A Comparative Analysis of the Effects of Dynamic Optical Circuit Provisioning on IP Routing

Mohit Chamania and Admela Jukan
Technische Universität Carolo-Wilhelmina zu Braunschweig
Email: {chamania, jukan}@ida.ing.tu-bs.de

***Abstract*—We analyze the effects of dynamic optical circuit setup on IP routing in general, and on two routing mechanisms in particular, i.e., explicit routing and shortest-path-first routing. We present analytical models for computing the size and placement of optical circuits, and propose model adaptations driven by the IP router system design. The results show that without careful consideration of intrinsic capabilities of IP routing protocol and forwarding, the size and location of optical circuits used can be vastly underestimated, also leading to significant disruptions in real networks. We present the Optical Bypass mechanisms and show that these methods, unlike traditional IP routing based solutions, affect a comparatively lower number of IP routes and can be computed near-optimally, even under unknown traffic matrix conditions, making them effective and feasible.**

I. Introduction

In the last decade, growing demand for Internet-based applications has led to significant innovations in networking. Among them is the widespread deployment of optical networking technologies. A number of IP-optical solutions already exist, differing in the way optical circuits are advertised in the IP routing. In large scientific networks, end-to-end optical circuits are setup between ingress and egress core routers and policy-based routing rules [1] are established at the routers to *offload* traffic from the scientific applications onto the optical circuits, while the IP routing protocol, and consequently the residual traffic in the network, remains unaffected. Alternately, in a typical telecom scenario, optical circuits created between routers are advertised as a new link in the IP routing protocol, initiating automatic convergence to a new routing configuration. Both scenarios have to carefully plan and control the introduction of new optical circuits to ensure that the impact on IP routing and network operations is minimal.

The transient nature and growth of Internet traffic is now driving the need for a more dynamic mechanism to setup optical circuits in IP networks, which, while conceptually tested, still represents a significant management challenge. In OSPF, for instance, routing re-convergence due to a newly added IP link (i.e., optical circuit advertised in IP routing) can take up to 30 seconds and can cause temporary routing loops, leading to disruption of network services. A change in routing is often also accompanied by significant reconfigurations in network management functions, such as the reconfiguration of alarm correlation functions [2]. Therefore, in addition to proper dimensioning of the circuit capacity, it is essential to fundamentally understand and quantify the interplay between capacity demands and the consequent effects on IP routing in order to deploy a dynamically reconfigurable IP-optical infrastructure. This in turn requires accurate modeling of key system parameters, such as IP routing protocols, forwarding capabilities and available IP traffic measurements.

In this paper, we present a comparative analysis of the effects of dynamic optical circuit provisioning on IP routing, by accurately modeling two major IP routing mechanisms, i.e., Explicit Routing (ER) and Shortest Path First (SPF) routing. Our analysis framework uses Integer Linear Programming (ILP) and incorporates features intrinsic to current routing protocols that have not been studied to date. As the application of ER and SPF schemes can affect a large number of IP routes as seen in our results, we compare these schemes with a newly proposed, namely Optical Bypass (BY) applicable to today's routers. Our model includes a mechanism to compute optical circuits using unknown IP traffic matrix, which is its main strength. The results also show that without consideration of intrinsic routing protocol features and forwarding capabilities, the size and location of optical circuits used can be vastly underestimated leading to degraded performance when applied to real networks. We show that optical bypass, unlike the ER and SPF schemes, affect a comparatively lower number of IP routes and can be computed near-optimally, even under unknown traffic matrix conditions.

The rest of the paper is organized as follows. In Section II, we describe the related work. Section III presents the network architecture and model assumptions. Section IV formulates the problem, objective function and constraints. Section V elaborates on the model formulation and adaptations driven by the system design. Section VI presents the performance evaluation study, while section VII concludes the paper.

II. Related Work and Our Contribution

Typical IP-optical network models assume that optical circuits are advertised as new IP links, and that IP routing follows a single(not necessarily shortest) path, which is essentially only applicable to MPLS-based networks with explicitly computed routes [3], [4], [5], [6], [7]. Explicit configurations are a challenge in real networks due to the management overhead and network operators usually prefer to use IP routing protocols with automated routing re-convergence, such as Shortest Path First routing (SPF). To model SPF routing, **two fundamental constraints must be enforced, i.e., destination-based forwarding and routing re-convergence, which are neglected in all models**. For instance, in a network where multiple shortest paths exist, the path computed in a model may use different outgoing links at intermediate hops to the same destination, which violates destination based forwarding and is therefore not applicable to typical IP networks. Also, in presence of multiple shortest paths, existing models typically compute a solution where traffic is re-routed over an alternate equal-cost shortest route. However, the SPF protocol re-converges to a new route only when a *shorter* route is found, and never over an alternative equal-cost route.

Most work also assume complete knowledge of traffic matrix, i.e., the traffic demand between all ingress-egress router pairs. However, to calculate the traffic matrix, the significant monitoring infrastructure and configuration effort is involved, which makes the computation impractical in large networks [8]. It is also for this reason, that the only successful deployment of dynamic circuits has been in science networks, where scientific applications with known capacity demands could request an end-to-end optical circuit in an IP network, with aid of a centralized controller, as seen in Lambdastation [9] and Phoebus [10]. The other Internet traffic however, typically consists of a large number of application flows with relatively low individual bandwidth requirement and hence solutions from scientific networks are not directly applicable. While flow detection techniques to identify long-lived high-bandwidth flows [11], [12] have been proposed for offloading special application flows in the Internet, they have only had limited success. It is therefore a question whether new analytical solutions can be developed for dynamic circuit setup using limited, easily measurable IP traffic parameters.

This paper is a major extension of our paper [13], including the consideration of destination-based forwarding and conditions on routing protocol re-convergence (Sec. IV-B2) which have not been considered elsewhere, and are essential. In our prior work [14], [15], [16], we advocated that methods are needed to setup optical circuit with minimum impact on IP routing, and proposed one such mechanism i.e., *optical bypass*. Though bypass was not a new term, our definition included a few distinct features. Akin to science networks, the optical bypass was not advertised in the IP routing protocol to limit routing changes and *explicit forwarding rules* were created at the ingress routers. However, we allowed the bypasses to be setup between any two routers in the network, and offload *aggregate traffic* (i.e., not only application-specific traffic), which was novel. Our optical bypass mechanism has since been studied by others in [17], [18], [19]. In this paper, we incorporate the aforementioned extensions to optical bypass (Sec. V-B) and compare its performance with the IP routing schemes. Finally, we present a generic ILP formulation in the absence of traffic matrix information (Sec. V-C) which is novel and different from [16], [3], [19], as it provides a guaranteed upper bound on the total optical capacity required.

III. Network Architecture and Assumptions

As illustrated in Fig. 1, we assume that IP/MPLS routers are co-located with optical switches, and an optical circuit established between two optical switches can be used to create

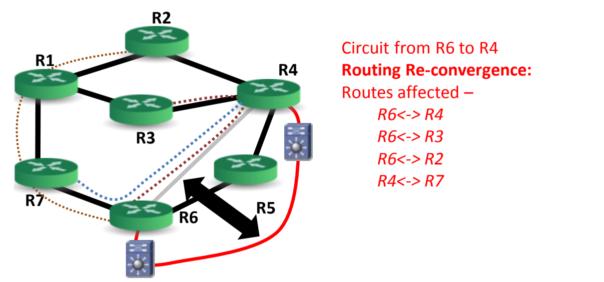

(a) Shortest Path First

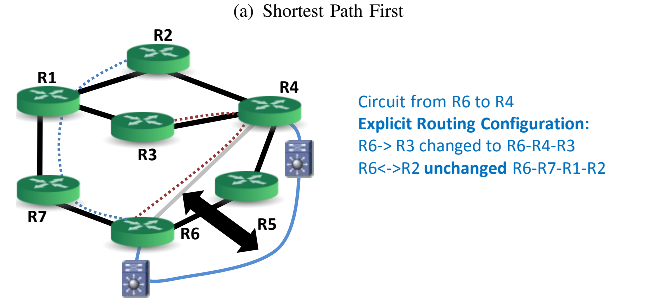

(b) Explicit Routing

Fig. 1. Schemes to introduce dynamic optical circuits into the IP network.

a link between the corresponding routers. Dynamic optical circuits can be used to boost the capacity of existing IP link or to create new IP links. Our application scenario is a core network, where every router acts as the edge router to other transit/stub networks and thus traffic is present between all pairs of routers. Also, IP routing always uses a single routing path established either explicitly (e.g., MPLS), or via a shortest path first protocol (e.g., OSPF). We assume that IP links are bi-directional, as it is commonly the case for IP interfaces and links to support functions such as OSPF neighbor discovery and Bidirectional Forwarding Detection [20]. At the same time, we use directed optical circuits so that IP links can have asymmetric capacities in different directions.

The **Shortest Path First (SPF)** mechanism assumes that the routing protocol automatically re-converges to a new configuration when IP topology changes. Instead of a specific SPF variant, we model typical features common to all SPF protocols and assume that 1) routing uses a single *shortest* path from source to destination, 2) forwarding is based on the destination IP address, 3) routing protocol re-converges only when a *shorter* route is found and 4) the upstream/downstream routes between a pair of routers are symmetric. In Fig. 1(a), the circuit from $R6$ to $R4$ is advertised as a new IP link. Under the constraints of SPF, this requires that the traffic from $R6$ to $R4$, $R3$ and $R2$, etc., use the new IP link from $R6$ to $R4$ due to the shortest path constraint.

The **Explicit Routing (ER)** (a.k.a static routing) supports any desired single-path routing configuration. In MPLS networks, routing can be defined based on the source as well as destination of traffic, using explicitly configured MPLS paths. While this flexibility makes load-balancing and traffic engineering feasible, the configuration effort increases significantly. Fig. 1(b) illustrates the explicit routing mechanism, where an optical circuit from $R6$ to $R4$ is advertised as a new IP link. Unlike the SPF example, the ER configuration allows that only the traffic from $R6$ to $R3$ is re-routed over the new link while the routing of other paths (e.g. $R6-R2$ shown in Fig. 1(b)) remains unchanged.

We now present the assumptions and notations used in our analysis. The notations used are summarized in Appendix A.

1) Network Topology Parameters: The IP network topology is represented as a directed graph $G(V,E)$ with routers $v_i \in V$ and directed links $e_{ij} \in E$ from v_i to v_j. IP links are bi-directional but can have different capacities in different directions. The representation contains *fixed* optical circuits, and *dynamic* optical circuits can be included to change capacity.

