

A Versatile Scheme for Routing Highly Variable Traffic in Service Overlays and IP Backbones

M. Kodialam T. V. Lakshman James B. Orlin Sudipta Sengupta
Bell Laboratories, Lucent Technologies, Holmdel, NJ, USA
Massachusetts Institute of Technology, Cambridge, MA, USA

Abstract—The emergence of new applications on the Internet like voice-over-IP, peer-to-peer, and video-on-demand has created highly dynamic and changing traffic patterns. In order to route such traffic with Quality-of-Service (QoS) guarantees *without requiring detection of traffic changes in real-time or reconfiguring the network in response to it*, we consider a routing and bandwidth allocation scheme that allows preconfiguration of the network such that all traffic patterns permissible within the network’s natural ingress-egress capacity constraints can be handled in a capacity efficient manner. The scheme routes traffic in two phases. In the first phase, incoming traffic is sent from the source to a set of intermediate nodes and then, in the second phase, from the intermediate nodes to the final destination. The traffic in the first phase is distributed to the intermediate nodes in predetermined proportions that depend on the intermediate nodes.

In this paper, we develop linear programming formulations and a fast combinatorial algorithm for routing under the scheme so as to maximize throughput (or, minimize maximum link utilization). We compare the throughput performance of the scheme with that of the optimal scheme among the class of all schemes that are allowed to even make the routing dependent on the traffic matrix. For our evaluations, we use actual Internet Service Provider topologies collected for the Rocketfuel project. We also bring out the versatility of the scheme in not only handling widely fluctuating traffic but also accommodating applicability to several widely differing networking scenarios, including (i) economical Virtual Private Networks (VPNs), (ii) supporting indirection in specialized service overlay models like Internet Indirection Infrastructure (i3), (iii) adding QoS guarantees to services that require routing through a network-based middlebox, and (iv) reducing IP layer transit traffic and handling extreme traffic variability in IP-over-Optical networks without dynamic reconfiguration of the optical layer. The two desirable properties of supporting indirection in specialized service overlay models and static optical layer provisioning in IP-over-Optical networks are not present in other approaches for routing variable traffic, such as direct source-destination routing along fixed paths.

I. INTRODUCTION

As the Internet continues to grow in size and complexity, it becomes increasingly difficult to predict future traffic patterns. Many emerging applications for the Internet are characterized by highly variable traffic behavior over time. Classical approaches to network design and planning rely on a model in which a single traffic matrix is estimated. When actual traffic does not conform to such assumptions (as is often the case), desired Quality-of-Service (QoS) cannot be guaranteed due to network congestion. Development of routing infrastructures that optimize network resources while accommodating extreme traffic unpredictability in a robust and efficient manner will be one of the defining themes in the next phase of expansion of the Internet.

In order to meet this requirement of robust and efficient network routing in a highly dynamic and changing traffic

environment, we consider a routing and bandwidth allocation scheme [9] that allows preconfiguration of the network such that all traffic patterns permissible within the network’s natural ingress-egress capacity constraints can be handled without network reconfiguration. Such preconfiguration simplifies network operation by avoiding the need to detect traffic changes in real-time and to reconfigure the network in response. The scheme routes traffic in two phases. In the first phase, incoming traffic is sent from the source to a set of intermediate nodes and then, in the second phase, from the intermediate nodes to the final destination. The traffic in the first phase is distributed to the intermediate nodes in predetermined proportions that depend on the intermediate nodes, as proposed in [9]. Throughout this paper, we will refer to this scheme as *two-phase routing*.

In order to fully comprehend the motivation behind the development of such a scheme, it is important to understand, from an Internet Service Provider’s (ISP) perspective, the difficulty of deploying and operating a more dynamic architecture that requires the measurement of possibly changing traffic in real-time as well as reconfiguring the network in response to such changes in order to provide QoS guarantees. We address these aspects in Sections II-A and IV.

The two-phase routing scheme is versatile not only in its ability to handle widely fluctuating traffic but also in its applicability to several widely differing networking scenarios. We illustrate this through example applications of the routing scheme to (i) economical Virtual Private Networks (VPNs), (ii) providing Internet Indirection Infrastructure (i3) [18] like functionality with QoS guarantees in a network, (iii) adding QoS guarantees to services that require routing through a network-based middlebox, and (iv) reducing IP layer transit traffic and handling extreme traffic variability in IP-over-Optical networks without dynamic reconfiguration of the optical layer.

A unique aspect of the i3 application arising from its indirection property is that unlike traditional networks, the final destination of a packet is not known at the network ingress. Hence, methods that need pre-provisioned paths to be set up between a network’s ingress and egress nodes for providing bandwidth guarantees are not usable. For the IP-over-optical network application, it is important that both paths and their associated bandwidths do not change with shifts in traffic. The scheme is well-suited to both these applications unlike existing routing methodologies.

We develop linear programming formulations and a fast combinatorial algorithm for routing under the scheme so as to maximize throughput (or, minimize maximum link utilization). We compare the throughput performance of two-phase routing with that of the optimal scheme among the class of all

schemes that are allowed to make the routing dependent on the traffic matrix. For our evaluations, we use actual ISP network topologies collected for the Rocketfuel project [17].

The combinatorial algorithm developed is a Fully Polynomial Time Approximation Scheme (FPTAS). An FPTAS is an algorithm that finds a solution with objective function value within $(1 + \epsilon)$ -factor of the optimal solution and runs in time that is a polynomial function of the input parameters and $\frac{1}{\epsilon}$. The input parameters in our problem are the number of nodes n and links m in the network, and the size (number of bits) of the input numbers (link capacities and node ingress-egress capacities). The value of ϵ can be chosen to provide the desired degree of optimality for the solution.

Throughput (which is the reciprocal of maximum link utilization) is an important but not the only optimization metric for network routing. For example, network capacity minimization has been considered in the context of two-phase routing in [9]. We focus on network throughput in this paper because it is one of the most common metrics used in the literature, it is used in capacity planning decisions by ISPs, it is directly related to other metrics like link congestion, and is useful for multi-period traffic planning when the traffic patterns scale (roughly) uniformly over time. When considering feasibility of a traffic matrix on (various what-if) capacitated network deployment scenarios, throughput is probably the most suitable metric to consider (feasibility is indicated by a throughput greater than or equal to 1).

The paper is structured as follows. In Section II, we discuss some aspects of the inherent difficulty in measuring traffic and introduce the traffic variation model. In Section III, we describe some application scenarios and their requirements. In Section IV, we argue why two of these requirements are not met by existing routing methodologies, thus making the case for two-phase routing. In Section V, we briefly discuss the two-phase routing scheme so as to provide context for this paper. We return to the application scenarios in Section VI and explain how two-phase routing meets the requirements outlined in Section III. Section VII introduces the throughput maximization problem for two-phase routing and provides linear programming formulations. In Section VIII, we develop a fast combinatorial algorithm for the problem. Performance evaluation of two-phase routing is presented in Section IX. Finally, we conclude in Section X. We briefly describe some notation before moving on to the next section.

