A Comparative Performance Analysis of IP Traffic Offloading Schemes over Dynamic Circuits

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

***Abstract*—We present an novel analytical framework for modeling current IP offloading schemes and show their comparative performance. We analyze relevant parameters, including IP routing re-convergence time, the number of affected IP routes, and the cost of circuit capacity used for offloading. Our results show that emerging offloading solutions based on *invisible bypasses* can better maintain network stability and reduce the routing reconfiguration effort compared to the traditional traffic engineering approaches that modify the IP routing.**

I. Introduction

Network operators currently operate their networks at about 30%-40% link utilization to deal with unexpected traffic surges as well as failures in the IP network [1]. The exponential rise in IP traffic however, requires link capacity upgrades every 6-12 months to maintain the required utilization levels which will soon limit the scalability of legacy IP networks. To cope with the same, carriers are exploring the use of Dynamic Circuit Services (DCS) to *offload* excess IP traffic onto dynamic circuits established on the fly. Recognizing the fact that most routers are commonly co-located with circuit switches based on optical transmission and that cost of optical switching is independent of the data rate [2], IP offloading may prove essential in paving the way for integration of IP services on new carrier-grade technologies.

However, coordinating operations between circuit and packet networks is non-trivial. Adding capacity via dynamic circuits in IP networks poses significant planning as well as operational complications. A new link in the IP network can trigger routing re-convergence, which in turn affects the traffic in other parts of the network, and must be taken into consideration in the planning phase. Moreover, it leads to a temporary network outage observed during the IP routing re-convergence and requires significant reconfiguration, most critically in the Event Correlation Database (ECD). In this reconfiguration, which needs to be fool-proof in order to ensure that any network outage can be identified and dealt with quickly, the ECD must not only examine the events whose triggers contain data loss in that particular IP flow, but must also update all failures associated with elements along the new data path of the IP flow. For these and other reasons, it is unlikely that any revolutionary dynamic "multi-layer" networking solution will be deployed in the near future.

In this paper, we develop a novel analytical planning model and present a comparative performance analysis of various offloading scenarios in the context of Internet architecture evolution. We assume a typical network, where routers are co-located with transport switches (e.g. SDH/SONET, WDM, MPLS-TP) and where dynamic circuits can be setup in two basic scenarios: 1) to upgrade the capacity of an existing IP link, 2) to create a new link between a pair of IP routers. In addition to this two basic scenarios, we analyze the case where setup circuits hidden from the IP routing service [4]. To this end, we developed an comprehensive[1] analytical framework, which can determine the optimal location and capacity required for provisioning dynamic circuits for a traffic offloading operation. The results show the that the latter approach carries significant potential in the evolution of the Internet.

II. Overview of IP Offloading Mechanisms

In a typical IP network, routers are interconnected with high-bandwidth leased lines forming the IP topology and use a shortest path first (SPF) routing variant, such as OSPF. Services on the IP network can also use explicitly-defined manually-configured routes defined via policy-based IP routing [6] or via statically-routed MPLS tunnels. Typically, due to the complexity of configuration, carriers only provision high guarantee services such as MPLS VPNs using statically-routed MPLS tunnels and regular Internet traffic is routed either over pure IP or via MPLS using LDP. The transport network cloud is assumed to be reconfigurable, such as a WDM network that consists of ROADMs (reconfigurable optical add-drop multiplexer) connected by optical fibers and supports a dynamic circuit service which can set up circuits on the fly. Since most routers are co-located with transport switches, and use leased line services on the transport network to generate the IP topology, we assume they can also use the dynamic circuit service to add additional capacity in the IP network. The latter is classified into two categories:

A) Offloading with circuits as new IP links: This is the traditional Traffic and Network Engineering scheme [5] where a new circuit is advertised as a new IP link in the SPF routing protocol; this action typically triggers routing re-convergence. Depending on the size of the SPF area, an ISP may experience a temporary service outage during the re-convergence. While efficient in terms of utilizing added capacity, this mechanism may lead to significant routing changes in the IP network, and it also triggers reconfiguration of monitoring and failure detection services.

[1] For readers novices to optimizations, the original long version of the analysis can be obtained from the authors directly.