2) Optical Transport Network Parameters: Commercial optical network systems support a number of discrete circuit granularities, such as OTU1, OTU2, etc. in Optical Transport Networks (OTN) [21], or, OC1, OC3 etc. in SONET/SDH systems. The optical network supports a set of T circuit granularities with bandwidth of type $t \in T$ given by C_t^O. Circuits are *unidirectional* and the integer variable X_{xy}^t indicates the number of circuits with capacity C_t^O from v_x to v_y. Multiple optical circuits (maximum N circuits) can be aggregated to create an IP link between a pair of routers, and the boolean variable X_{xy} indicates if at least one dynamic circuit is established from v_x to v_y. We assume that dynamic circuits can always be established and the cost associated with establishing a circuit C_t^O from v_x to v_y is given by the constant $Cost_{xy}^t$. Since our focus is on the impact of optical circuit setup on IP routing, we do not discuss the routing of optical circuits in detail, but present how routing of optical circuits can be easily incorporated within this model in Appendix B.

3) IP Routing Parameters: In the IP network, all links that use fixed optical circuits (i.e. before creation of new optical circuits) are indicated using the constant $\hat{L}_{ij}$, where $\hat{L}_{ij} = \hat{L}_{ji} = 1$ if a link exists between v_i and v_j, while the boolean variable L_{ij} indicates the existence of a link after optical circuit setup. A boolean routing variable r_{ij}^{sd} indicates if the route from v_s to v_d uses the link from v_i to v_j after optical circuit creation. For example, in Fig. 1(b), $L_{R6R4} = 1$ indicates that a link exists between $R6$ and $R4$, while $r_{R6R4}^{R6R3} = 1$ indicates that the route from $R6$ to $R3$ uses the link from $R6$ to $R4$ in the final configuration. The SPF mechanism requires the definition of link weights, which are equal to $\hat{\omega}_{ij}$ if the link exists, or are otherwise substituted with a large constant ω_∞. Constant link weights imply an inherent assumption that optical circuit installed between already connected routers only boost link capacity and do not affect the IP routing. The routing cost variable RC_{sd} is used to enforce shortest path routing constraints while the boolean variable $FT_i^d(j)$ mimics a forwarding table to enforce destination-based forwarding.

4) IP Traffic Matrix: When the complete traffic matrix is known, traffic from source (ingress) v_s to destination (egress) v_d is given by $\hat{\lambda}_{sd}$. When the traffic matrix is not known, traffic from v_s to v_d is represented using a variable λ_{sd} and we present a formulation in Section V-C that uses basic measurements like IP link loads and traffic in virtual output queues to compute optical circuits.

IV. Problem Formulation

The objective function (1) minimizes the total cost of introducing optical circuits in the IP network while subject to network configuration constraints, IP routing and traffic constraints, as defined next.

$$Min: \sum_t \sum_{xy} X_{xy}^t \cdot Cost_{xy}^t \quad (1)$$

A. Network Configuration Constraints

As stated before, at most N optical circuits may be established between a pair of routers to boost the capacity of an existing IP link or to create a new IP link (2). (3) ensures that the boolean $X_{xy} = 0$ if no circuits are established $(\sum_t X_{xy}^t = 0)$ while (4) ensures that $X_{xy} = 1$ if at least one circuit is established. Finally, (5) ensures that IP links are bi-directional, but need not have symmetric capacity in both directions. All routers are assumed to be capable of establishing dynamic circuits, unless otherwise defined. For example, $(\sum_j X_{ij} = 0)$ can be used in case v_i cannot setup a circuit. The total final capacity of an IP link C_{ij} (6) is computed using the fixed capacity (installed in the base topology) $\hat{C}_{ij}$ and the capacity of the installed dynamic circuits $X_{ij}^t \cdot C_t^O$.

$$\forall v_x, v_y \in V, x \neq y \;:\; \sum_t X_{xy}^t \leq N \quad (2)$$

$$\forall v_x, v_y \in V, x \neq y, t \in T \;:\; X_{xy} \leq \sum_t X_{xy}^t \quad (3)$$

$$\forall v_x, v_y \in V, x \neq y, t \in T \;:\; X_{xy} \geq N^{-1} \cdot \sum_t X_{xy}^t \quad (4)$$

$$\forall v_x, v_y \in V, x \neq y \;:\; X_{xy} = X_{yx} \quad (5)$$

$$\forall v_i, v_j \in V \;:\; C_{ij} = \hat{C}_{ij} + \sum_t X_{ij}^t \cdot C_t^O \quad (6)$$

B. IP Routing and Traffic Constraints

1) Explicit Routing (ER): A link L_{ij} exists if either fixed optical circuits exist between the routers $(\hat{L}_{ij} = 1)$ (7), or if a new optical circuit was established $(X_{ij} = 1)$ (8). In the absence of both (9) ensures that $L_{ij} = 0$. As the topology for IP routing may change during operation, (10) constrains r_{ij}^{sd} to only use a link if it exists.

$$\forall v_i, v_j \in V,\ i \neq j \quad : \quad L_{ij} \geq \hat{L}_{ij} \quad (7)$$

$$\forall v_i, v_j \in V,\ i \neq j \quad : \quad L_{ij} \geq X_{ij} \quad (8)$$

$$\forall v_i, v_j \in V,\ i \neq j \quad : \quad L_{ij} \leq \hat{L}_{ij} + X_{ij} \quad (9)$$

$$\forall v_s, v_d, v_i, v_j \in V \quad : \quad r_{ij}^{sd} \leq L_{ij} \quad (10)$$

Explicit routing provides the operator complete control on IP routing, and the constraints only ensure routing continuity i.e. the existence of a single valid route between each source-destination router pair. (11), (12) ensure that the route from v_s to v_d uses exactly one outgoing link from v_s and one incoming link into v_d, while (13) ensures that equal number of incoming and outgoing links (≤ 1) on v_j $(j \neq s, d)$ are used to route traffic from v_s to v_d. Finally, (14) uses the routing and (known) traffic matrix information ($\hat{\lambda}_{sd}$) to ensure that no link is overloaded, i.e., the capacity utilization of all links is below the defined threshold α. These constraints are necessary and sufficient for single path routing in IP/MPLS architectures.

$$\forall v_s, v_d \in V \ : \ \sum_j r_{sj}^{sd} \ = \ \sum_j r_{jd}^{sd} \ = \ 1 \tag{11}$$

$$\forall v_s, v_d \in V \ : \ \sum_j r_{js}^{sd} \ = \ \sum_j r_{dj}^{sd} \ = \ 0 \tag{12}$$

$$\forall v_s, v_d, v_j \in V, j \neq s, d \quad : \quad \sum_i r_{ij}^{sd} = \sum_k r_{jk}^{sd} \leq 1 \tag{13}$$

$$\forall v_s, v_d, v_i, v_j \in V \quad : \quad \sum_{sd} r_{ij}^{sd} \cdot \hat{\lambda}_{sd} \leq \alpha \cdot C_{ij} \tag{14}$$

2) Shortest Path First (SPF) mechanism: IP routing under SPF also follows single path routing which means that (7), (8), (9), (10), (11), (12),(13), (14) from the ER formulation are also applicable here. However, automatic routing convergence within SPF constrains the routing to be symmetric, i.e.,

$$\forall v_i, v_j, v_s, v_d \in V \ : \ r_{ij}^{sd} = r_{ji}^{ds} \tag{15}$$

As specified before, IP link weight is given by $\hat{w}_{ij}$ if the link exists, and is otherwise infinite (w_∞). The actual link weight w_{ij} is given by (16) and route cost from v_s to v_d (RC_{sd}) is computed as the sum of the weights of all links on the route (17). The term $w_{ij} \cdot r_{ij}^{sd}$ is non-linear, but is equivalent to $\hat{w}_{ij} \cdot r_{ij}^{sd}$: in case a link does not exist ($L_{ij} = 0, \ w_{ij} = w_\infty$), the routing variable r_{ij}^{sd} is already constrained to be 0 in (10) and when $L_{ij} = 1$, we have $w_{ij} = \hat{w}_{ij}$ (16).

$$\forall v_i, v_j, \in V \ : \ w_{ij} = L_{ij} \cdot \hat{w}_{ij} + (1 - L_{ij}) \cdot w_\infty \tag{16}$$

$$\forall v_s, v_d, \in V \ : \ RC_{sd} = \sum_{ij} w_{ij} \cdot r_{ij}^{sd} = \sum_{ij} \hat{w}_{ij} \cdot r_{ij}^{sd} \tag{17}$$

(18) ensures that the routing cost should be less than or equal to the routing cost to any immediate neighbor (of the destination) plus the link weight to the destination. In this formulation, any router can be an immediate neighbor as we establish new links in the IP network. This presents a sufficient condition to ensure the use of the shortest path in IP routing.

$$\forall v_s, v_d, \in V, s \neq d : RC_{sd} \leq RC_{sx} + w_{xd} \tag{18}$$

SPF routing also follows destination-based forwarding and the variable $FT_i^d(j)$ mimics a forwarding table, with $FT_i^d(j) = 1$ indicating that traffic at v_i to destination v_d is forwarded over the link from v_i to v_j. (19) ensures that v_i can only have v_j as a next hop if a link exists between v_i and v_j, (20) ensures that v_i only has one next hop router choice for destination v_d, and (21) constrains the relationship between the routing r_{ij}^{sd} and $FT_i^d(j)$, ensuring that the route from v_s to v_d can only use a link from v_i to v_j if $FT_i^d(j) = 1$, i.e.,

$$\forall v_s, v_d, v_i, v_j \in V, \ i \neq d : FT_i^d(j) \leq L_{ij} \tag{19}$$

$$\forall v_d, v_i \in V, i \neq d : \sum_j FT_i^d(j) = 1 \tag{20}$$

$$\forall v_s, v_d, v_i, v_j \in V, \ i \neq d : FT_i^d(j) \geq r_{ij}^{sd} \tag{21}$$

SPF changes a route only if a *shorter* path is found to a destination. To enforce the same, we assume knowledge of the routing cost before establishment of optical circuits, $\hat{RC}_{sd}$ and (22) ensures that in case the new routing cost RC_{sd} is equal to the original routing cost $\hat{RC}_{sd}$, all routing variables remain unchanged ($\hat{r}_{ij}^{sd} = r_{ij}^{sd}$). If routing cost after establishment of optical circuits is different, even a minor change in routing cost can result in significant change in the route used, and a large positive constant P is used to ensure that the expression $P \cdot | \left(\hat{RC}_{sd} - RC_{sd} \right) |$ is always larger than or equal to the change in the number of links used for the route from v_s to v_d. Since (22) is non-linear, we use a boolean variable V_{sd} and a large positive constant P_∞, to linearize the same using equations (23) and (24). The use of the boolean variable V_{sd} means that the positive constant P_∞ would only be used either in (23) or in (24). As the expression $\sum_{ij} \hat{r}_{ij}^{sd} + r_{ij}^{sd} - 2\hat{r}_{ij}^{sd} r_{ij}^{sd} \geq 0$, a necessary condition for both the constraints to be valid is equivalent to the one presented in (22).

$$\forall v_s, v_d \in V, \ s \neq d :$$

$$P \cdot | \left(\hat{RC}_{sd} - RC_{sd} \right) | \geq \sum_{ij} \hat{r}_{ij}^{sd} + r_{ij}^{sd} - 2\hat{r}_{ij}^{sd} r_{ij}^{sd} \tag{22}$$

$$P \cdot \left(\hat{RC}_{sd} - RC_{sd} \right) + P_\infty \cdot V_{sd} \geq \sum_{ij} \hat{r}_{ij}^{sd} + r_{ij}^{sd} - 2\hat{r}_{ij}^{sd} r_{ij}^{sd} \tag{23}$$

$$P \cdot \left(RC_{sd} - \hat{RC}_{sd} \right) + P_\infty \cdot (1 - V_{sd}) \geq \sum_{ij} \hat{r}_{ij}^{sd} + r_{ij}^{sd} - 2\hat{r}_{ij}^{sd} r_{ij}^{sd} \tag{24}$$

C. Dynamic Release of Optical Circuits

In a network deploying *dynamic* circuits, the model needs to consider scenarios where circuits may be *released* or switched off. To this end, the objective function is updated under the following assumptions: 1) Re-use of existing (already setup) dynamic circuits incurs no additional cost, 2) Introduction of new optical circuits has the same cost as defined in (1) and 3) decommissioning of an optical circuit is associated with a profit. The objective function under these assumptions while also considering only single circuit deployment ($N = 1$) was presented in [13] and we now present the generic objective function for any positive value of N.