A. Notation

We assume that we are given a network $G = (N, E)$ with node set N and (directed) edge set E where each node in the network can be a source or destination of traffic. Let $|N| = n$ and $|E| = m$. The nodes in N are labeled $\{1, 2, \dots, n\}$. The sets of incoming and outgoing edges at node i are denoted by $E^-(i)$ and $E^+(i)$ respectively. We let (i, j) represent a directed link in the network from node i to node j . To simplify the notation, we will also refer to a link by e instead of (i, j) . The capacity of link (i, j) will be denoted by u_{ij} . The *utilization* of a link is defined as the traffic on the link divided by its capacity.

II. TRAFFIC MEASUREMENT AND VARIABILITY

In an utopian network deployment scenario where complete traffic information is known and does not change over time, we can optimize the routing for that single traffic matrix – a large volume of research has addressed this problem. The most important innovation of the two-phase routing scheme is the handling of traffic variability in a capacity efficient manner through static preconfiguration of the network and without requiring either (i) measurement of traffic in real-time or (ii) reconfiguration of the network in response to changes in it. We address the difficulties associated with (i) in this section and then introduce the traffic variation model. The difficulties associated with (ii) for IP-over-Optical networks are addressed in Section IV.

A. Difficulties in Measuring Traffic

Network traffic is not only hard to measure in real-time but even harder to predict based on past measurements. Direct measurement methods do not scale with network size as the number of entries in a traffic matrix is quadratic in the number of nodes. Moreover, such direct real-time monitoring methods lead to unacceptable degradation in router performance. In reality, only aggregate link traffic counts are available for traffic matrix estimation. SNMP (Simple Network Management Protocol) provides this data via incoming and outgoing byte counts computed per link every 5 minutes. To estimate the traffic matrix from such link traffic measurements, the best techniques today give errors of 20% or more [13].

The emergence of new applications on the Internet like voice-over-IP, peer-to-peer, and video-on-demand has reduced the time-scales at which traffic changes dynamically, making it impossible to extrapolate past traffic patterns to the future. Currently, ISPs handle such unpredictability in network traffic by gross over-provisioning of capacity. This has led to ISP networks being under-utilized to as low as 20% [13].

B. Traffic Variation Model

We consider a traffic variation model where the total amount of traffic that enters (leaves) an ingress (egress) node in the network is bounded by the total capacity of all external ingress links at that node. This is known as the *hose model* and was proposed by Fingerhut et al. [7] and subsequently used by Duffield et al. [6] as a method for specifying the bandwidth requirements of a Virtual Private Network (VPN). Note that the hose model naturally accommodates the network's ingress-egress capacity constraints.

We denote the upper bounds on the total amount of traffic entering and leaving the network at node i by R_i and C_i respectively. The point-to-point matrix for the traffic in the network is thus constrained by these ingress-egress link capacity bounds. These constraints are the only known aspects of the traffic to be carried by the network, and knowing these is equivalent to knowing the row and column sum bounds on the traffic matrix. That is, any allowable traffic matrix $T = [t_{ij}]$ for the network must obey

$$\sum_{j \in N, j \neq i} t_{ij} \leq R_i, \quad \sum_{j \in N, j \neq i} t_{ji} \leq C_i \quad \forall i \in N$$

For given R_i and C_i values, denote the set of all such matrices that are partially specified by their row and column sums by $\mathcal{T}(\vec{R}, \vec{C})$, that is

$$\mathcal{T}(\vec{R}, \vec{C}) = \{[t_{ij}] \mid \sum_{j \neq i} t_{ij} \leq R_i \text{ and } \sum_{j \neq i} t_{ji} \leq C_i \ \forall i\}$$

We will use $\lambda \cdot \mathcal{T}(\vec{R}, \vec{C})$ to denote the set of all traffic matrices in $\mathcal{T}(\vec{R}, \vec{C})$ with their entries multiplied by λ .

Note that the traffic distribution T could be any matrix in $\mathcal{T}(\vec{R}, \vec{C})$ and could change over time. Two-phase routing provides a routing architecture that does not make any assumptions about T apart from the fact that it is partially specified by row and column sum bounds and can provide QoS guarantees for routing all matrices in $\mathcal{T}(\vec{R}, \vec{C})$ without requiring any detection of changes in traffic patterns or dynamic network reconfiguration in response to it.

III. MOTIVATING NETWORKING APPLICATIONS

We discuss some motivating networking architectures and applications that need to handle traffic variation and identify the requirements of a suitable routing scheme for each scenario. In the next section, we then argue why such requirements are not met by existing routing methodologies.

A. IP Backbones

Core (long-haul) networks of ISPs form the backbone of the Internet and span vast geographical areas (countries and continents). Each node in such a network, also called a Point-of-Presence (PoP), connects an access network (or, regional/metro network) to the core network. Internet backbones are often deployed by interconnecting routers over a switched optical backbone, also called an *IP-over-Optical network*. Because a router line card is typically 3-4 times more expensive than an optical switch card, an IP-over-Optical network architecture reduces network cost by keeping traffic mostly in the optical layer [15]. By removing transit traffic from the routers to the optical switches, the requirement to upgrade router PoP configurations with increasing traffic is minimized (since optical switches are more scalable with increasing port count than routers). Also, since optical switches are known to be much more reliable compared to routers [12], this makes the architecture more robust and reliable.

Routing in IP-over-Optical networks needs to make a compromise between keeping traffic at the optical layer (for the above reasons) and using intermediate routers for packet grooming in order to achieve efficient statistical multiplexing of data traffic. In addition, the routing must be able to handle traffic variability. The (current) traffic matrix is not only difficult to estimate but changes in the same may not be detectable in real time. Moreover, dynamic changes in routing in the network may be difficult or prohibitively expensive from a network operations perspective. In spite of the continuing research on IP-Optical integration, network deployments are far away from utilizing the optical control plane to provide bandwidth provisioning in real-time to the IP layer. These translate to the following requirements on the routing methodology:

- IP traffic must be routed “mostly” at the optical layer from source to destination routers. Intermediate IP layer transit may be required for grooming purposes.

- The optical layer (circuits and their bandwidth) must be statically provisioned a priori to provide bandwidth guarantees for end-to-end IP traffic. Routing at the IP layer cannot also be adaptive to traffic changes.
- Bandwidth guarantees must be provided for routing all traffic matrices.

B. Specialized Service Overlays

The Internet Indirection Infrastructure (i3) was proposed in [18] to ease the deployment of services – like mobility, multicast and anycast – on the Internet. i3 provides a rendezvous-based communication abstraction through indirection – sources send packets to a logical identifier, and receivers express interest in packets sent to an identifier. The rendezvous points are provided by i3 servers that forward packets to all receivers that express interest in a particular identifier. The communication between senders and receivers is through these rendezvous points over an overlay network. The i3 infrastructure does not store packets but only forwards them. It is important to note that i3 provides only a best-effort service like today’s Internet – it neither implements reliability nor guarantees ordered delivery on top of IP.

Two-phase routing can support indirection and provide *QoS guarantees* for variable traffic in specialized service overlays like i3. This is discussed in Section VI-B. Three important requirements in this context are:

- The routing from the source node(s) to the rendezvous points cannot depend on the final destination(s) of the packet, since this is unknown at the source.
- The traffic from the source nodes to the rendezvous points and from the latter to the destination nodes must be routed along bandwidth-guaranteed paths.
- These paths cannot be re-routed in response to changes in traffic patterns, and must have sufficient bandwidth to handle all possible traffic patterns subject to network ingress-egress constraints.