B) Offloading with circuits not advertised in IP routing service: This mechanism was first introduced in Lambdastation [7], a circuit setup system used in large scientific networks. In order to provide DCS setup control to the operator, either and API or flow-based selective offloading approaches like [8] can be used, wherein providers attempt to identify high-bandwidth long-lived flows and offload them onto dynamic circuits established between edge routers. The effectiveness of this approach is limited by the quality of the flow classifier used. The edge-to-edge constraint in this approach however does not allow for optimal utilization of resources, especially those in the middle of the network where a circuit can be used for aggregated flows also. In our past work [4], we proposed a mechanism to establish dynamic circuits between any two IP routers as *invisible* optical bypasses between any pair of routers, not just edge routers. The bypasses of this kind do not interfere with the IP routing service to offload traffic which contains multiple multiplexed flows. To ensure that bypasses do not affect IP routing, we create bypasses between routers only for flows that use these routers on the original IP routing paths[2]. Specialized forwarding mechanisms like policy-based routing are then used to divert flows onto the bypass at the ingress router.

III. Comparative Analysis

In this section, we provide a novel analytical framework based on Integer Linear Programming (ILP) for modeling IP traffic offloading which incorporates the above-mentioned scenarios, dynamic circuits with and without IP link advertisement in IP routing. To this end, we first present the generic framework description, followed by modeling.

In the framework, we assume to have information about the IP and the transport network topology and their interconnections, and the traffic and routing information in the network[3] The transport network topology is defined by the graph $G^p(V^p, E^p)$, with vertices $v_i^p \in V^p$ and edges $e_{ij}^p \in E^p$, and the IP network topology is defined by $G^l(V^l, E^l)$, with vertices $v_i^l \in V^l$ and edges $e_{ij}^l \in E^l$. Note that the edges in E^l represent the static leased lines used in the network, and the constant L_{ij} indicates if a static link exists from v_i^l to v_j^l.

The topology information also contains information about the capacity of the links in the IP and the transport network, given by C_{ij}^l and C_{ij}^p respectively. In case the link does not exist, the capacity is set to 0. We note that there may be different services running simultaneously in the network. Given that there are Q services in the network, the total traffic for a service $q \in Q$ entering the core network at v_s^l and leaving the network at v_d^l is given by λ_{sd}^q in terms of the required capacity (in Gb/s). It is assumed that for each service q, a unique (single) routing path is used between a pair of (source, destination) routers. The routing in the IP network is represented using the boolean routing constants $\psi(q)_{xy}^{sd}$ and $\psi(q)_{xy}^{sd}(ij)$. $\psi(q)_{xy}^{sd}$ indicates if the IP route for service $q \in Q$ from v_s^l to v_d^l crosses the routers v_x^l and v_y^l (in that order, $x \neq y$). Note that v_x^l and v_y^l need not be neighboring routers in the network. The boolean $\psi(q)_{xy}^{sd}(ij)$ indicates if the link e_{ij}^l lies on the IP route for service $q \in Q$ from v_s^l to v_d^l between the routers v_x^l and v_y^l ($x \neq y$). Note that $\psi(q)_{xy}^{sd}(ij) \leq \psi(q)_{xy}^{sd}$.

Given that most transport networks and upcoming DCSs traditionally only support fixed granularities, we assume that there are T different granularities available for establishing dynamic circuits, and for each granularity type $t \in T$, capacity of the circuit is given by C_{DC}^t. A binary variable X_{xy}^t indicates formation of a dynamic circuit of type $t \in T$ from switch v_x^p to v_y^p. In this formulation, for simplicity of presentation, router v_i^l in the IP network is assumed to be co-located with the switch with the same index v_i^p. All IP routers are also assumed to have the capability to offload traffic over DCs. Theese are the basic constraints of our model:

$$\forall v_i^l, v_j^l \in V^l : \sum_t X_{ij}^t \leq 1 \qquad (1)$$

$$\forall t \in T, v_x^p, v_y^p \in V^p \forall e_{ij}^p \in E^p : r_{xy}^t(ij) \leq X_{xy}^t \qquad (2)$$

$$\sum_i r_{xy}^t(xi) = \sum_i r_{xy}^t(iy) \ = \ X_{xy}^t \qquad (3)$$

$$v_i^p \in V^p, i \neq x, i \neq y : \sum_k r_{xy}^t(ki) = \sum_j r_{xy}^t(ij) \qquad (4)$$