We assume knowledge of already existing links with $\hat{X}_{xy} = 1$ indicating that one or more optical circuits already exist from v_x to v_y and $\hat{X}_{xy}^t$ indicating the number and type of these optical circuits. If no optical circuits exist previously ($\hat{X}_{xy} = 0$), the cost of optical circuits is similar to that in (1). In case one or more optical circuits exist ($\hat{X}_{xy} = 1$), we need to consider the difference between the initial ($\hat{X}_{xy}^t$) and final (X_{xy}^t) number of optical circuits of a specific type $t \in T$: if additional circuits are included ($X_{xy}^t > \hat{X}_{xy}^t$) then the cost associated should be computed as specified in (1), and a profit should be associated with the release of optical circuits ($X_{xy}^t < \hat{X}_{xy}^t$). In the model, the parameter SW_{xy}^t gives the profit for switching off a single optical circuit (type $t \in T$) from v_x to v_y. It is necessary to determine for scenarios with ($\hat{X}_{xy} = 1$), if the final configuration contains more or less optical circuits of the same type. We introduce a boolean variable Y^t_{xy} which is constrained in (25), (26). As the total number of circuits is bounded by N, this implies $-N \leq (X^t_{xy} - \hat{X}^t_{xy}) \leq N$. (25) ensures that if $(X^t_{xy} > \hat{X}^t_{xy})$, $Y^t_{xy} = 1$, while (26) ensures that $Y^t_{xy} = 0$ otherwise.

$$\forall v_x, v_y \in V, x \neq y \;:\; Y^t_{xy} \geq N^{-1} \cdot (X^t_{xy} - \hat{X}^t_{xy}) \quad (25)$$

$$\forall v_x, v_y \in V, x \neq y \;:\; Y^t_{xy} \leq 1 + N^{-1} \cdot (X^t_{xy} - \hat{X}^t_{xy}) \quad (26)$$

The objective function is now given by (27), where the first term gives the cost of adding new optical circuits when no circuits existed initially. In case circuit exist already, the second term incorporates cost for adding additional circuits while the third term incorporates profit for releasing existing optical circuits. The use of Y^t_{xy} ensures that only one of the second or the third term will be incorporated when computing for a specific circuit type between a pair of routers.

$$\begin{aligned} Min: \; & \sum_{xy}\sum_{t}(1 - \hat{X}_{xy}) \cdot X^t_{xy} \cdot Cost^t_{xy} \\ & + \sum_{xy}\sum_{t} \hat{X}_{xy} \cdot Y^t_{xy} \cdot (X^t_{xy} - \hat{X}^t_{xy}) \cdot Cost^t_{xy} \\ & - \sum_{xy}\sum_{t} \hat{X}_{xy} \cdot (1 - Y^t_{xy}) \cdot (\hat{X}^t_{xy} - X^t_{xy}) \cdot SW^t_{xy} \quad (27) \end{aligned}$$

(27) is non-linear as it contains the product $Y^t_{xy} \cdot X^t_{xy}$. We introduce an integer variable Z^t_{xy} $(0 \leq Z^t_{xy} \leq N)$ to substitute the term $Y^t_{xy} \cdot X^t_{xy}$ in (27) which is constrained using (28), (29), (30). As Y^t_{xy} is boolean, (28) constrains $Z^t_{xy} = 0$ if $Y^t_{xy} = 0$ while (29), (30) ensure $Z^t_{xy} = X^t_{xy}$ otherwise.

$$\forall v_x, v_y \in V, t \in T \;:\; Z^t_{xy} \leq N \cdot Y^t_{xy} \quad (28)$$

$$\forall v_x, v_y \in V, t \in T \;:\; Z^t_{xy} \leq X^t_{xy} \quad (29)$$

$$\forall v_x, v_y \in V, t \in T \;:\; Z^t_{xy} \geq X^t_{xy} - (1 - Y^t_{xy}) \cdot N \quad (30)$$

V. Model Adaptations driven by System Design

A. ER and SPF Routing

The actual system capabilities available in networks can significantly affect the applicability of the model presented. In explicit routing, for instance, we only enforce constraints for a single unique path, with routing decision made on the basis of the source and the destination address. As mentioned earlier, configuring routing in this fashion can only be achieved in MPLS networks. Alternately, in pure IP networks, static routing rules can be used to configure explicit routes, and to this end the model should incorporate constraints on destination-based forwarding, i.e., (19),(20),(21). This scenario is termed Explicit Routing with Destination Forwarding (ER-D).

While no system-specific constraints are enforced on SPF, incomplete modeling of SPF leads to significant overestimation of the performance, which is shown in the results. We note that the proposed formulation enforces conditions to mimic routing convergence as accurately as possible, but still, in case of a tie for route costs of a *shorter* route, the ILP will only choose one of the possible configurations and route preferences may need to be configured on some routers to ensure routing convergence to the configuration computed by the ILP. The exact configuration is vendor specific, and is out of the scope of this work. The proposed analysis can also be extended to support Equal Cost Multipath Routing (ECMP). Typical implementations of the ECMP in OSPF limit the maximum number of paths to be used by the protocol, and attempt to divide the flow equally between the same. In the model, this can be incorporated by including multiple routing variables based on the number of multiple routes supported, all constrained similarly. However, given that there is inconsistency in how different vendors manage traffic distribution, especially in the presence of non-IP MPLS traffic, such as PseudoWire services, ECMP is typically not used in core networks [22] and is not studied in this paper.

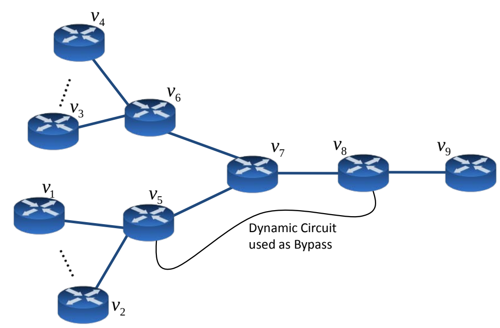

Fig. 2. Optical circuit introduced in IP networks as optical bypass.

B. "Hiding" Optical Circuits from IP Routing

To limit IP routing changes, new circuits may not be advertised in IP routing, i.e., can be "hidden"; we refer to such a new link as Optical Bypass (BY). This can be implemented by either using a very high link weight or by using IP addresses for the link endpoints that are outside the IP routing subnet. IP traffic can be offloaded onto bypass by configuring specific forwarding rules, such as static routing, at the ingress router. The bypass mechanism, as described in [14], [16], [13], is most effective when the IP traffic can only be *offloaded* onto a bypass if the ingress and egress routers of the bypass lie on the original routing path of the offloaded IP traffic.

An example can be seen in Fig. 2, where a bypass from v_5 to v_8 is used to offload traffic from v_5 and routers upstream from v_5 to v_8 and routers downstream from v_8. In order to offload traffic on a bypass from v_x to v_y it is necessary to know if the original routing path from a source v_s to destination v_d traverses the segment from v_x to v_y. We use two boolean routing constants ψ^{sd}_{xy} and $\psi^{sd}_{xy}(ij)$ along with original IP routing protocol configuration $\hat{r}^{sd}_{ij}$ to constrain the same. ψ^{sd}_{xy} indicates if the original routing path from v_s to v_d traverses the segment from v_x to v_y, while $\psi^{sd}_{xy}(ij)(\leq \psi^{sd}_{xy})$ identifies links (v_i to v_j) that are included in the segment from v_x to v_y. For example, in Fig. 2, $\psi^{29}_{58} = 1$ while $\psi^{49}_{58} = 0$ as the routing path of the former goes over the segment from v_5 to v_8 but not for the latter, and the parameters $\psi^{29}_{58}(57) = \psi^{29}_{58}(78) = 1$ indicate that the link $v_5 - v_7$ and $v_7 - v_8$ are traversed by the route from v_2 to v_9 over the segment from v_5 to v_8.