C. Middlebox Routing

Intermediate network elements (so called middleboxes), such as firewalls and transparent caches, are now commonplace. They provide important services like caching, load-balancing, and content filtering (for network security). To be effective, the services provided by such middleboxes are required to be comprehensive in the sense that every packet routed through the network must pass through at least one middlebox providing the service. In order to support a middlebox routing architecture, the routing scheme needs to not only provide bandwidth guarantees for variable traffic but also handle the additional constraint that all network traffic must pass through at least one intermediate network element node. Two-phase routing can naturally accommodate such an architecture in ISP networks. This is discussed in Section VI-C.

D. Other Scenarios of Interest

Another example application where the traffic matrix is unknown is the provisioning of network-based VPN services [3] to enterprise customers. VPNs typically provide network

connectivity among different sites of an enterprise. The traffic distribution between the sites is not known a priori - it may also change depending on time-of-day, day-of-week, special activities, etc. The enterprise customer specifies to the ISP only the total traffic volume and the peak rate out of a given site (e.g., if a site is connected to the ISP through a T1 link, this peak rate is about 1.5 Mbps). It is the ISP's task to transport all of the offered VPN traffic to the network and carry the traffic in accordance with the bandwidth guarantees provided in the Service Level Agreement (SLA). The traffic originating from or destined to a VPN node is limited only by the aggregate bandwidth connection of that node to the VPN.

Networks for grid computing also need to handle highly variable traffic patterns. In grid computing, a complex computational task is partitioned amongst different computing nodes that can be geographically distributed and are connected by a network. The communication patterns amongst grid computing nodes are highly unpredictable and also can require high burst rates. Since the traffic matrix is not known, one option is to dynamically reserve capacity over an underlying network but this approach will be too slow for grid computing applications.

IV. RELATED ROUTING METHODOLOGIES

We briefly review related work on routing with traffic variability and point out why such existing methods cannot meet the requirements outlined above for the various application scenarios.

Direct routing from source to destination (instead of in two phases) along *fixed* paths for the hose traffic model has been considered by Duffield et al. [6] and Kumar et al. [11]. In related work, Azar et al. [2] consider direct source-destination routing along fixed paths and show how to compute *relative guarantees* for routing an arbitrary traffic matrix with respect to the best routing for that matrix. However, they do not provide *absolute bandwidth guarantees* for routing variable traffic under the hose model.

In both these approaches, *direct source-destination paths* are fixed a priori for routing the traffic between each source-destination pair. Thus, the source needs to *know the destination of a packet* for routing it, without which the source cannot determine the path along which the packet should be forwarded. In specialized service overlay models like i3, the *final destination of a packet is not known at the source*. Thus, any of the above approaches cannot be used for routing in service overlay networks.

Direct source-destination routing, when applied to IP-over-Optical networks, routes packets from source to destination along direct paths in the optical layer. Note that even though the paths are fixed a priori and do not depend on the traffic matrix, their *bandwidth requirements change* with variations in the traffic matrix. Thus, bandwidth needs to be deallocated from some paths and assigned to other paths as the traffic matrix changes. (Alternatively, paths between every source-destination pair can be provisioned a priori to handle the maximum traffic between them, but this leads to gross overprovisioning of capacity, since all source-destination pairs cannot simultaneously reach their peak traffic limit in the hose traffic model.) This necessitates *dynamic reconfiguration* of the provisioned optical layer circuits (i.e., change in bandwidth)

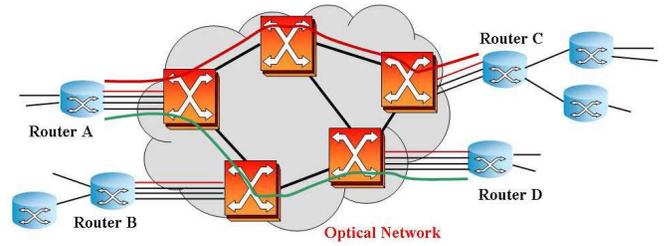


Fig. 1. Routing through direct optical layer circuits in IP-over-Optical networks.

in response to traffic variations, thus making direct source-destination routing unsuitable for IP-over-Optical networks.

To illustrate this last point, consider the scenario in Figure 1 for direct source-destination routing in IP-over-Optical networks. Here, router A is connected to router C using 3 OC-48 connections and to Router D using 1 OC-12 connection, so as to meet the traffic demand from node A to nodes C and D of 7.5 Gbps and 600 Mbps respectively. Suppose that at a later time, traffic from A to C decreases to 5 Gbps, while traffic from A to D increases to 1200 Mbps. Then, the optical layer must be reconfigured so as to delete one OC-48 connection between A and C and creating a new OC-12 connection between A and D. As such, the *requirement of static provisioning at the optical layer is not met*.

In contrast, two-phase routing has the following properties. They address both of the above issues.

- The source routes packets independent of their intended (or, unknown) destination, and
- Both the paths and their bandwidth are fixed a priori and do not need to be changed as traffic patterns change over time.

V. OVERVIEW OF TWO-PHASE ROUTING

In this section, we give an overview of the two-phase routing scheme from [9]. As mentioned earlier, the scheme does not require the network to detect changes in the traffic distribution or reconfigure the network in response to it. The only assumption about the traffic is the limits imposed by the ingress-egress constraints at each node, as outlined in Section II-B.

As is indicative from the name, the routing scheme operates in two phases:

- **Phase 1:** A predetermined fraction α_j of the traffic entering the network at any node is distributed to every node j *independent of the final destination of the traffic*.
- **Phase 2:** As a result of the routing in Phase 1, each node receives traffic destined for different destinations that it routes to their respective destinations in this phase.

This is illustrated in Figure 2. Note that the traffic split ratios $\alpha_1, \alpha_2, \dots, \alpha_n$ in Phase 1 of the scheme are such that $\sum_{i=1}^n \alpha_i = 1$. A simple method of implementing this routing scheme in the network is to form *fixed bandwidth paths between the nodes*. In order to differentiate between the paths carrying Phase 1 and Phase 2 traffic, we will refer to them as Phase 1 and Phase 2 paths respectively. The critical reason the two-phase routing strategy works is that the *bandwidth required for these tunnels depends on the ingress-egress capacities R_i, C_i and the traffic split ratios α_j but*

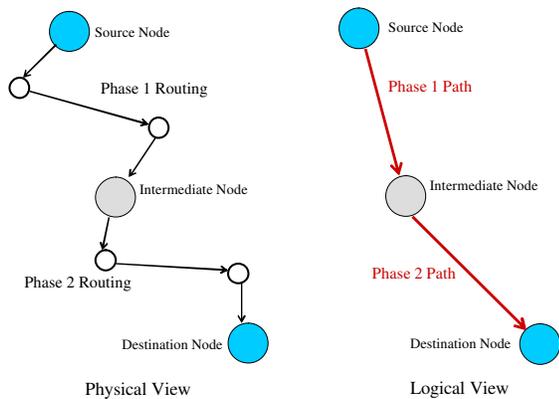


Fig. 2. Two-Phase Routing.

not on the (unknown) individual entries in the traffic matrix. Depending on the underlying routing architecture, the Phase 1 and Phase 2 paths can be implemented as IP tunnels, optical layer circuits, or Label Switched Paths in Multi-Protocol Label Switching (MPLS) [14].