$$\forall e_{ij}^p \in E^p : \sum_t \left(C_{DC}^t \sum_{xy} r_{xy}^t(ij) \right) \leq C_{ij}^p \qquad (5)$$

$$\forall e_{ij}^l \in E^l : \hat{C}_{ij}^l = C_{ij}^l + \sum_t X_{ij}^t \cdot C_{DC}^t \qquad (6)$$

$$Cost_{xy}^t = Cost^t \cdot X_{xy}^t + \sum_{ij} P_{ij}^t \cdot r_{xy}^t(ij) \qquad (7)$$

$$Minimize: \ \sum_t \sum_{xy} Cost_{xy}^t \qquad (8)$$

We constrain the dynamic circuits such that only *one circuit is instantiated between a source-destination switch pair* during an offloading operation (1). Additional constraints can also be introduced to limit the offloading capability of IP routers. Note that in both IP offloading scenarios presented in Section II, a new circuit between *neighboring routers* only increases the existing link capacity and the capacity of an existing link e_{ij}^l after traffic offloading is given by $\hat{C}_{ij}^l$ (6). To determine the cost of the IP traffic offloading, we assign a cost to a) bandwidth used in the transport network by the dynamic circuits and b) the interfaces used at the DC end-points. To determine the cost of the bandwidth, for each link $e_{ij}^p \in E^p$, we assign a cost P_{ij}^t for provisioning a circuit of type $t \in T$ on the link and use this in conjunction with the routing path to determine the total bandwidth cost. The interface cost is associated with activating a new interface or reserving capacity on an existing interface at the IP router and for is given by a

[2]This condition is important, see [4].

[3]Topology information can be obtained via the network management system and interconnection information is configured manually. Routing information for SPF variants can be extracted via passive monitoring [9] typically used in commercial IP/MPLS networks, or by probing individual routers using SNMP. Obtaining the traffic matrix is however fundamentally difficult [10], and in this paper version, we do not consider the traffic matrix estimation methods.

constant $Cost^t$ for a circuit of type $t \in T$. In this formulation, the total cost of a circuit of type $t \in T$ from v_x^p to v_y^p given by $Cost_{xy}^t$ (7). Here, if a dynamic circuit is not established (i.e. $X_{xy}^t = 0,\ r_{xy}^t(ij) = 0$) $Cost_{xy}^t = 0$. We minimize the total cost of the IP offloading operation with (8)

We now provide the relevant formulation for both offloading strategies. Given the numerous suggestions for IP offloading, it is not possible to present all facets of different proposals to date for space reasons. In this framework, we show how to account for IP routing changes when operating under both SPF and static routing, and also present a generic framework for offloading using bypasses, which can be adapted to both our proposal [4] and to the end-to-end DC approach.

A. *Traffic Offloading with IP Routing Changes*

In this approach, a dynamic circuit between non-neighboring routers is represented as a new link, which in turn impacts the IP routing. To model this behavior, a binary variable $\hat{L}_{ij}$ is used to indicate if a link exists in the IP topology after offloading operation. Constraints (9) and (10) ensure that $\hat{L}_{ij}$ is equal to 1 if a link (L_{ij}) or a dynamic circuit ($\sum_t X_{ij}^t$) exist between the nodes v_i^l and v_j^l. If neither exist, (11) constrains $\hat{L}_{ij} = 0$. As dynamic circuits are introduced as new links in this scenario, the updated link capacity ($\hat{C}_{ij}^l$) after offloading operation for all router pairs is given by (12).

$$\forall v_i^l, v_j^l \in V^l,\ i \neq j \quad : \quad \hat{L}_{ij} \geq L_{ij} \tag{9}$$

$$\forall v_i^l, v_j^l \in V^l,\ i \neq j \quad : \quad \hat{L}_{ij} \geq \sum_t X_{ij}^t \tag{10}$$

$$\forall v_i^l, v_j^l \in V^l,\ i \neq j \quad : \quad \hat{L}_{ij} \leq \sum_t X_{ij}^t + L_{ij} \tag{11}$$

$$\forall v_i^l, v_j^l \in V^l,\ i \neq j \quad : \quad \hat{C}_{ij}^l = C_{ij}^l + \sum_t X_{ij}^t \cdot C_{DC}^t \tag{12}$$