Formulations for Optical Bypasses (BY), as presented in [13], [16] choose traffic to be offloaded based on the destination as well as source address, which is typical in scientific networks [9]. However, such routing rules have high lookup complexity and are not applied in core networks. We present a formulation where only destination-based forwarding rules typical to IP routers are used, and show the difference between the two approaches in our results. The formulation uses the objective function (1) and is subject to the same network configuration constraints presented in Sec. IV-A. A boolean variable f^d_{xy} indicates if traffic to a destination v_d at bypass ingress v_x and all routers upstream from v_x is offloaded over a circuit from v_x to v_y. (31) ensures that traffic is only offloaded on a circuit if it exists while (32) ensures that traffic from v_x and all routers upstream from v_x to v_d can only use a circuit from v_x to v_y, if the egress v_y lies on the original routing path using the constant ψ^{xd}_{xy}. The number of times a specific route is may be offloaded is not constrained, but the formulation ensures that no two *overlapping* optical circuits are used to offload the same traffic with (33), where $\psi^{sd}_{xy}(ij)$ is used to ensure that for a specific route, no single link is bypassed by two distinct optical circuits. Here, f^d_{xy} indicates if traffic is offloaded from v_x to v_y while $\psi^{sd}_{xy}(ij)$ indicates if the segment traverses the link e_{ij} and the sum over all possible bypasses ensures that a route can only use one optical circuit that bypasses a specific link. For example, in Fig. 2, traffic from v_5 to v_9 cannot simultaneously use an optical circuit from v_5 to v_8 and from v_7 to v_9, as the link from v_7 to v_8 would be bypassed twice.

$$\forall v_d, v_x, v_y \in V,\ \hat{L}_{xy} = 0 : f^d_{xy} \leq X_{xy} \quad (31)$$

$$\forall v_d, v_x, v_y \in V,\ \hat{L}_{xy} = 0 : f^d_{xy} \leq \psi^{xd}_{xy} \quad (32)$$

$$\forall v_s, v_d, v_i, v_j \in V,\ \hat{L}_{ij} = 1\ :\ \sum_{xy:\hat{L}_{xy}=0} \psi^{sd}_{xy}(ij) \cdot f^d_{xy} \leq 1 \quad (33)$$

Individual constraints are defined for link capacity utilization on existing IP links and on optical bypasses. (34) constrains the traffic on an existing IP link and uses the boolean expression $(1 - \sum_{xy:\hat{L}_{xy}=0} \psi^{sd}_{xy}(ij) \cdot f^d_{xy})$ (33) to ensure that traffic from v_s to v_d is not included in the computation of load on link e_{ij} if the traffic was offloaded from this link. The traffic on the optical bypass from v_x to v_y is computed using f^d_{xy} as presented in (35).

$$\forall e_{ij} \in E,\ \hat{L}_{ij} = 1\ :\ \sum_{sd} \hat{\lambda}_{sd} \cdot \hat{r}^{sd}_{ij} \left(1 - \sum_{xy:\hat{L}_{xy}=0} \psi^{sd}_{xy}(ij) \cdot f^d_{xy}\right) \leq \alpha \cdot C_{ij} \quad (34)$$

$$\forall v_x, v_y \in V,\ \hat{L}_{xy} = 0 : \sum_{sd} \hat{\lambda}_{sd} \cdot \psi^{sd}_{xy} \cdot f^d_{xy} \leq \alpha \cdot C_{xy} \quad (35)$$

C. Unknown Traffic Matrix Information

Unknown traffic matrix presents a significant challenge, as all capacity constraints presented till now become non-linear in the absence of traffic information. However, if routing changes in the network are limited, it is possible to use traffic measurements as a substitute for actual traffic values. Unlike for the ER scheme, this idea can be applied effectively to the SPF and the BY schemes that affect fewer routes. In the model, traffic from ingress router v_s to egress router v_d is represented as the *variable* λ_{sd}, and two types of measurements, namely IP link loads $LinkLoad_{ij}$ and the virtual output queue γ^k_{ij} are used which are presented with an example in Fig. 3. $LinkLoad_{12}$ is the traffic measured on the directed link e_{12}, while γ^2_{13} is the traffic measured on v_2 that arrives from v_1 and is forwarded to v_3. There are two special cases for the measurement of γ^i_{jk}, with γ^1_{12} indicating the traffic entering the network from the router v_1 with the next hop as v_2 and γ^3_{23} indicating the traffic exiting the network at router v_3 with the previous hop as router v_2. In commercial routers, the link load measurements are easily available via the Simple Network Management protocol (SNMP), while measurements for virtual output queues can be obtained either from the underlying switch fabric or by configuring custom ingress-interface based queuing [23] on routers and measuring the traffic in these queues via SNMP. These measurements can be expressed as a sum of traffic variables λ_{sd} , i.e.,

$$\forall e_{ij} \in E\ :\ \sum_{sd} \lambda_{sd} \cdot \hat{r}^{sd}_{ij} = LinkLoad_{ij} \quad (36)$$

$$\forall e_{ij}, e_{jk} \in E, i \neq j \neq k\ :\ \sum_{sd} \lambda_{sd} \cdot \hat{r}^{sd}_{ij} \cdot \hat{r}^{sd}_{jk} = \gamma^j_{ik} \quad (37)$$

$$\forall e_{ij} \in E\ :\ \sum_{d} \lambda_{id} \cdot \hat{r}^{id}_{ij} = \gamma^i_{ij} \quad (38)$$

$$\forall e_{ij} \in E\ :\ \sum_{s} \lambda_{sj} \cdot \hat{r}^{sj}_{ij} = \gamma^j_{ij} \quad (39)$$

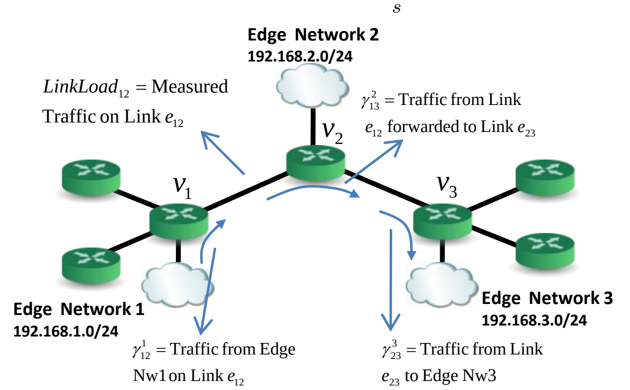

Fig. 3. Traffic measurements available in IP networks

As there are no routing loops in the network, the coefficient of λ_{sd} in the expressions for traffic are binary and we define a set of expressions D generated using available measurements, and each $D_i \in D$ consists of a set of binary routing parameters d^i_{sd} and measured traffic B_i. The relation between these parameters can be expressed as:

$$D_i\ :\ \sum_{sd} d^i_{sd} \cdot \lambda_{sd} = B_i \quad (40)$$

The expressions in set D are unique, i.e. if for any $D_i,\ D_j \in D$, if all $d^i_{sd} = d^j_{sd}$, then i must be equal to j, or that no two expressions in D have same coefficients for all λ_{sd}. For example, in the network in Fig. 3, an expression D_1 generated using the measurement for $LinkLoad_{12}$ $(= \lambda_{12} + \lambda_{13})$ would have $d^1_{12} = d^1_{13} = 1$, $B_1 = LinkLoad_{12}$ and all other $d^1_{xy} = 0$. Also, as expressions should be unique, if the set D already contains the expression for $LinkLoad_{12}$, the expression for γ^1_{12} (which is also given by the sum of λ_{12} and λ_{13}) would not be included as a new expression in the set D.

In the ILP, under known traffic matrix conditions, all capacity utilization constraints (14)(35)(34) are of the form

$$\forall v_i, v_j \in V, L_{ij} = 1 : \sum_{sd} a^{sd}_{ij} \cdot \lambda_{sd} \leq \alpha \cdot C_{ij} \tag{41}$$

where a^{sd}_{ij} is a boolean variable indicating if traffic from v_s to v_d uses the link from v_i to v_j. We propose to generate constraints to check if the link load expression in (41) matches a known expression $D_i \in D$. In case a match is found, we substitute the value of the expression $\sum_{sd} a^{sd}_{ij}\lambda_{sd}$ with B_i. We also compute upper bounds on each λ_{sd} using the known expressions in D which are represented as λ^{max}_{sd}. In case no expression in D matches the link load expression, we substitute λ_{sd} with λ^{max}_{sd} to linearize (41). For an IP link from v_i to v_j, we use a binary variable S^x_{ij} to indicate if the link load expression on link from v_i to v_j is the same as the expression $D_x \in D$ using (42) and (43). (42) ensures that $S^x_{ij} = 1$ if a coefficient for any λ_{sd} in D_x does not match between the sum of traffic routed on a link from v_i to v_j (a^{sd}_{ij}) while (43) ensures that $S^x_{ij} = 0$ otherwise. Both constraints are non-linear, and we use the fact that both a^{sd}_{ij} and d^x_{sd} are binary and consequently $|a^{sd}_{ij} - d^x_{sd}|$ is also binary to linearize them as (44) and (45) respectively.

$$\forall D_x \in D, v_i, v_j \in V$$

$$\forall v_s, v_d \in V \ : \ S^x_{ij} \geq |a^{sd}_{ij} - d^x_{sd}| \tag{42}$$

$$S^x_{ij} \leq \sum_{sd} |a^{sd}_{ij} - d^x_{sd}| \tag{43}$$

$$\forall v_s, v_d \in V \ : \ S^x_{ij} \geq a^{sd}_{ij} + d^x_{sd} - 2 \cdot d^x_{sd} \cdot a^{sd}_{ij} \tag{44}$$

$$S^x_{ij} \leq \sum_{sd} \left(a^{sd}_{ij} + d^x_{sd} - 2 \cdot d^x_{sd} \cdot a^{sd}_{ij} \right) \tag{45}$$

S^x_{ij} identifies if a traffic expression in D matches the link load expression, and the binary variable S_{ij} defined in (46) is used to identify if none of the expressions in D match the link load expression for link from v_i to v_j. As all expressions in D are unique, there can be a maximum of one expression that will match the routing on the link from v_i to v_j so the parameter S_{ij} will be 0 in case no expressions match ($S^x_{ij} = 1 \ \forall x$) and will be 1 in case an expression matches the traffic on the link.

$$\forall v_i, v_j \in V \ : \ S_{ij} = \sum_x \left(1 - S^x_{ij}\right) \tag{46}$$

Using these variables, the constraint for traffic on an IP link (as presented in (14)) can now be represented as:

$$\forall v_i, v_j \in V, i \neq j \ : \ \sum_{sd} a^{sd}_{ij} \cdot \lambda^{max}_{sd} - S_{ij} \cdot \lambda_\infty \leq \alpha \cdot C_{ij} \tag{47}$$

$$\forall v_i, v_j \in V, i \neq j \ : \ \sum_x (1 - S^x_{ij}) \cdot B_x \leq \alpha \cdot C_{ij} \tag{48}$$

Here, λ_∞ is a very large positive constant used to ensure that (47) is redundant if the traffic on a link matches an expression in D ($S_{ij} = 1$). In this scenario, (48) will substitute the value of an expression in D which matches the traffic scenario to enforce the traffic constraint. In case none of the expressions in D match the traffic expression, ($S_{ij} = 0$) (47) uses the computed bounds on individual traffic values instead to enforce constraints on IP link capacity. Note that in this scenario, (48) is redundant as all $S^x_{ij} = 1$.

As an example, consider the network in Fig. 3, where we assume shortest path routing. If we consider the traffic on link e_{12} and assume that the final solution does not contain any new links from v_1, the routing variables $r^{12}_{12} = r^{13}_{12} = 1$ while all other $r^{sd}_{12} = 0$. This matches the expression generated using the measurement $LinkLoad_{12}$ (D_1) where $d^1_{12} = d^1_{13} = 1$ and all other $d^1_{xy} = 0$. Here, the coefficient $S^1_{12} = 0$, while all other $S^x_{12} = 1$ and $S_{12} = 1$. Therefore, (47) becomes redundant, while (48) uses the measurement $B_1 = LinkLoad_{12}$ to constrain the capacity C_{12}. On the other hand, if for a particular configuration an expression for the link load is not found in D, i.e. $S_{12} = 0$ and all $S^x_{12} = 1$, (48) becomes redundant, while the constraint in (47) uses the measured traffic bounds to constrain the capacity C_{12}.