We now derive the bandwidth requirement for the Phase 1 and Phase 2 paths. Consider a node i with maximum incoming traffic R_i . Node i sends $\alpha_j R_i$ amount of this traffic to node j during the first phase for each $j \in N$. Thus, the traffic demand from node i to node j as a result of Phase 1 routing is $\alpha_j R_i$. At the end of Phase 1, node i has received $\alpha_i R_k$ traffic from any other node k . Out of this, the traffic destined for node j is $\alpha_i t_{kj}$ since all traffic is initially split without regard to the final destination. The traffic that needs to be routed from node i to node j during Phase 2 is $\sum_{k \in N} \alpha_i t_{kj} \leq \alpha_i C_j$. Thus, the traffic demand from node i to node j as a result of Phase 2 routing is $\alpha_i C_j$.

Hence, the maximum demand from node i to node j as a result of routing in Phases 1 and 2 is $\alpha_j R_i + \alpha_i C_j$. Note that this does not depend on the matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$. The scheme handles variability in traffic matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$ by effectively routing the fixed matrix $D = [d_{ij}] = [\alpha_j R_i + \alpha_i C_j]$ that depends only on aggregate ingress-egress capacities and the traffic split ratios $\alpha_1, \alpha_2, \dots, \alpha_n$, and not on the specific matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$. This is what makes the routing scheme oblivious to changes in the traffic distribution.

An instance of the scheme requires specification of the traffic split ratios $\alpha_1, \alpha_2, \dots, \alpha_n$ and routing of the Phase 1 and Phase 2 paths. Computation of the above so as to maximize network throughput is the main focus of this paper.

The traffic split ratios α_i can be generalized to depend on source and/or destination nodes of the traffic, as proposed in [9]. While this does not meet the indirection requirement of specialized service overlays like i3, it can potentially increase the throughput performance of the two-phase routing scheme for other application scenarios like IP-over-Optical networks. We consider the problem of maximum throughput two-phase routing with generalized traffic split ratios in [10].

VI. APPLICATIONS OF TWO-PHASE ROUTING

We now return to the application scenarios described in Section III and discuss how our routing scheme can be applied to each scenario.

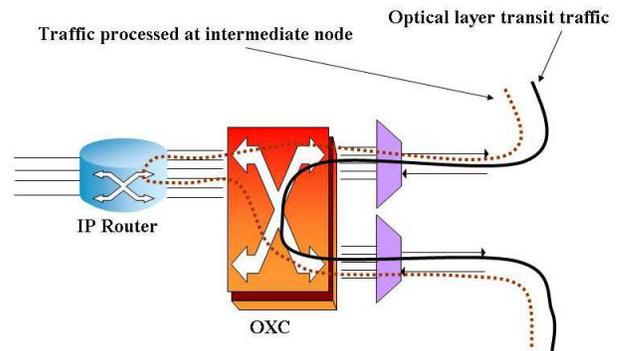


Fig. 3. Intermediate node packet processing for Two-Phase Routing in IP-over-Optical networks.

A. IP Backbones

Two-phase routing, as envisaged for IP-over-Optical networks, establishes the fixed bandwidth Phase 1 and Phase 2 paths at the optical layer. Thus, the *optical layer is statically provisioned* and does not need to be reconfigured in response to traffic changes. IP packets are routed end-to-end with *IP layer processing at a single intermediate node only*. While in transit at the optical layer inside either Phase 1 or Phase 2 tunnels, packets do enter the router but appear as transit traffic at the Optical Cross-Connect (OXC) only. The IP layer packet processing at an intermediate node works as follows. The optical layer circuit is dropped at the IP router at the node (through OXC-to-router links), wherein the packets are multiplexed back to the OXC (through router-to-OXC links) to be routed through direct optical layer circuits to their final destinations. Figure 3 illustrates optical layer transit traffic and intermediate node packet processing functionality at a node.

This architecture provides the desirable statistical multiplexing properties of packet switching for handling highly variable traffic *without significantly increasing the IP layer transit*. Compare this with the high levels of IP layer transit traffic in IP-over-WDM architecture where routers are directly connected to WDM systems and need to process packets at each hop.

B. Specialized Service Overlays

Two-phase routing can be used to provide QoS guarantees for variable traffic and support indirection in intra-ISP deployments of specialized service overlays like i3. (Note that we are not considering Internet-wide deployment here.) The intermediate nodes in the two-phase routing scheme are ideal candidates for locating i3 servers. Because we are considering a network whose topology is known, two-phase routing can be used to not only pick the i3 server locations (intermediate nodes) but also traffic engineer paths for routing with bandwidth guarantees between sender and receiver through i3 server nodes. Because the two-phase routing scheme can route the Phase 1 and Phase 2 paths with protection, this can also provide network level reliability of the services provided.

The ingress-egress traffic constraints R_i, C_j in the two-phase routing scheme now apply to network nodes to which hosts attach for using the services provided. For example, the host could be a laptop and a node could be a corporate site or

an ISP PoP. Mobility of the hosts manifest itself as changes in traffic originating from or destined to the network points of attachment (nodes), since mobile hosts will attach themselves to different nodes over time. The Phase 1 and Phase 2 paths of the specified bandwidth will provide bandwidth guarantees across all 13 applications described in [18], including mobility, multicast, and anycast. This is because the traffic arising from such applications obey, by default, the aggregate ingress-egress constraints at each node.

C. Middlebox Routing

Two-phase routing can naturally accommodate a middlebox routing architecture in ISP networks and also provide QoS guarantees for variable traffic. The intermediate nodes in two-phase routing are ideal locations for deploying middleboxes that provide functionalities like caching and content filtering. Because all traffic passes through one of the intermediate nodes in the scheme, the requirement of the middleware service to be comprehensive (in the sense that every packet routed through the network must be examined at least once) is also met. The routing can now provide end-to-end bandwidth guarantees for variable traffic patterns. Experiments on actual ISP topologies for maximum throughput two-phase routing in Section IX indicate that the number of intermediate nodes in two-phase routing is small compared to the total number of nodes in the network. Given that the deployment of services like content filtering are expensive (from a hardware perspective), a smaller number of intermediate nodes can lead to cost-effective deployment of such services.

VII. LINEAR PROGRAMMING FORMULATIONS

Given a network with link capacities u_e and constraints R_i, C_j on the ingress-egress traffic, we consider the problem of two-phase routing so as to maximize throughput. The throughput is the maximum multiplier λ such that all matrices in $\lambda \cdot \mathcal{T}(\vec{R}, \vec{C})$ can be feasibly routed.

In this section, we describe linear programming (LP) formulations for the above problem. Note that for the case of equal split ratios, i.e., $\alpha_i = \frac{1}{n}$ for all $i \in N$, the demand between nodes i and j is $(R_i + C_j)/n$, and the problem reduces to the maximum concurrent flow problem [16].