As routing can change with introduction of new links, a binary routing variable $r(q)_{ij}^{sd}$ is introduced which indicates if the route from v_s^l to v_d^l for service $q \in Q$ uses the IP link from v_i^l to v_j^l. The routing variable $r(q)_{ij}^{sd}$ is constrained by the requirement that a link must exist from router v_i^l to v_j^l ($\hat{L}_{ij} = 1$) for $r(q)_{ij}^{sd} = 1$ (13).

$$\forall q \in Q, v_s^l, v_d^l, v_i^l, v_j^l \in V^l : r(q)_{ij}^{sd} \leq \hat{L}_{ij} \tag{13}$$

The routing variable is also used in (14) to ensure that the link utilization does not exceed a fixed utilization threshold α.

$$\forall v_i^l, v_j^l \in V^l : \sum_q \sum_{sd} r(q)_{ij}^{sd} \cdot \lambda_{sd}^q \leq \alpha \hat{C}_{ij} \tag{14}$$

The two commonly used routing schemes, namely static routing and SPF routing (with re-convergence), differ in the manner in which IP routing is configured after dynamic circuit setup or release. We now present the unique constraints on the routing in these models separately. Note however, that in a single network, different services q can use different routing schemes, and static routing and SPF routing can co-exist in the same network.

1) Static (Policy-Based) Routing: Static routing allows the operator complete control over the routing in the IP layer. As mentioned before, we constrain the routing from v_s^l to v_d^l for a service q to use a single path using (15) and (16) to ensure routing continuity in the IP network, where constraint (15) ensures that for any given service q, the route from v_s^l to v_d^l uses one outgoing link from v_s^l and one incoming link into v_d^l while (16) ensures the routing continuity from v_s^l to v_d^l.

$$q \in Q, \forall v_s^l, v_d^l \in V^l : \sum_j r(q)_{sj}^{sd} = \sum_j r(q)_{jd}^{sd} \ = \ 1 \tag{15}$$

$$q \in Q, \forall v_s^l, v_d^l, v_j^l \in V^l, j \neq s, d : \sum_i r(q)_{ij}^{sd} = \sum_k r(q)_{jk}^{sd} \tag{16}$$

Note that while the equations sum over all routers in the IP network which may or may not be connected to the router in question, (13) ensures that only routers which are directly connected after offloading operation are considered.

Static routing allows very flexible re-routing of traffic during offloading operations, and *service constraints* may be applied in order to meet different service requirements.

2) SPF based routing: The SPF based IP routing follows the basic tenets of routing continuity (15) (16) with two additional requirements: 1) A route always use the shortest path from source to destination, and 2) The routing decision at a router is destination-based.

In this formulation, we define link weights for all IP links and routing cost from a source to a destination router. We assume that the link weight is pre-defined as w_{ij}^l based on a static metric. So when a dynamic circuit is established from v_i^l to v_j^l, the IP link weight is given by w_{ij}^l, and is otherwise infinite (w_∞). Using this definition, the actual link weight $\hat{w}_{ij}^l$ is given by the expression in (17) and the route cost from v_s^l to v_d^l for the service q is given be the expression in (18). Note that (18) is a non-linear expression. However, if a link does not exist ($\hat{L}_{ij} = 0, \hat{w}_{ij}^l = w_\infty$), the routing variable $r(q)_{ij}^{sd}$ is already constrained to be 0 by (13), and when $\hat{L}_{ij} = 1$, we have $\hat{w}_{ij}^l = w_{ij}^l$ (17). Using this, (18) is reduced to (19).

$$\forall v_i^l, v_j^l, \in V^l : \hat{w}_{ij}^l = \hat{L}_{ij} \cdot w_{ij}^l + \left(1 - \hat{L}_{ij}\right) \cdot w_\infty \tag{17}$$

$$q \in Q, \forall v_s^l, v_d^l, \in V^l : RC_{sd}^q = \sum_{ij} \hat{w}_{ij}^l \cdot r(q)_{ij}^{sd} \tag{18}$$

$$= \sum_{ij} w_{ij}^l \cdot r(q)_{ij}^{sd} \tag{19}$$

Using the cost function (19), we introduce the shortest path routing constraints. (20) presents an upper bound on the cost for reaching the destination v_d^l from v_s^l via an immediate neighboring router v_x^l which is equal to the shortest path cost. This constraint in conjunction with the routing continuity constraints (15) and (16) ensure that only the shortest routing path from the source to the destination is selected.