To adapt this approach to the proposed mechanisms SPF and BY, we map the coefficients in the corresponding capacity constraints to the variable a^{sd}_{ij}. In the case of the SPF, we substitute a^{sd}_{ij} with r^{sd}_{ij} to generate capacity utilization constraints while for BY the values of a^{sd}_{ij} when applied to an existing IP link ($\hat{L}_{ij} = 1$) is given by $\hat{r}^{sd}_{ij}\left(1 - \sum_{xy:\hat{L}_{xy}=0} \psi^{sd}_{xy}(ij) \cdot f^d_{xy}\right)$ and in case of an optical bypass is given by $f^d_{xy} \cdot \psi^{sd}_{xy}$.

Note that the mechanism used to populate the traffic expression set D is critical to the quality of the solutions under unknown traffic matrix conditions. Just the use of all possible link load measurements guarantees a solution where optical circuits are only used to increase the capacity of an existing link, as in this scenario, routing in the network does not change, and for the same routing expression, there must be a unique match for the traffic expression on every link in the expression set D. In the simulation study, we include expressions for the IP link loads $LinkLoad_{ij}$ the expressions for the virtual output queue measurements γ^j_{ij} and traffic expressions obtained by subtracting the traffic measured on the virtual output queues from the link loads, i.e. $LinkLoad_{ij} - \gamma^j_{ik}$ and $LinkLoad_{jk} - \gamma^j_{ik}$. The last expressions can be used to compute the accurate traffic (in many cases) for two-hop optical circuits under the different mechanisms. For example, for the topology in Fig. 3, if we include three more expressions, given by γ_{13}^2, $LinkLoad_{12} - \gamma_{13}^2$ and $LinkLoad_{23} - \gamma_{13}^2$, the model can compute the optimal solution where a new circuit is established from v_1 to v_3. Here, traffic from v_1 to v_3 is shifted onto the new link from v_1 to v_3, and this traffic is eliminated from the links e_{12} and e_{23}.

VI. Performance Evaluation

We evaluate the performance of the explicit routing with and without destination forwarding (ER and ER-D), SPF and optical bypass schemes with destination forwarding (BY-D). In addition, we present results from BY formulations (BY) from our paper [13] to show the effect of the extensions proposed here. Results also include the computations under unknown traffic matrix (SPF-NoTM, BY-D-NoTM, BY-NoTM). Effectiveness of the different schemes is measured as the installed optical capacity ($\sum X_{xy}^t \cdot C_t^O$) and the number of routes (and consequently network services) affected. The Gurobi Optimizer v5.0 was used to solve the ILP and computations under 30 minutes were solved to optimality (MIP gap $1e^{-4}$) while longer simulations were terminated with a MIP gap of 0.01 (1%). The solver used all processors on a PC with 4 core Intel i-5 processor (2.60 GHz) and 4 GB RAM. The results were averaged over approximately 400 iterations with 95% of confidence intervals for all results presented.

The study assumes three circuit granularities with capacities 2.5G, 10G, and 40G. Optical circuit cost $Cost_{xy}^t$ for the objective function is computed using normalized interface cost $IFCost^t$ with values (10, 30, 90) and the $CostPerHop^t$ with values (1, 3, 9) for each granularity. Interface costs have been modeled based on list prices for long reach optical interfaces for commercial routers, where a common trend indicates that the cost of high-capacity, e.g., 40G interface, is typically 3-4 times the cost of a 10G interface [24]. As we are primarily concerned with the optical circuit placement, the cost of the bandwidth $CostPerHop^t$ is assumed to be a factor lower than $IFCost^t$, thereby making the number of interfaces the deciding factor, but preferring shorter optical circuits to longer ones. The hop count of the shortest path in the original IP network $HopCount_{xy}$ is used as the measure of distance in the objective, and consequently $Cost_{xy}^t$ is given by

$$Cost_{xy}^t = IFCost^t + HopCount_{xy} * CostPerHop^t \quad (49)$$

If $CostPerHop^t = 0$, the total installed capacity is not affected but the location and placement of circuits, especially in the case of ER and ER-D can be different. At the same time, the location and placement are unlikely to be affected in SPF and BY mechanisms due to strict constraints on routing. Also, if higher capacity interfaces were to become dramatically cheaper (1-2 times the cost of smaller interfaces), the optimization may compute solutions with fewer higher capacity circuits. However, for small changes in price, the total installed capacity remains unaffected.

The performance study is divided into three parts: The first study evaluates the performance under different traffic conditions, while the second study analyzes the effect of network topologies. The final study evaluates the dynamic circuit release over time. We use the NSFNet topology (14 routers, 40 directed links) as the base topology with a representative (scaled) traffic as provided in [25]. Initial traffic for each iteration is computed by varying the traffic for each s-d pair in the representative traffic matrix with a uniform random scaling factor between [0.9, 1.1]. In the second study, we use a ring (14 routers, 28 directed links) and a 4x4 grid topology (16 routers, 48 directed links). Here, initial traffic for all s-d pairs is generated using a uniform random distribution between [0.4,1.2] Gbps. Initial link capacities and routing are computed using the SPF model, where all $\hat{L}_{ij} = 0$ (no links) and the parameter L_{ij} is constrained to be 1 if a link exists in the base topology used, and 0 otherwise. The computed solution only contains edges as seen in the base topology, along with the optimal capacity distribution and the routing configuration. The maximum link utilization threshold $\alpha = 0.7$ and the link weight $\hat{w}_{ij} = 1$, which implies that SPF uses the shortest hop path for all computations. In the results, unless specified otherwise, the number of circuits that can be aggregated between the same pair of routers (N) is set to 1.

A. *Performance under Different Traffic Conditions*

In this study, we apply two different traffic overloading scenarios, the first one illustrative of temporary traffic churns affecting a relatively small number of routes, while the second one illustrative of a more uniform traffic increase. Thus, first scenario randomly selects a small number of router pairs and increases the traffic between them by 150%; the degree of overloading is differentiated by changing the number of router pairs selected. The second scenario increases the traffic on all routers, and the percentage increase of traffic on all IP routes is varied to generate different overloading conditions.

Fig. 4(a) presents the variation in the average optical circuit capacity required to overcome overloading. The ER mechanism, with fully flexible IP routing has the lowest requirements for additional capacity followed by ER-D. Under known traffic matrix conditions, the required capacities for SPF, BY and BY-D in this topology are comparable. BY-D exhibits an almost insignificant increase in capacity when compared to BY, and as it only uses destination-based static routing rules, is preferable for application in core IP networks. The ability of the bypass schemes to *selectively* divert IP routes over optical circuits is indicated in the slightly higher requirement on capacity by the SPF mechanism.

We also observe that the methods BY-NoTM and BY-D-NoTM have a significantly better performance than SPF-NoTM. This is due to the mechanism used for generating the expressions in D for computation in unknown traffic matrix conditions. The combination of traffic expressions γ_{ik}^j, $LinkLoad_{ij} - \gamma_{ik}^j$ and $LinkLoad_{jk} - \gamma_{ik}^j$ when used with optical bypasses ensures that an expression match is always found for traffic on existing links as well as bypasses when an optical bypass created from v_i to v_k (2 hop bypass). The choice of the expressions used also leads to the near equal capacity requirements for BY-NoTM and BY-D-NoTM, as even though the former can offload traffic from individual s-d

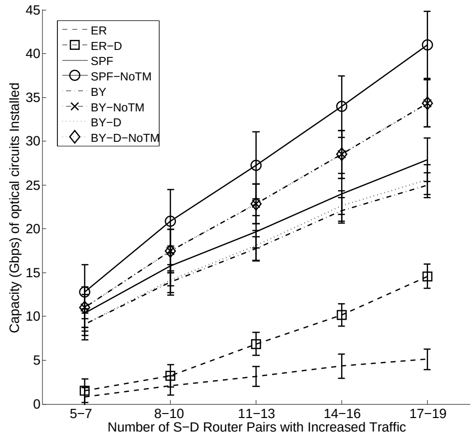

(a) Optical circuit capacity required in Scenario 1

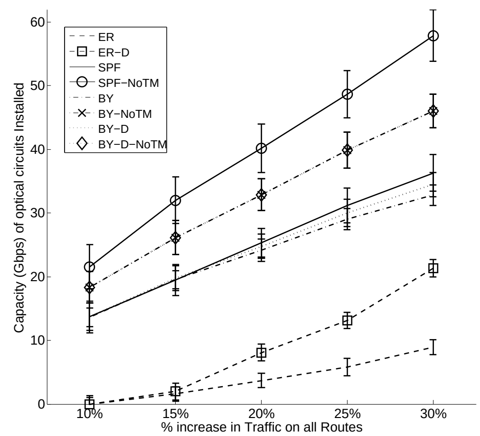

(b) Optical circuit capacity required in Scenario 2

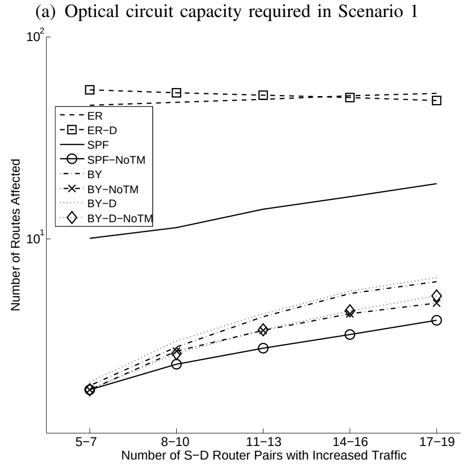

(c) Number of routes affected in Scenario 1

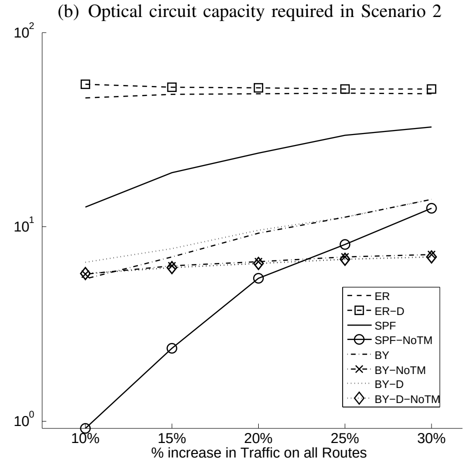

(d) Number of routes affected in Scenario 2

Fig. 4. Performance for the NSFNet for two scenarios: (Scenario 1) Traffic on Randomly Selected Router Pairs is increased by 150%; (Scenario 2) Traffic on all Router Pairs is increased by a factor as indicated (As confidence intervals on the results for routing changes were small, the same have not been shown in the figures for clarity.)

pairs, the expressions available in D provide optimal solutions when aggregate traffic flows are offloaded. In case of SPF, a new link established from v_i to v_k can cause routing changes where routes not traversing the links e_{ij} and e_{jk} may also use this new link, and consequently a matching expression may not be found for the new link in D, leading to optical circuit capacity. A similar linear trend is observed for the Scenario 2 in Fig. 4(b), albeit at a larger scale, indicating that the relative performance of the different schemes remain independent of the overloading mechanism used.