Suppose we relax the requirement that the traffic split ratios α_j sum to 1 in a feasible solution of the problem. Consider the sum $\lambda = \sum_{i \in N} \alpha_i$. The traffic split ratios can be divided by λ (normalized) so that they sum to 1, in which case all matrices in $\lambda \cdot \mathcal{T}(\vec{R}, \vec{C})$ can be feasibly routed. Thus, the appropriate measure of throughput is the quantity $\sum_{i \in N} \alpha_i$ when the traffic split ratios are not constrained to sum to 1.

A. Link Flow Based Formulation

We adopt the standard network flow terminology from [1]. Let x_e^{ij} denote the flow value on link e for routing $\alpha_j R_i + \alpha_i C_j$ amount of flow from source node i to destination node j . Then, the problem of two-phase routing so as to maximize throughput can be expressed as the following link indexed linear program:

$$\text{maximize } \sum_{i \in N} \alpha_i$$

subject to

$$\sum_{e \in E^+(k)} x_e^{ij} - \sum_{e \in E^-(k)} x_e^{ij} = \begin{cases} \alpha_j R_i + \alpha_i C_j & \text{if } k = i \\ -\alpha_j R_i - \alpha_i C_j & \text{if } k = j \\ 0 & \text{otherwise} \end{cases} \quad \forall i, j, k \in N \quad (1)$$

$$\sum_{i, j \in N} x_e^{ij} \leq u_e \quad \forall e \in E \quad (2)$$

$$\alpha_i \geq 0 \quad \forall i \in N \quad (3)$$

$$x_e^{ij} \geq 0 \quad \forall e \in E, \forall i, j \in N \quad (4)$$

Constraints (1) corresponding to the routing of $\alpha_j R_i + \alpha_i C_j$ amount of flow from node i to node j . Constraints (2) are the link capacity constraints. By using per-source flow variables x_e^i instead of per source-destination variables x_e^{ij} , the number of flow variables in the above linear program can be reduced by a factor of n .

The above linear program is of polynomial size and is amenable for solution with LP solvers like CPLEX [5]. However, it is well known that running times of general linear programming based algorithms for network problems do not scale well with increasing network size. So we propose to design a fast combinatorial algorithm (FPTAS) with performance guarantees for the problem.

B. Incorporating Node Capacity Constraints in IP-over-Optical Networks

Consider the deployment of our routing scheme in IP-over-Optical networks as discussed in Section VI-A. The end-to-end IP traffic traverses router-to-OXC links not only at the source and destination nodes but also at the intermediate nodes. This router-to-OXC traffic at a node is bounded by the aggregate connectivity of the IP router to the OXC at that node. Thus, we need to model such node capacity constraints in our problem formulation.

This is done by transforming the graph representation of the network as follows. Split each node into two sub-nodes, one representing the IP router and another the OXC at that node. All links incident at each node in the original graph are now incident at corresponding OXC sub-node. Add links in either direction connecting the router and OXC sub-nodes with capacity equal to the given router-to-OXC connectivity at that node. Traffic originates and terminates at the router sub-nodes in this transformed graph. Transit traffic traverses the OXC sub-nodes only, except at the intermediates nodes where it uses the router-to-OXC links to enter and leave the router sub-nodes. With this graph transformation, we can apply the problem formulation from Section VII-A in the context of IP-over-Optical networks.

C. Path Flow Based Formulation

In this section, we present a path indexed linear programming formulation for the above problem. This will be subsequently used to develop the fast combinatorial algorithm (FPTAS) in Section VIII.

Let \mathcal{P}_{ij} denote the set of all paths from node i to node j . Let $x(P)$ denote the traffic on path P .

$$\begin{aligned}
& \text{maximize} && \sum_{i \in N} \alpha_i \\
\text{subject to} &&& \\
& \sum_{P \in \mathcal{P}_{ij}} x(P) = \alpha_j R_i + \alpha_i C_j \quad \forall i, j \in N && (5) \\
& \sum_{P \ni e} x(P) \leq u_e \quad \forall e \in E && (6) \\
& \alpha_i \geq 0 \quad \forall i \in N && (7) \\
& x(P) \geq 0 \quad \forall P && (8)
\end{aligned}$$

In Section VIII, we state the dual of the linear program. In general, a network can have an exponential number of paths (in the size of the network). Hence, this (primal) linear program can have possibly exponential number of variables and its dual can have an exponential number of constraints – they are both not suitable for running on medium to large sized networks. The usefulness of the primal and dual formulation is in designing a fast combinatorial algorithm for the problem.

VIII. FAST COMBINATORIAL ALGORITHM

In this section, we develop a fast combinatorial algorithm (FPTAS) that computes the traffic split ratios and routing of Phase 1 and Phase 2 paths up to $(1 + \epsilon)$ -factor of the optimal objective function value (maximum throughput) for any $\epsilon > 0$. We begin with the dual formulation of the linear program discussed above. The primal-dual approach we develop is adapted from the technique in Garg and Könemann [8] for solving the maximum multicommodity flow problem, where flows are augmented in the primal solution and dual variables are updated in an iterative manner.

The dual formulation of the linear program outlined in Section VII-C associates a variable π_{ij} with each demand constraint in (3), and a non-negative variable $w(e)$ with each link capacity constraint in (4). Let $SP(i, j)$ denote the cost of the shortest path $P \in \mathcal{P}_{ij}$ under weights $w(e)$. That is,

$$SP(i, j) = \min_{P \in \mathcal{P}_{ij}} \sum_{e \in P} w(e)$$

After simplification and removal of the dual variables π_{ij} , the dual linear program can be written as below:

$$\begin{aligned}
& \text{minimize} && \sum_{e \in E} u_e w(e) \\
\text{subject to} &&& \\
& \sum_{i \neq k} R_i SP(i, k) + \sum_{j \neq k} C_j SP(k, j) \geq 1 \quad \forall k \in N && (9) \\
& w(e) \geq 0 \quad \forall e \in E && (10)
\end{aligned}$$

For a given node k and weights $w(e)$, let $V(k)$ denote the left-hand-side (LHS) of constraint (5). Given the weights $w(e)$, note that the values $V(k)$ for all $k \in N$ can be computed in polynomial time using a single all-pairs shortest path computation with link costs $w(e)$.

Given a set of weights $w(e)$, it is a feasible solution for the dual program if and only if

$$\min_{k \in N} V(k) \geq 1$$

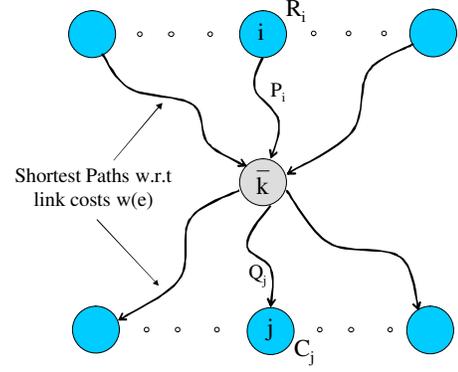


Fig. 4. One Step in the Primal-Dual Computation.

