal{D}_\alpha$ such that $\lambda \notin \Lambda_{\mathcal{P}}^{\text{sh}}$. Similarly, for $\alpha > \alpha_{\mathcal{P}}^{\text{mh}}$, there exists $\lambda \in \mathcal{D}_\alpha$ such that $\lambda \notin \Lambda_{\mathcal{P}}^{\text{mh}}$.

Applying Theorems 4.2 and 4.3, we can use results from the recent RWA literature on k-allowable traffic [13]–[17] to characterize the values $\alpha_{\mathcal{P}}^{\text{sh}}, \alpha_{\mathcal{P}}^{\text{mh}}$ for various physical topology configurations.

V. Conclusions

In this paper, we have studied the optimal throughput performance properties of reconfigurable WDM-based packet networks. We considered networks having arbitrary physical topologies, and general node architectures. Our architectural assuptions were deliberately made general to admit a variety of network features that can emerge in a future agile Terabit optical-based communication network, including wavelength conversion capability, tunable transceivers, optical switches, and multiplexers/demultiplexers.

Under our general network setting, we have presented optimal throughput-achieving algorithms for networks employing single-hop and multi-hop electronic routing. These on-line algorithms make use of queue backlog information at network nodes to simultaneously schedule packet routing at the electronic layer as well as WDM layer reconfiguration. In general, the stability region of arrival rates that can be supported in a particular network is described as a convex combination of available service configurations in that network.

In this work, we used RWA decompositions to establish the entire stability region under any physical topology in terms of the RWA properties of the same physical topology graph. The RWA problem with no conversion was tied to the single-hop capacity region of the reconfigurable network, while the RWA problem with conversion was tied to the multi-hop capacity region. 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.

These geometric properties provide a measure of the optimal achievable throughput under any physical topology. 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. Many additional details as well as the proofs of all theorems presented here can be found in [12].

References

[1] 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.

[2] 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.

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

[4] 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.

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

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

[7] 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.

[8] 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.

[9] 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.

[10] G. Weichenberg, V.W.S. Chan, and M. Medard, "On the capacity of optical networks: A framework for comparing different transport architectures," in *IEEE Proc. Information Communications (INFOCOM)*, Apr. 2006.

[11] A. Brzezinski and E. Modiano, "Dynamic reconfiguration and routing algorithms for IP-over-WDM networks with stochastic traffic," *J. Lightw. Technol.*, vol. 23, no. 10, pp. 3188–3205, Oct. 2005.

[12] A. Brzezinski and E. Modiano, "Achieving 100% throughput in reconfigurable optical networks," MIT/LIDS Technical Report #2677, Jan. 2006.

[13] P. Saengudomlert, *Architectural Study of High-Speed Networks with Optical Bypassing*, Ph.D. thesis, Massachusetts Institute of Technology, Sep. 2002.

[14] 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.

[15] P. Saengudomlert, E. Modiano, and R. Gallager, "On-line routing and wavelength assignment for dynamic traffic in WDM ring and torus networks," *IEEE/ACM Trans. Netw.*, Nov. 2005.

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

[17] 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.