$$\forall v_s^l, v_d^l, \in V^l, s \neq d : RC_{sd}^q \leq RC_{sx}^q + \hat{w}_{xd}^l \tag{20}$$

In order to model destination-based forwarding, the boolean variable $FT(q)_i^d(j)$ mimics a forwarding table, with $FT(q)_i^d(j)$ indicating if at v_i^l, the traffic for service q to v_d^l

is forwarded over a link from v_i^l to v_j^l. (21) ensures that the forwarding table at v_i^l can only have v_j^l as a next hop for any v_d^l if a link exists from v_i^l to v_j^l. (22) describes the relationship between the routing $r(q)_{ij}^{sd}$ and $FT(q)_i^d(j)$, ensuring that for a service q, if the route from v_s^l to v_d^l uses the link between v_i^l and v_j^l, $FT(q)_i^d(j) = 1$, while (23) ensures that the forwarding table for a service q at any v_i^l has one-and-only one next hop choice for a destination v_d^l.

$$q \in Q, \forall v_s^l, v_d^l, v_i^l, v_j^l \in V^l,\ i \neq d : FT(q)_i^d(j) \leq \hat{L}_{ij} \quad (21)$$

$$q \in Q, \forall v_s^l, v_d^l, v_i^l, v_j^l \in V^l,\ i \neq d : FT(q)_i^d(j) \geq r(q)_{ij}^{sd} \quad (22)$$

$$q \in Q, \forall v_d^l, v_i^l \in V^l, i \neq d : \sum_j FT(q)_i^d(j) = 1 \quad (23)$$

B. Bypass-based IP offloading without routing re-convergence

In the bypass-based IP offloading model, a dynamic circuit between non-neighboring routers is not advertised as a new IP link. A boolean variable $f(q)_{xy}^{sd}$ indicates if the traffic for service q from v_s^l to v_d^l is offloaded onto a bypass from v_x^l to v_y^l, and (24) ensures that traffic is only offloaded when a dynamic circuit exists from v_x^l to v_y^l. From the basic tenets of bypass-based IP offloading as described in Section II, it also follows that for a service q, traffic from v_s^l to v_d^l can only use a bypass from v_x^l to v_y^l, if both v_x^l and v_y^l lie on the original routing path, which is ensured in (25) using the original routing information defined in the constant $\psi(q)_{xy}^{sd}$.

$$q \in Q, \forall v_s^l, v_d^l, v_x^l, v_y^l \in V^l, L_{xy} = 0 : f(q)_{xy}^{sd} \leq \sum_t X_{xy}^t \quad (24)$$

$$q \in Q, \forall v_s^l, v_d^l, v_x^l, v_y^l \in V^l, L_{xy} = 0 : f(q)_{xy}^{sd} \leq \psi(q)_{xy}^{sd} \quad (25)$$

The formulation does not constrain the number of bypasses used along the routing path from v_s^l to v_d^l. However, two bypasses used along the same route cannot overlap. For a given service and a source-destination pair, overlapping bypasses will have at least one IP link common in the overlapped segments, and (26) uses $\psi(q)_{xy}^{sd}(ij)$ to ensure multiple bypasses which bypass e_{ij}^l are not used for service q from v_s^l to v_d^l.

$$q \in Q, \forall v_s^l, v_d^l \in V^l, e_{ij}^l \in E^l : \sum_{xy: L_{xy}=0} \psi(q)_{xy}^{sd}(ij) \cdot f(q)_{xy}^{sd} \leq 1 \quad (26)$$

Constraints for the capacity utilization of the IP links and the bypasses are defined in (27) and (28) respectively. In (27), the original traffic on the link e_{ij}^l is given by $\lambda_{sd}^q \cdot \psi(q)_{ij}^{sd}$ and the additional factor excludes traffic bypassed over e_{ij}^l.