The algorithm works as follows. Start with equal initial weights $w(e) = \delta$ (the quantity δ depends on ϵ and is derived later). Repeat the following until the dual feasibility constraints are satisfied:

- 1) Compute the node $k = \bar{k}$ for which $V(k)$ is minimum. This identifies a node \bar{k} as well as shortest paths P_i from node i to node \bar{k} for all i and paths Q_j from node \bar{k} to node j for all j (see Figure 4).
- 2) For each $e \in E$, let $N_P(e)$ be the set of nodes i for which path P_i contains link e and $N_Q(e)$ be the set of nodes j for which path Q_j contains link e . Compute

$$\alpha = \min_{e \in E} \frac{u_e}{\sum_{i \in N_P(e)} R_i + \sum_{j \in N_Q(e)} C_j} \quad (11)$$

- 3) Send αR_i amount of flow on path P_i for all i and αC_j amount of flow on path Q_j for all j , and compute the total flow $\Delta(e)$ that is sent on link e for all $e \in E$. Increment the flow on link e by $\Delta(e)$.
- 4) Update the weights $w(e)$ for all $e \in E$ as

$$w(e) \leftarrow w(e) \left(1 + \frac{\epsilon \Delta(e)}{u_e}\right)$$

- 5) Increment the split ratio $\alpha_{\bar{k}}$ associated with node \bar{k} by α .

When the above procedure terminates, dual feasibility constraints will be satisfied. However, primal capacity constraints on each link will be violated, since we were working with the original (and not residual) link capacities at each stage. To remedy this, we scale down the flows and traffic split ratios α_i uniformly so that capacity constraints are obeyed.

Note that since the algorithm maintains primal and dual solutions at each step, the optimality gap can be estimated by computing the ratio of the primal and dual objective function values. The computation can be terminated immediately after the desired closeness to optimality is achieved.

The pseudo-code for the above procedure, called Algorithm MAX-THROUGHPUT, is provided in the box below. Array $flow(e)$ keeps track of the traffic on link e as the algorithm progresses. The variable G is initialized to 0 and remains < 1 as long as the dual constraints remain unsatisfied. After the **while** loop terminates, the maximum factor by which the capacity constraint is violated on any link is computed into $scale$. Finally, the α_i values are divided by the maximum capacity violation factor and the resulting values output.

Algorithm MAX-THROUGHPUT:

```

 $\alpha_k \leftarrow 0 \quad \forall k \in N;$ 
 $w(e) \leftarrow \delta \quad \forall e \in E;$ 
 $flow(e) \leftarrow 0 \quad \forall e \in E;$ 
 $G \leftarrow 0;$ 

while  $G < 1$  do
  Compute shortest path of cost  $SP(i, j)$  from
   $i$  to  $j$  under link costs  $w(e) \quad \forall i, j \in N;$ 
   $V(k) \leftarrow \sum_{i \neq k} R_i SP(i, k) + \sum_{j \neq k} C_j SP(k, j);$ 
   $G \leftarrow \min_{k \in N} V(k);$ 
   $\bar{k} \leftarrow \arg \min_{k \in N} V(k);$ 
  if  $G \geq 1$  break;
  (Denote shortest path from  $i$  to  $\bar{k}$  by  $P_i$  for all  $i$ 
  and shortest path from  $\bar{k}$  to  $j$  by  $Q_j$  for all  $j$ .)
   $N_P(e) \leftarrow \{i : P_i \text{ contains } e\}$  for all  $e$ ;
   $N_Q(e) \leftarrow \{j : Q_j \text{ contains } e\}$  for all  $e$ ;
   $\alpha \leftarrow \min_{e \in E} \frac{u_e}{\sum_{i \in N_P(e)} R_i + \sum_{j \in N_Q(e)} C_j};$ 
  Send  $\alpha R_i$  flow on path  $P_i$  for all  $i$  and
   $\alpha C_j$  flow on path  $Q_j$  for all  $j$  and compute
  resulting capacity usage  $\Delta(e)$  on link  $e$  for all  $e$ ;
   $flow(e) \leftarrow flow(e) + \Delta(e)$  for all  $e$ ;
   $w(e) \leftarrow w(e)(1 + \epsilon \Delta(e)/u_e)$  for all  $e$ ;
   $\alpha_{\bar{k}} \leftarrow \alpha_{\bar{k}} + \alpha;$ 
end while

 $scale \leftarrow \max_{e \in E} flow(e)/u_e$  for all  $e \in E$ ;
 $\alpha_k \leftarrow \alpha_k / scale$  for all  $k \in N$ ;
Output traffic split ratios  $\alpha_k$  for all  $k \in N$ ;

```

Let $L = (n - 1)(\sum_{i \in N} R_i + \sum_{j \in N} C_j)$ and let L' denote the minimum non-zero value of the R_i 's and C_j 's. The values of ϵ and δ are related, in the following theorem, to the approximation factor guarantee of Algorithm MAX-THROUGHPUT.

Theorem 1: For any given $\epsilon' > 0$, Algorithm MAX-THROUGHPUT computes a solution with objective function value within $(1 + \epsilon')$ -factor of the optimum for

$$\delta = \frac{1 + \epsilon}{L'[(1 + \epsilon) \frac{L}{L'}]^{1/\epsilon}} \quad \text{and} \quad \epsilon = 1 - \frac{1}{\sqrt{1 + \epsilon'}}$$

We end this section with a bound on the running time of Algorithm MAX-THROUGHPUT.

Theorem 2: For any given $\epsilon > 0$ chosen to provide the desired approximation factor guarantee in accordance with Theorem 1, Algorithm MAX-THROUGHPUT runs in time polynomial in the input size and $1/\epsilon$, that is

$$O\left(\frac{nm}{\epsilon^2}(m + n \log n) \log \frac{L}{L'}\right)$$

IX. EVALUATION ON ISP TOPOLOGIES

In this section, we evaluate the performance of two-phase routing. We first define a quantity called *throughput efficiency* that will be used to measure the effectiveness of two-phase routing against a general class of routing schemes that can handle traffic variability.

A. Throughput Efficiency

Given a network with link capacities u_e and bounds R_i, C_j on the traffic matrix, an output λ^* of the problem formulation in Section VII provides a guarantee that all matrices in $\lambda^* \cdot \mathcal{T}(\vec{R}, \vec{C})$ can be routed by two-phase routing. The highest

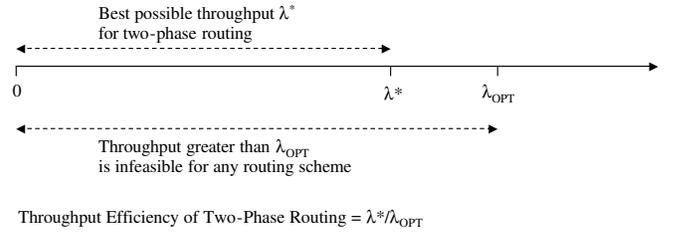


Fig. 5. Schematic Illustrating Throughput Efficiency of Two-Phase Routing.

possible throughput λ_{OPT} is admitted by the optimal scheme among the class of schemes that is allowed to make the routing dynamically dependent on the traffic matrix. This is illustrated in Figure 5. The ratio $\frac{\lambda^*}{\lambda_{OPT}} (\leq 1)$ is defined as the *throughput efficiency* of two-phase routing.