The large difference between the ER and the SPF schemes indicates the importance of accurately modeling IP routing. We provide an example to quantify the effects of these parameters: We use the traffic scenario where traffic of 17-19 randomly selected source-destination router pairs is increased by 150%, and compute the required optical capacity when only enforcing a) the ER constraints, b) constraints on routing to use the shortest path (SP) (17)(18), c) constraints on routing to use shortest path as well as destination based forwarding (19)(20)(21) (SPDF) and d) constraints on re-convergence (SPF). From Table I, we can see that all the three constraints have a significant affect on the required capacity demand; just the inclusion of the constraint on routing re-convergence can lead to an increase in capacity demand by a factor of 2.6, which is significant.

Fig. 4(c) and Fig. 4(d) show the number of routes affected for different schemes and overloading conditions. BY schemes affect a very small number of routes, which tend to stay almost constant with varying network load. The SPF mechanism affects a comparatively larger number of routes, and the number of routes affected increases with the overloading increase. This is expected, as constraints on shortest path routing and destination-based forwarding lead to a significant number of routing changes. Under unknown traffic matrix conditions, the SPF mechanism tends to increase the existing link capacity as the traffic on new links cannot be computed accurately in all instances, leading to reduced routing changes. Finally, the ER mechanism, as expected, affects a large number of routes, as it typically traffic engineers the networks, with ER-D affecting a comparably larger number of routes. The large number of affected routes is one of the reasons that explicitly configured paths are rarely used for all traffic in commercial IP networks.

The number of routing changes observed in BY and SPF mechanisms is linked to placement of optical circuits. We study placement of optical circuits in the scenario where 17-19 s-d router pairs have increased traffic. As shown in

TABLE I
OPTICAL CAPACITY REQUIRED DEPENDING ON IP ROUTING CONSTRAINTS CONSIDERED IN DIFFERENT MECHANISMS

ER	SP	SPDF	SPF
5.61	14.63	16.79	26.12

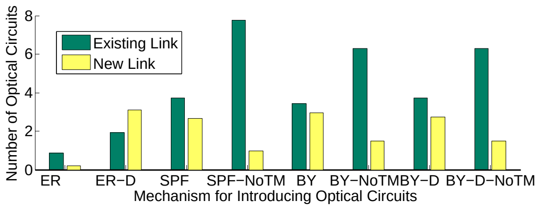

Fig. 5. Optical circuits as new links or capacity boost for existing links when traffic of 17-19 s-d pairs in NSFNet is increased by 150%

Fig. 5, ER has the lowest number of circuits used as new links, since the cost of a multi-hop (longer) optical circuit is higher than the cost of a single hop (shorter) circuits. ER-D uses comparatively larger number of circuits, with a high percentage of circuits used as new links, which combined with the low capacity requirement highlights the routing flexibility available in ER-D as compared to other schemes. The SPF, BY and BY-D mechanisms again have comparable distribution of optical circuits used as new links versus optical circuits used on existing links. We also see that the SPF-NoTM cannot estimate traffic on multi-hop (new links) effectively, and as a result mostly uses circuits to boost the capacity of existing IP links. BY-NoTM and BY-D-NoTM demonstrate a reduced tendency to create new links as compared to BY and BY-D, but make more new links as compared to SPF-NoTM, with lower overall capacity requirements.

We also study the difference in the installed capacity when multiple optical circuits can be aggregated to form a single link. As seen in Fig. 6, the aggregation of 4 circuits ($N = 4$) reduces the overall installed capacity in each case, but the difference under optimal traffic matrix conditions is minimal due to the prevalent use of low capacity interfaces at N=1. The solutions for unknown traffic matrix use more high-capacity interfaces and as a result, the aggregation of multiple lower capacity optical circuits can better match the capacity demand, leading to a decrease of almost 5 Gbps as seen in the case of SPF-NoTM. The difference in the case of BY-NoTM and BY-D-NoTM is less pronounced but still significant.

B. Performance in Different Topologies

For the ring topology (Fig. 7(a)), we observe that all mechanisms demonstrate an almost linearly increasing optical capacity requirement with the increase in the number of overloaded s-d pairs. As expected, ER and ER-D have

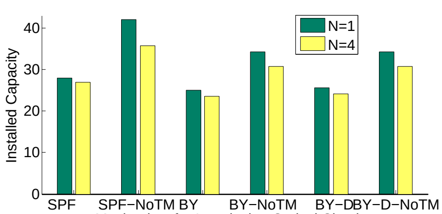

Fig. 6. Difference in installed capacity with larger N when traffic on 17-19 s-d pairs in NSFNet is increased by 150%

the lowest capacity requirement, and the difference between the two is less pronounced than in the NSFNet topology. In a ring topology, i.e., with a low degree of connectivity, the introduction of a new multi-hop circuit can significantly disrupt the routing under shortest path constraints, and as a result, large and unwanted volumes of traffic are diverted onto the new multi-hop link. Since the BY based schemes can control the forwarding of traffic flows onto new multi-hop optical circuits, they require a significantly lower capacity as compared to the SPF scheme. The BY schemes under unknown traffic matrix conditions do not perform very well in sparse topologies as they can only guarantee expressions in D for two hop bypasses. As a result, the difference between the optimal BY schemes and the schemes under unknown traffic matrix in the ring topology is greater than what was observed in the NSFNet topology.

An interesting observation here is that BY solutions under unknown traffic conditions are still better than the optimal SPF solution, while the SPF-NoTM solution uses significantly larger resources. This gap may reduce for smaller size rings but for large ring-like topologies, the constraint on shortest path routing significantly hampers the performance of SPF. The trend for affected routes as seen in Fig. 7(c) also highlights the difference in the performance of the SPF and the BY schemes, with the latter exhibiting routing changes which are almost a factor lower than those exhibited by the SPF schemes. The BY-NoTM and the By-D-NoTM schemes also exhibit significantly larger routing changes as compared to the optimal schemes, but the routing changes are still lower than those of SPF. The ER and ER-D schemes affect an almost constant number of routes over the different loading conditions and the SPF-NoTM scheme, being the closest to the real world scenario, has the worst performance, especially for high overloading.

The performance of the mechanisms in the mesh topology also follows a similar pattern as seen in Fig. 7(b). As the initial topology provides a higher path choice, the total average capacity installed on each link is lower and as a result the same traffic scenario can lead to multiple overloading conditions in the network. As a result, when compared to the ring topology, the ER, ER-D and the optimal BY based schemes have larger capacity demands in the mesh topology. In comparison, the SPF scheme performs marginally better in the mesh topology, as even though the topology is symmetric, the number of routes that are affected by a new link are smaller in the mesh topology than in the ring topology. The BY schemes under unknown traffic conditions are still seen to perform better than the optimal SPF scheme with fewer routing changes, as both prefer the use of 1 or 2 hop circuits in most cases but the BY mechanisms can better control routing changes. In the scenario where traffic on 17-19 router pairs is increased by 150%, the use of ≥ 3 hop circuits by SPF (25 %) BY-NoTM (10%) and BY-D-NoTM (10%) is lower than BY (30%) and BY-D(29%), while ER and ER-D only employ 3 or more hop circuits. As the overloading increases, the former use more circuits in the network which leads to increase in the required capacity. Compared to the BY schemes, SPF employs fewer

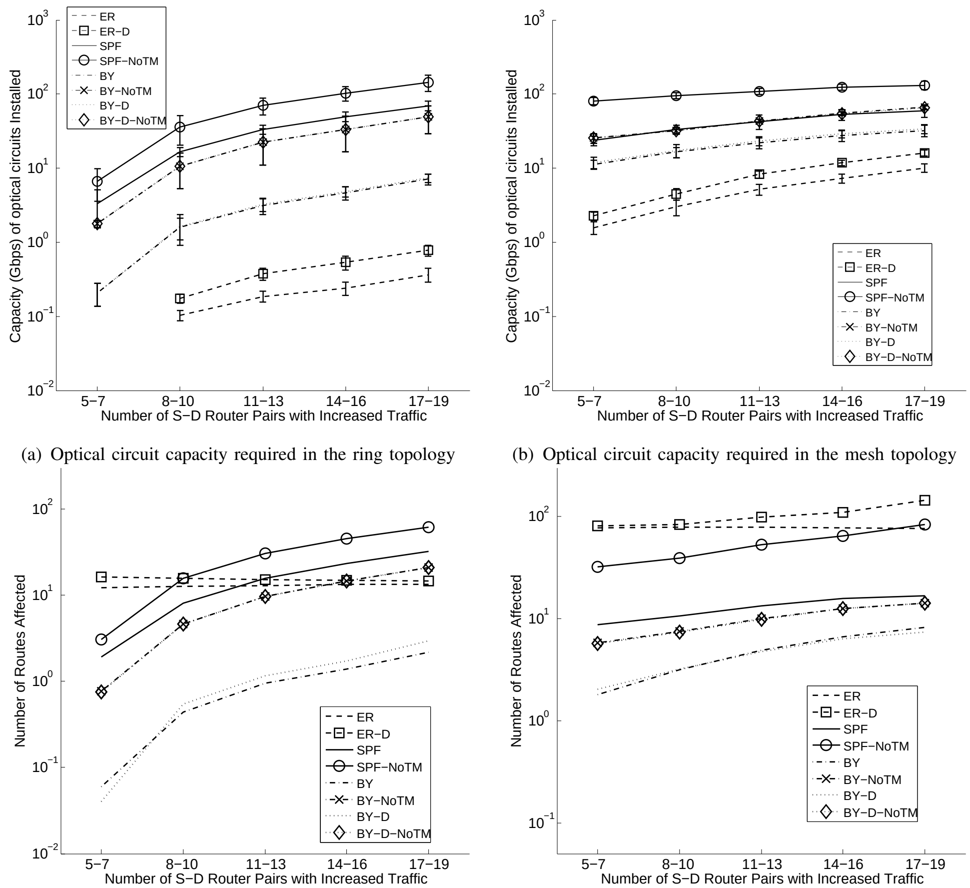

(a) Optical circuit capacity required in the ring topology

(b) Optical circuit capacity required in the mesh topology

(c) Number of routes affected in the ring topology

(d) Number of routes affected in the mesh topology

Fig. 7. Performance parameters for the 14 node Ring and the 4x4 grid (mesh) topology for the different mechanisms when traffic overloading is achieved by increasing the traffic on randomly selected source-destination pairs (as indicated on the X axis) by 150 %

circuits, but constraint on shortest path routing leads to more routing changes which drive more traffic than desired onto the new circuits increasing the demands for optical capacity.