$$\forall e_{ij}^l \in E^l : \sum_q \sum_{sd} \lambda_{sd}^q \cdot \psi(q)_{ij}^{sd} \cdot \left(1 - \sum_{xy: L_{xy}=0} \psi(q)_{xy}^{sd}(ij) \cdot f(q)_{xy}^{sd}\right) \leq \alpha \cdot \hat{C}_{ij}^l \quad (27)$$

$$\forall v_x^l, v_y^l \in V^l s.t. L_{xy} = 0 : \sum_q \sum_{sd} \lambda_{sd}^q \cdot f(q)_{xy}^{sd} \leq \alpha \sum_t X_{xy}^t \cdot C_{DC}^t \quad (28)$$

TABLE I
DATA OF AVAILABLE DYNAMIC CIRCUITS

Capacity (Gb/s)	1	2	5	10	20	50	100
Interface Cost	20	30	50	70	100	150	250
Cost/Hop	2	3	5	7	10	15	25

C. Decommission of Dynamic Circuits

After an IP offloading operation which installs dynamic circuits, we now describe how to decommission dynamic circuits in the network while ensuring that routing and capacity utilization constraints are met. We define a boolean constant $\hat{X}_{xy}^t$ which indicates if a dynamic circuit of type t is currently provisioned from v_x^p to v_y^p. The objective function (8) is then modified to incorporate the following: 1) no additional cost incurred to keep existing dynamic circuits and 2) Decommissioning of a dynamic circuit is associated with a profit $SwProf_{xy}^t$ (29)(SW^t is a profit associated with decommissioning a dynamic circuit of type $t \in T$). Note that $SwProf_{xy}^t$ is non-zero only when an existing dynamic circuit ($\hat{X}_{xy}^t = 1$) is decommissioned ($X_{xy}^t = 0$).

$$SwProf_{xy}^t = \hat{X}_{xy}^t \cdot SW^t \cdot \left(1 - X_{xy}^t\right) \quad (29)$$

The objective function is now given by (30), where $(1 - \hat{X}_{xy}^t) \cdot Cost_{xy}^t$ ensures that the cost of establishing a dynamic circuit is not included if it was already commissioned while the term $SwProf_{xy}^t$ incorporates the profit incurred by decommissioning a dynamic circuit.

$$Min: \sum_t \sum_{xy} \left[\left(1 - \hat{X}_{xy}^t\right) \cdot Cost_{xy}^t - SwProf_{xy}^t\right] \quad (30)$$

IV. PERFORMANCE EVALUATION

In the performance study, we use the NSFNet topology (14 nodes, 40 unidirectional links) as a reference topology. Each IP router is assumed to be co-located with a transport network switch, with unlimited inter-layer (IP-transport network) interface capacity. Each unidirectional transport network link is assumed to be 100 GB/s carrier Ethernet and 7 dynamic circuit granularities are supported as shown in Table I. We use the traffic matrix from [11] which is scaled by 10^5 (all λ_{ss} set to 0). Initial IP routing is computed using SPF routing, and is used to compute the initial traffic on all IP links. The static IP link capacities are then set so that IP link initial utilization is 60%. The max. utilization threshold $\alpha = 0.7$.

The following mechanisms are used for comparison, i.e.,

1) Dynamic circuits added to existing IP links: `IncrCap`,
2) New IP link advertised with SPF routing: `SPF`,
3) New IP link advertised with static routing: `Static`,
4) New IP link advertised with static routing with a policy restriction on hop count: `StaticHR`,
5) Bypass-based IP offloading: `BBR`

`IncrCap` is introduced to give a baseline costs involved in offloading operations commonly used in current network. To trigger a traffic offloading operation, two traffic loading scenarios are used: 1) increase the traffic of five randomly chosen source-destination pairs, while varying the percentage increase in traffic and 2) vary the number of source-destination pairs selected and increase the traffic on each selected pair by

70%. Results for each scenario are averaged over 100 random cases using the Gurobi Optimizer [12].

Fig. 1 presents the average cost of offloading for the different schemes. We see that `IncrCap` has the maximum cost, while `Static` has the minimum cost. Due to a 10% gap between the initial and max IP link utilization, re-routing in the IP layer can reduce the capacity requirements for the established dynamic circuits, and the scheme with highest routing flexibility (`Static`) consequently has the lowest cost requirements. Increasing the constraints(`StaticHR`) increases the resource cost which is comparable to `SPF`. The stringent constraints on offloading enforced in `BBR` leads to higher costs as compared to `SPF` and the static routing based methods. It is clear from the results if traffic increase is permanent, the `SPF`, `Static` and `StaticHR` outperform the bypass-based traffic offloading scheme significantly. It should also be mentioned that a simple capacity increase cannot find solutions when only smaller DCS granularities (upto 20 Gb/s) are used while all other mechanisms can find a solution.