Note that the throughput efficiency, as defined above, is different from the oblivious ratio of Azar et al. [2]. In the latter case, the routing is compared with the *best routing for a single traffic matrix*. In our case, the routing is compared with *best scheme for routing all matrices in $\mathcal{T}(\vec{R}, \vec{C})$* .

It is shown in [10] that the throughput efficiency of two-phase routing is at least 0.5 (or, 50%) when the ingress-egress capacities are symmetric, i.e., $R_i = C_i$ for all i . The latter assumption holds for all the ISP topologies we use in our experiments because network routers and switches have bidirectional ports (line cards). We will see that the throughput efficiency of two-phase routing on the evaluated topologies is significantly better than this theoretical lower bound of 50%.

The value λ_{OPT} is hard to compute. Suppose that we take any single matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$ and compute the maximum multiplier $\lambda(T)$ (using a maximum concurrent flow formulation [16]) such that $\lambda(T) \cdot T$ can be feasibly routed in the network with given link capacities. Then, $\lambda_{OPT} \leq \lambda(T)$, and hence $\frac{\lambda^*}{\lambda(T)} \leq \frac{\lambda^*}{\lambda_{OPT}} \leq 1$. Thus, for any traffic matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$, the quantity $\frac{\lambda^*}{\lambda(T)}$ is a lower bound on the throughput efficiency of two-phase routing. To obtain a tight lower bound, we would like to identify a matrix $T \in \mathcal{T}(\vec{R}, \vec{C})$ for which $\lambda(T)$ is minimum. This matrix T is hard to compute. We use a heuristic approach to find a matrix that gives tight lower bounds.

B. Topologies and Link/Ingress-Egress Capacities

For our experiments, we use six ISP topologies collected by Rocketfuel, an ISP topology mapping engine [17]. These topologies list multiple intra-PoP (Point of Presence) routers and/or multiple intra-city PoPs as individual nodes. We coalesced PoPs into nodes corresponding to cities so that the topologies represent geographical PoP-to-PoP ISP topologies. Some data about the original Rocketfuel topologies and their coalesced versions is provided in Table I.

Link capacities, which are required to compute the maximum throughput, are not available for these topologies. Rocketfuel computed OSPF/IS-IS link weights for the topologies so that shortest cost paths match observed routes. In order to deduce the link capacities from the weights, we assumed that the given link weights are the default setting for OSPF weights in Cisco routers, i.e., inversely proportional to the

Topology	Routers (original)	Links (inter-router)	PoPs (coalesced)	Links (inter-PoP)
Telstra (Australia) 1221	108	306	57	59
Sprintlink (US) 1239	315	1944	44	83
Ebone (Europe) 1755	87	322	23	38
Tiscali (Europe) 3257	161	656	50	88
Exodus (Europe) 3967	79	294	22	37
Abovenet (US) 6461	141	748	22	42

TABLE I

ROCKETFUEL TOPOLOGIES: ORIGINAL NUMBER OF ROUTERS AND INTER-ROUTER LINKS, AND NUMBER OF COALESCED POPS AND INTER-POP LINKS.

Topology	Throughput Efficiency of Two-Phase Routing	Throughput Efficiency of Point-to-Point Pipe Model
Telstra (Australia) 1221	100%	5.39%
Sprintlink (US) 1239	97.71%	3.76%
Ebone (Europe) 1755	98.90%	7.33%
Tiscali (Europe) 3257	95.65%	5.97%
Exodus (Europe) 3967	100%	13.15%
Abovenet (US) 6461	94.82%	10.44%

TABLE II

THROUGHPUT EFFICIENCY OF TWO-PHASE ROUTING AND POINT-TO-POINT PIPE MODEL FOR ROCKETFUEL TOPOLOGIES.

Topology	Number of Intermediate Nodes
Telstra (Australia) 1221	1
Sprintlink (US) 1239	5
Ebone (Europe) 1755	4
Tiscali (Europe) 3257	7
Exodus (Europe) 3967	3
Abovenet (US) 6461	7

TABLE III

NUMBER OF INTERMEDIATE NODES IN TWO-PHASE ROUTING FOR ROCKETFUEL TOPOLOGIES.

link capacities [4]. The link capacities obtained in this manner turned out to be symmetric, i.e., $u_{ij} = u_{ji}$ for all $(i, j) \in E$.

There is also no available information on the ingress-egress traffic capacities at each node. Because ISPs commonly engineer their PoPs to keep the ratio of add/drop and transit traffic approximately fixed, we assumed that the ingress-egress capacity at a node is proportional to the total capacity of network links incident at that node. We also assume that $R_i = C_i$ for all nodes i – since network routers and switches have bidirectional ports (line cards), hence the ingress and egress capacities are equal. Thus, we have $R_i (= C_i) \propto \sum_{e \in E^+(i)} u_e$.

C. Experiments and Results

To obtain the maximum throughput for two-phase routing for purposes of comparison with that of the optimal scheme, we used the exact linear programming formulation from Section VII and solved it using CLPEx [5].

1) *Throughput Efficiency*: In Table II, we list the throughput efficiency of two-phase routing for the six Rocketfuel topologies. We compare this with the throughput efficiency of the *point-to-point pipe provisioning model* in which a fixed demand of $\min(R_i, C_j)$ is provisioned from node i to node j for all $i, j \in N$ to handle the maximum possible traffic from i and j under the given ingress-egress capacities. Similar to that for two-phase routing, the throughput efficiency of the point-to-point pipe model is measured relative to the throughput of the optimal scheme.

Table II clearly shows that the throughput of two-phase routing is very close to that of the best possible scheme for routing with traffic variability on all six Rocketfuel topologies. Thus, two-phase routing, surprisingly, is able to meet the requirements of Section III without any appreciable decrease in throughput compared to the optimal scheme. Table II also brings out the poor throughput performance of the point-to-

point pipe model, the throughput efficiency of which is in the range of 3-14%.

2) *Number of Intermediate Nodes*: In Table III, we list the number of intermediate nodes i with $\alpha_i > 0$ for maximum throughput two-phase routing on the six Rocketfuel topologies. Interestingly, the number of such intermediate nodes, especially for the larger topologies, is small compared to the total number of nodes. This may have favorable implications in the adaptation of the scheme to specialized service overlays and middlebox routing as explained in Section VI. In these two application scenarios, the intermediate nodes are sites for locating overlay routing servers and middleboxes respectively.

3) *Equal vs. Unequal Traffic Split Ratios*: For the two-phase routing scheme, we denote the throughput for equal traffic split ratios by λ_{equal} and the throughput for our general problem formulation that allows unequal traffic split ratios by $\lambda_{unequal}$. It is easy to see that $\lambda_{unequal} \geq \lambda_{equal}$. In Table IV, we give the throughput of two-phase routing with equal and unequal split ratios. The percentage increase in throughput $\frac{\lambda_{unequal} - \lambda_{equal}}{\lambda_{equal}}$ when we go from equal to unequal split ratios is also shown. When either the link capacities or ingress-egress capacities are scaled by a constant, the throughput values are scaled by the same constant. Hence, for comparison purposes, we have normalized the values so that the throughput for the unequal traffic split ratios case is $\lambda_{unequal} = 1.0$.