As seen here, the SPF protocols are ill suited to highly regular topologies, and the performance of SPF is much closer to the BY schemes in irregular topologies where fewer routes are affected by the introduction of a new link. The result also highlights the importance of the routing re-convergence constraint (22): if not enforced, the computed performance is significantly better than the BY schemes, as multiple alternate (equal-cost) routes exist in the mesh network and the ILP model can re-route traffic from a route over overloaded links to an alternate route, which does not commonly happens according to the current routing protocols.

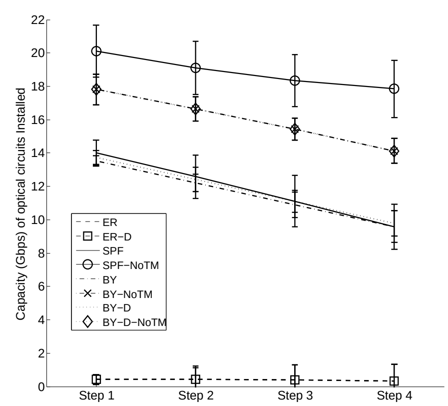

Fig. 8. Average Installed Capacity at each Step for Switch Off in the NSFNet topology.

C. Effects of Dynamic Optical Circuit Release

In this study, we quantify the capacity and the configurations required when addressing a transient traffic increase with a subsequent circuit release. We use the NSFNet configuration for the topology and initial traffic. In the first step, traffic on 18 source-destination pars is increased by 120% and in the next 3 steps (step 2-4), we select 3 of the router pairs selected in step 1 and reset the traffic on these pairs to the initial value. The scenario mimics a temporary traffic increase which reverts back to the original traffic matrix over time. We use (27) as the objective function, while $SW^t_{xy} = 0.7 \cdot Cost^t_{xy}$.

As expected, the required capacity at each step is lower or equal to the previous step. The required capacity is the lowest for ER and ER-D, and does not change significantly with each step. SPF, BY and BY-D display comparable performance with a gradual reduction in the installed capacity. The BY-NoTM and the BY-D-NoTM have exactly the same performance over the different steps, while the SPF-NoTM is the worst

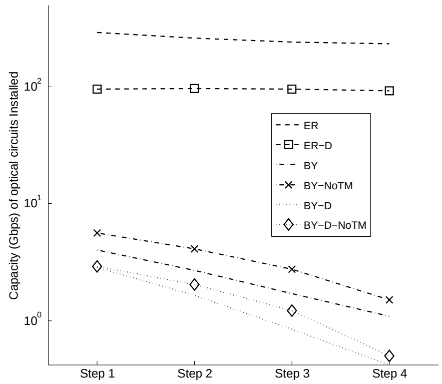

Fig. 9. Average routing configurations at each step for switch-off in the NSFNet topology.

performer. For this study, the generation of D is updated to include scenarios where a temporary link may be switched-off. Therefore, for a temporary optical link from v_x to v_y, we also include the expression $LinkLoad_{ij} + LinkLoad_{xy}$ for all e_{ij} where $\psi_{ij}^{xy} = 1$ in the original routing topology, and hence if a circuit from v_x to v_y is decommissioned, traffic is likely to go back on e_{ij}.

A significant factor in such scenarios is the number of configurations at each step. For ER, each r_{ij}^{sd} that is changed requires a configuration and the total routing configurations are given by $\sum_{sd} |r_{ij}^{sd} - \hat{r}_{ij}^{sd}|$. For ER-D, the total number of configurations would include all changes to the forwarding rules for a unique destination, i.e. all changes to the variable $FT_i^d(j)$. For BY, a configuration is made for each traffic flow at each bypass ingress used by the flow the total configurations would be given by the difference in the offloading variable f for BY and by $f_{xy}^d \cdot \psi_{xy}^{sd}$ for the BY-D mechanism.

The routing configurations at each step are shown in Fig. 9, and provide a measure of the configuration overhead involved. The routing flexibility of the ER and ER-D schemes also results in a very high configuration overhead. The configuration effort in ER-D is seen to be lower, even when the total number of routes affected are higher. This is due to the fact that routing rules need to be set up for every s-d pair in ER, while ER-D only configures rules based on the destination of traffic on each router. The configuration effort for all BY based schemes including the schemes with known and unknown traffic matrix is < 10, with the destination-based forwarding schemes demonstrating a marginally lower configuration effort as compared to the BY and BY-NoTM schemes. However, the significantly lower configurations required for all steps make the BY based schemes comparably favorable for applications in IP networks with dynamic traffic fluctuations.

D. Computation Complexity

We illustrate the measured computation times (95% confidence) in Table II. In the ER and SPF schemes, the routing variable r_{ij}^{sd} contributes primarily to the complexity with $O(|V|^4)$ variables ($|V|$ number of routers in the network). Unlike traditional routing problems, where the edges in the network are fixed, the model here can include multiple new links and hence all routing variables must be considered. The ER and ER-D model exhibit the highest time complexity due to the large flexibility in routing, which is indicated by the high average computation times for both these models. The SPF model also has a high complexity: variants of the unsplittable shortest path routing problem have been shown to be NP-hard in fixed network topologies [26], and the use of optical circuits as new links further increases the complexity. As a result, while the average computation time is around 270 seconds (N=4), the worst case time observed here was $\tilde{6}500$ seconds.

TABLE II
COMPUTATION TIMES (IN SECONDS) IN THE NSFNET TOPOLOGY

N	ER	ER-D	SPF	SPF-NoTM
1	1471.68 ± 83.52	945 ± 80.23	190.65 ± 45.92	61.35±8.53
4	-	-	270.71± 64.86	83.42±10.27
N	BY	BY-NoTM	BY-D	BY-D-NoTM
1	0.14 ± 0.01	5.27 ± 0.03	0.05 ± 0.01	4.41 ± 0.07
4	0.22 ± 0.04	7.34 ± 0.05	0.09 ± 0.01	6.82 ± 0.13

The BY models are significantly simpler as offloading is constrained by the original routing in the network. The variable count for f_{xy}^d is of the order of $O(|V|^2 \cdot H)$ where H is the average hop count in the routing configuration, which is significantly lower than SPF. Even here, if a bypass spans links that are not overloaded, it is never utilized in the optimal solution which further reduces the complexity. The same is reflected in the average computation times, which are significantly lower for BY and BY-D. Also, the variation in the computation times for SPF is significantly large, while the same for bypass was very small, indicating that the latter is better suited for computations in real-time. The complexity of computation under unknown traffic matrix conditions is affected by the size of set D which determines the number of additional variables used in the model. Here, the choice of circuits that can be used is severely limited based on the available measurements in D, and as a result the computation time for SPF-NoTM is significantly lower than for SPF.

The SPF scheme does not scale well to large topologies, as the introduction of a new link can cause large-scale routing changes. Network operators typically introduce routing areas with a relatively small number of routers within the core network in order to manage the complexity of the same. The BY mechanisms have a very small computation time and can be used as is for moderate size topologies. Also, a large network graph could be cut along the edges into multiple small subgraphs, with the constraint that a route can only enter/exit a subgraph once. The BY mechanism could then be applied to different sub-graphs in parallel as any bypass established within a sub-graph will not affect the traffic condition beyond this subgraph, while overloading on links between the subgraphs is solved by increasing their capacity. The computation under unknown traffic matrix conditions does not scale well due to the significantly higher number of variables involved, and efficient splitting mechanisms must be developed to employ the same in large topologies.

VII. Conclusion

We analyzed the effects of dynamic optical circuit setup on IP routing and presented models for computing the size and placement of optical circuits. Our results show for the first time that inaccurate modeling of essential IP routing protocol parameters leads to significant under-estimation of required optical capacity. We showed that explicit routing, focus of most past work, affects a comparatively large number of routes and network services. In operational networks, it is likely that the governing factor for introducing optical circuits will be to ensure minimal disruptions, and therefore the shortest path first and the optical bypass based approaches are more suited. Another notable result was the almost same performance of the bypass schemes with and without destination based forwarding consideration, ensuring that the mechanism may be deployed not only MPLS, but also in traditional IP networks.

We also studied the scenarios under unknown traffic matrix, where only basic traffic measurements are used to compute the size and placement of optical circuits. Our results indicate that the performance of the bypass schemes is significantly better than the shortest path first scheme, and less likely to be adversely affected by the topology in place. As complete traffic matrix information is mostly unavailable in real networks, the optical bypass schemes carry potential for dynamic traffic scenarios. The bypass schemes are also comparatively easier to configure in multi-service systems, typical to most commercial ISPs, by ensuring critical services not be affected by the introduction of optical circuits.

We addressed the dynamic circuit switch-off mechanisms and present the optimal formulation in the form of an updated objective function. This problem, however, requires further study, especially in determining the *trigger* or the network conditions/thresholds at which the optimization should be initialized in order to switch-off optical circuits. Our current model also does not consider the issues arising from the behavior of other common IP protocols. For example, the re-routing of a high-bandwidth TCP/IP flow over a new optical circuit can lead to packet re-ordering and losses, which in turn can temporarily decrease the network throughput. Another interesting future study would be to analyze this behavior in order to understand the time scales for which the use of dynamic optical circuits may be beneficial in IP networks. Future study also entails the study of the different mechanisms in case of network failures, with both, i.e., the use of dynamic optical circuits during link failures and the effect of link failures in a network with dynamic optical circuits, presenting an important challenge.

References

[1] Policy based routing. [Online]. Available: http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml

[2] R. Hao, D. Lee, J. Ma, and J. Yang, "Fault management for networks with link state routing protocols," in *Network Operations and Management Symposium, 2004. NOMS 2004. IEEE/IFIP*, vol. 1, april 2004, pp. 103 –116 Vol.1.

[3] K. Shiomoto, I. Inoue, and E. Oki, "Multi-layer network operation and management for future carrier backbone networks," in *Global Telecommunications Conference, 2008. IEEE GLOBECOM 2008. IEEE*, 30 2008-dec. 4 2008, pp. 1 –5.

[4] T. Miyamura, E. Oki, I. Inoue, and K. Shiomoto, "Enhancing bandwidth on demand service based on virtual network topology control," in *Network Operations and Management Symposium Workshops, 2008. NOMS Workshops 2008. IEEE*, april 2008, pp. 201 –206.

[5] T. Engel, A. Autenrieth, and J.-C. Bischoff, "Packet layer topologies of cost optimized transport networks multi-layer netwok optimization," in *Optical Network Design and Modeling, 2009. ONDM 2009. International Conference on*, feb. 2009, pp. 1 –7.

[6] S. Xu, K. Sezaki, and Y. Tanaka, "A two-stage simulated annealing logical topology reconfiguration in ip over wdm networks," in *Telecommunications Network Strategy and Planning Symposium. NETWORKS 2004, 11th International*, june 2004, pp. 327 – 332.

[7] T. Engel, A. Autenrieth, and J.-C. Bischoff, "Packet layer topologies of cost optimized transport networks multi-layer netwok optimization," in *Optical Network Design and Modeling, 2009. ONDM 2009. International Conference on*, feb. 2009, pp. 1 –7.

[8] K. Papagiannaki, N. Taft, and A. Lakhina, "A distributed approach to measure ip traffic matrices," in *ACM SIGCOMM conference on Internet measurement*, ser. IMC '04, 2004, pp. 161–174.

[9] The lambdastation project. [Online]. Available: http://www.lambdastation.org

[10] Internet2 phoebus. [Online]. Available: http://e2epi.internet2.edu/phoebus.html

[11] C. Xin, C. Qiao, Y. Ye, and S. Dixit, "A hybrid optical switching approach," in *Global Telecommunications Conference, 2003. GLOBECOM '03. IEEE*, vol. 7, dec. 2003, pp. 3808 – 3812 vol.7.

[12] I. de Miguel, J. C. Gonzlez, T. Koonen, R. Durn, P. Fernndez, and I. T. Monroy, "Polymorphic architectures for optical networks and their seamless evolution towards next generation networks," *Photonic Network Communications*, vol. 8, pp. 177–189, 2004.

[13] M. Chamania, M. Caria, and A. Jukan, "A comparative performance analysis of ip traffic offloading schemes over dynamic circuits," in *INFOCOM, 2011 Proceedings IEEE*, april 2011, pp. 136 –140.

[14] ——, "Achieving ip routing stability with optical bypass," in *Advanced Networks and Telecommunication Systems (ANTS), 2009 IEEE 3rd International Symposium on*, dec. 2009, pp. 1–3.

[15] ——, "Effective usage of dynamic circuits for ip routing," in *Communications (ICC), 2010 IEEE International Conference on*, may 2010, pp. 1–6.

[16] ——, "Achieving ip routing stability with optical bypass," vol. 7, no. 4, 2010, pp. 173 – 184, selected Papers from the Third International Symposium on Advanced Networks and Telecommunication Systems (ANTS 2009).

[17] M. Karol, "A distributed algorithm for optimal (optical) bypass of ip routers to improve network performance and efficiency," in *Information Sciences and Systems (CISS), 2011 45th Annual Conference on*, march 2011, pp. 1 –6.

[18] J. Fernandez-Palacios, L. Perez, J. Rodriguez, J. Dunne, and M. Basham, "Ip offloading over multi-granular photonic switching technologies," in *Optical Communication (ECOC), 2010 36th European Conference and Exhibition on*, sept. 2010, pp. 1 –3.

[19] M. Jinno, Y. Sone, H. Takara, A. Hirano, K. Yonenaga, and S. Kawai, "Ip traffic offloading to elastic optical layer using multi-flow optical transponder," in *Optical Communication (ECOC), 2011 37th European Conference and Exhibition on*, sept. 2011, pp. 1 –3.

[20] Bidirectional forwarding detection for ospf. [Online]. Available: http://www.cisco.com/en/US/technologies/tk648/tk365/tk480/technologies_white_paper0900aecd80244005.pdf

[21] *Interfaces for the Optical Transport Network (OTN)*, ITU-T Std. Rec. G.709/Y.1331, Feb. 2012.

[22] B. S. A. L. Swallow, G., "Avoiding equal cost multipath treatment in mpls networks," *IETF RFC 4928*, June 2007. [Online]. Available: http://tools.ietf.org/rfc/rfc4928

[23] Configuring custom queueing, cisco ios 12.0,. [Online]. Available: http://www.cisco.com/en/US/docs/ios/12_0/qos/configuration/guide/qccq.html

[24] Router-switch.com world's leading cisco supplier,. [Online]. Available: http://www.router-switch.com

[25] B. Mukherjee, *Optical WDM Networks (Optical Networks)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006, p. 422.

[26] A. Bley, “On the approximability of the minimum congestion unsplittable shortest path routing problem,” in *International conference on Integer Programming and Combinatorial Optimization*, ser. IPCO’05, 2005, pp. 97–110.

Appendix A
Summary of the Notation

General Constants

Symbol	Description
$G(V,E)$	Directed graph of the IP network topology with routers in set V and edges given by set E
v_i	Router i in the IP network topology ($v_i \in V$)
e_{ij}	Directed IP link from v_i to v_j ($e_{ij} \in E$)
w_∞	Large link weight used when a link does not exist
$\hat{L}_{ij}$	Boolean to indicate if an IP link exists between v_i and v_j before optical circuits are established and if it is advertised in the routing protocol
$\hat{w}_{ij}$	Pre-defined link weight metric to define the weight of a link from v_i to v_j if it exists
$\hat{X}_{xy}$	Boolean indicating one or more optical circuits from v_x to v_y before starting computation
$\hat{X}^t_{xy}$	Integer indicating number of optical circuits of type $t \in T$ from v_x to v_y before starting computation
C^O_t	Constant indicating the capacity of a dynamic circuit of type $t \in T$ in the transport network
T	Set of granularities for all dynamic circuit granularities supported in the Transport Network
N	Max. number of optical circuits between router pairs
$Cost^t_{xy}$	Cost of establishing a circuit of type $t \in T$ from v_x to v_y in the transport network
$IFCost^t$	Normalized Cost of Optical interface of type $t \in T$
$HopCount_{xy}$	Number of edges in the shortest path computed on the base topology in our performance evaluation
SW^t_{xy}	Profit for switching off a single dynamic optical circuit of type $t \in T$ from v_x to v_y
$CostPerHop^t$	Cost of circuit of type t per optical link
H	Average hop-count between $s-d$ pairs in a topology

Routing Constants

Symbol	Description
$\hat{r}^{sd}_{ij}$	Boolean to indicate if the route from v_s to v_d uses e_{ij} before establishment of optical circuits
$\hat{RC}_{sd}$	Route cost from v_s to v_d before establishment of optical circuits
ψ^{sd}_{xy}	Boolean indicator if route from v_s to v_d traverses v_x and v_y before establishment of optical circuits
$\psi^{sd}_{xy}(ij)$	Boolean indicating if route from v_s to v_d traverses v_x and v_y and the segment v_x to v_y traverses link e_{ij} before establishment of optical circuits
P	Large positive constant $>> \sum \hat{w}_{ij}$
P_∞	Large positive constant to linearize (22)

Traffic Constants

Symbol	Description
α	Link load threshold in the IP network
$\hat{\lambda}_{sd}$	(Known) IP Traffic from v_s to v_d
$LinkLoad_{ij}$	Traffic measured on link e_{ij} before establishment of optical circuits
γ^j_{ik}	Virtual Output Queue Traffic measured on v_j incoming from v_i and routed to v_k before establishment of optical circuits
D	Set of linear traffic expressions obtained from link load and VoQ traffic measurements
D_i	Unique Traffic Expression in D
d^i_{sd}	Boolean coefficient of λ_{sd} in D_i
B_i	Value of expression D_i
λ_∞	Very large positive constant $>>$ sum of all traffic in the IP network
λ^{max}_{sd}	Upper bound computed on traffic from v_s to v_d using the set of traffic expressions in D

Variables

Symbol	Description
λ_{sd}	Variable indicating IP Traffic from v_s to v_d
X^t_{xy}	Integer variable to indicate if a dynamic circuit of capacity C^O_t is provisioned from v_x to v_y
X_{xy}	Boolean to indicate an optical circuit from v_x to v_y
L_{ij}	Boolean to indicate if an IP link exists between v_i and v_j after optical circuits are established and is advertised in the routing protocol.
r^{sd}_{ij}	Boolean to indicate if route from v_s to v_d uses the link from v_i to V_j
$FT^d_i(j)$	Boolean to indicate if forwarding table at v_i has link to v_j as next hop to destination v_d
RC_{sd}	Routing cost for the route from v_s to v_d
w_{ij}	Link weight metric to define the weight of a link from v_i to v_j
V_{sd}	Boolean to constrain routing re-convergence in SPF
Y^t_{xy}	Boolean to indicate if $X^t_{xy} > \hat{X}^t_{xy}$
Z^t_{xy}	Integer to linearize the product $X^t_{xy} \cdot Y^t_{xy}$
f^d_{xy}	Boolean to indicate if traffic to v_d uses a bypass from v_x to v_y
C_{ij}	Total link capacity on IP link e_{ij}
a^{sd}_{ij}	Boolean to indicate the contribution of λ_{sd} for traffic on link from v_i to v_j for formulation in unknown traffic matrix conditions
S^x_{ij}	Boolean to indicate if traffic on e_{ij} matches the traffic expression $D_x \in D$
S_{ij}	Boolean to indicate if no $D_x \in D$ matches the traffic expression on link e_{ij}

Appendix B
Routing of Circuits in Optical Networks

The number, granularity and location of optical circuits required in the IP network are indicated by the variable X^t_{xy}. We assume an optical network with nodes $v^{OT}_x \in V^{OT}$ and directional links $e^{OT}_{ij} \in E^{OT}$ from v^{OT}_i to v^{OT}_j with *available* capacity C^{OT}_{ij}. Note that we use available capacity and not total capacity so as to exclude any capacity in the optical network that is used to provision fixed links. For simplicity of representation, we assume that the router v_x in the IP network is connected to the switch v^{OT}_x in the optical network. Therefore, X^t_{xy} optical circuits of type $t \in T$ must be routed in the optical network from v^{OT}_x to v^{OT}_y. A positive integer variable $R^{xy}_{ij}(t)$ is used to indicate the number of circuits of type $t \in T$ are reserved on the link e^{OT}_{ij} for connections from v^{OT}_x to v^{OT}_y. Routing in the optical network is subject to the following typical constraints

$$\forall t \in T, V^{OT}_x, V^{OT}_y \in V^{OT} \quad :$$

$$\sum_i r^{xy}_{xi}(t) = \sum_i r^{xy}_{iy}(t) = X^t_{xy} \tag{B.1}$$

$$\sum_i r^{xy}_{ix}(t) = \sum_i r^{xy}_{yi}(t) = 0 \tag{B.2}$$

$$V^{OT}_i \in V^{OT}, i \neq x, i \neq y \quad : \quad \sum_k r^{xy}_{ki}(t) = \sum_j r^{xy}_{ij}(t) \tag{B.3}$$

$$\forall E^{OT}_{ij} \in E^{OT} \ : \ \sum_{xy}\sum_t r^t_{xy}(ij) \cdot C^{OT}_t \leq C^{OT}_{ij} \tag{B.4}$$

(B.1) ensures that the total number of circuits starting from the source and terminating at the destination are equal to the number of optical circuits required, (B.2) ensures that no circuits terminate at the source or begin at the destination and (B.3) ensures routing continuity for the circuits in the rest of the network. Capacity usage of the links in the optical network is constrained in (B.4) using the routing variable and the available capacity information C^{OT}_{ij} of the links in the transport network. The formulation presents the provisioning of circuits in a generic optical network, and technology specific constraints can also be added to map to specific optical technologies.