The results clearly bring out the increase in network throughput when the split ratios α_i are allowed to be unequal. The average savings for the six Rocketfuel topologies is 57.89% and the range is from 12% to as high as 152%. We conclude that by allowing the traffic split ratios to be unequal, network throughput for two-phase routing can be increased significantly over the equal traffic split ratios case.

X. RELATED WORK

In Section IV, we reviewed related work in [11], [2] for direct source-destination routing along fixed paths. We pointed

Topology	$\lambda_{unequal}$	λ_{equal}	$\frac{\lambda_{unequal} - \lambda_{equal}}{\lambda_{equal}}$
Telstra (Australia) 1221	1.0	0.7756	28.93%
Sprintlink (US) 1239	1.0	0.3978	151.38%
Ebone (Europe) 1755	1.0	0.6137	62.95%
Tiscali (Europe) 3257	1.0	0.6625	50.95%
Exodus (Europe) 3967	1.0	0.8908	12.26%
Abovenet (US) 6461	1.0	0.7098	40.89%

TABLE IV

THROUGHPUT OF TWO-PHASE ROUTING WITH UNEQUAL AND EQUAL TRAFFIC SPLIT RATIOS FOR ROCKETFUEL TOPOLOGIES.

out two aspects of these approaches that do not meet the requirements of application scenarios discussed in Section III, namely (i) the source needs to *know the final destination of a packet* for routing it, and (ii) the bandwidth requirements of the (fixed) paths change with traffic variations.

Because of (i), these methods cannot be used in specialized service overlay models like i3 where the *final destination of a packet is not known at the source*. Because of (ii), the adaptation of these methods for IP-over-Optical networks necessitates detection of changes in traffic patterns and dynamic reconfiguration of the provisioned optical layer circuits in response to it, a functionality that is not present in current IP-over-Optical network deployments.

The origin of two-phase routing can be traced back to Valiant's randomized scheme for communication among parallel processors interconnected in a hypercube topology [19].

Our current work is a sequel to [9]. In [20], a restricted version of the scheme with *equal traffic split ratios of $\frac{1}{n}$* and *equal ingress-egress capacities ($R_i = C_i = c$ for all i)* is considered. The authors in [20] further assume that the IP layer topology is a full-mesh (fully connected complete graph), so that the Phase 1 and Phase 2 paths are one hop in length. These paths need to be routed (via multi-hop paths) on the physical WDM topology (which is a sparse graph), an important aspect which they do not consider. Also, if the IP topology is not full-mesh, the Phase 1 and Phase 2 paths will be multi-hop at the IP layer itself. Our problem formulation for two-phase routing in [9] (and in this paper) models the multi-hop routing of Phase 1 and Phase 2 paths and can be applied to a general IP layer topology and a physical WDM topology.

XI. CONCLUSION

The two-phase routing scheme was recently proposed for routing highly dynamic and changing traffic patterns on the Internet with QoS guarantees. If deployed, it will allow service providers to operate their networks in a quasi-static manner where both intra-domain paths and the bandwidths allocated to these paths is robust to extreme traffic variation. The scheme has the desirable properties of supporting (i) indirection in specialized service overlay models like i3, and (ii) static optical layer provisioning in IP-over-Optical networks. To our knowledge, this routing scheme is the only one that is sufficiently versatile to handle the needs of such diverse applications that we studied while also being robust to extreme traffic fluctuation.

In this paper, we developed linear programming formulations and a fast combinatorial algorithm for routing under the scheme so as to maximize throughput. We compared the

throughput performance of two-phase routing with that of the optimal scheme among the class of all schemes that are allowed to make the routing dynamically dependent on the traffic matrix. Experiments on actual ISP topologies taken from the Rocketfuel project show that the throughput of two-phase routing with intermediate node dependent traffic split ratios is within 6% of the optimal scheme on all evaluated topologies. Thus, two-phase routing achieves its robustness to traffic variation and its versatility in being applicable to the discussed networking scenarios without any significant over-provisioning of capacity.

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, February 1993.
- [2] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke, "Optimal oblivious routing in polynomial time", *35th ACM Symposium on the Theory of Computing (STOC)*, 2003.
- [3] W. Ben-Ameur and H. Kerivin, "New Economical Virtual Private Networks", *Communications of the ACM*, Vol. 46, No. 6, pp. 69-73, June 2003.
- [4] "Configuring OSPF", Cisco Systems Product Documentation, <http://www.cisco.com/univercd/home/home.htm>.
- [5] ILOG CPLEX, <http://www.ilog.com>.
- [6] N. G. Duffield, P. Goyal, A. G. Greenberg, P. P. Mishra, K. K. Ramakrishnan, J. E. van der Merwe, "A flexible model for resource management in virtual private network", *ACM SIGCOMM 1999*, August 1999.
- [7] J. A. Fingerhut, S. Suri, and J. S. Turner, "Designing Least-Cost Nonblocking Broadband Networks", *Journal of Algorithms*, 24(2), pp. 287-309, 1997.
- [8] N. Garg and J. Könemann, "Faster and Simpler Algorithms for Multi-commodity Flow and other Fractional Packing Problems", *39th Annual Symposium on Foundations of Computer Science (FOCS)*, 1998.
- [9] M. Kodialam, T. V. Lakshman, and S. Sengupta, "Efficient and Robust Routing of Highly Variable Traffic", *Third Workshop on Hot Topics in Networks (HotNets-III)*, November 2004.
- [10] M. Kodialam, T. V. Lakshman, and S. Sengupta, "Maximum Throughput Routing of Traffic in the Hose Model", *IEEE Infocom 2006*, April 2006.
- [11] A. Kumar, R. Rastogi, A. Silberschatz, B. Yener, "Algorithms for provisioning VPNs in the hose model", *ACM SIGCOMM 2001*, August 2001.
- [12] C. Labovitz, A. Ahuja, and F. Jahanian, "Experimental Study of Internet Stability and Backbone Failures", *Proceedings of 29th International Symposium on Fault-Tolerant Computing (FTCS)*, Madison, Wisconsin, pp. 278-285, June 1999.
- [13] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, C. Diot, "Traffic Matrix Estimation: Existing Techniques and New Directions", *ACM SIGCOMM 2002*, August 2002.
- [14] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [15] S. Sengupta, D. Saha, and V. P. Kumar, "Switched Optical Backbone for Cost-effective Scalable Core IP Networks", *IEEE Communications Magazine*, June 2003.
- [16] F. Shahrokhi and D. Matula, "The Maximum Concurrent Flow Problem", *Journal of ACM*, 37(2):318-334, 1990.
- [17] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP Topologies with Rocketfuel", *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 2-16, February 2004.
- [18] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, S. Surana, "Internet Indirection Infrastructure", *ACM SIGCOMM 2002*, August 2002.
- [19] L. G. Valiant, "A scheme for fast parallel communication", *SIAM Journal on Computing*, 11(7), pp. 350-361, 1982.
- [20] R. Zhang and N. McKeown "Designing a Predictable Internet Backbone Network", *Third Workshop on Hot Topics in Networks (HotNets-III)*, November 2004.