In case traffic fluctuations in the network are frequent, we must also consider the operational stability of the network which includes overheads for reconfiguration of routing. Each service path that is re-routed observes a downtime during routing re-convergence (in 10s of seconds [3]) and can affect the performance of the services running on this network. We must also consider the reconfiguration effort in the ECD. For a single service, the reconfiguration is given by the the links limited to only the original or the new routing path.

To illustrate the re-configuration effort involved with different offloading schemes, especially in networks with frequent changes, we perform a simple study. We first select 9 router pairs (Step 1) and increase the traffic between each of these pairs by 70%. In the next steps, we reset the traffic on two of the selected router pairs to the initial value. Therefore, Step 2 will have 7 , Step 3 will have 5 and Step 4 will have 3 router pairs with increased traffic. For computing the optimal solution at each step we use (30) which models the decommissioning scenario and set $SW^t = 0.1 \cdot Cost^t$. Fig. 2 shows the reconfiguration effort with the labels indicating the average number of dynamic circuits deployed. In the (BBR) scenario, each consecutive step involves the decommissioning of approx. one circuit. Even with this highly dynamic behavior, the total number of reconfigurations required is very low. On the other hand, the (Static) offloading scenario exhibits a large reconfiguration overhead even with few circuit decommissions per step. We observe a decreasing trend in the number of reconfigurations required per step as the number of possible routing configurations decrease with the number of active dynamic circuits thereby limiting the number of routing modifications in the (Static), (StaticHR) and (SPF) scenarios.

Our results indicate a trade-off between the offloading cost and the reconfiguration overhead. Based on the expected frequency of offloading, an operator may choose the best scheme: networks requiring frequent offloading would employ `BBR` to reduce reconfiguration overhead, while using `StaticHR`, `SPF` or `Static` in case of infrequent offloading.

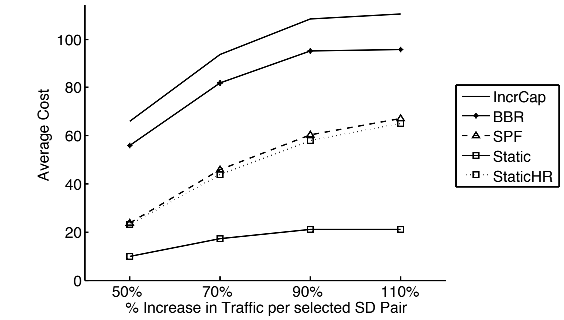

Fig. 1. Cost of IP offloading vs % increase in load for 5 S-D pairs.

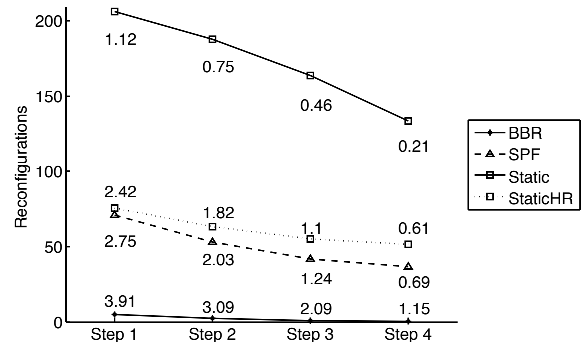

Fig. 2. Reconfigurations Effort with subsequently reducing no. of S-D pairs with increased traffic. Labels indicate the average number of dynamic circuits

V. Conclusion

In this paper, we presented the first analytical framework for performance comparison of various IP traffic offloading schemes using dynamic circuits. Our results show that while traditional capacity upgrade solutions which advertise a dynamic circuit as a new IP link are the most cost efficient, the emerging bypass-based offloading solutions significantly improve network stability and reduce the routing reconfiguration effort. From our results, it is evident that bypass-based offloading schemes are better suited for frequent short-lived offloading operations due to their low operational overhead, while traditional capacity upgrade solutions are better suited for quasi-static offloading operations.

References

[1] Internet2 Headroom : *https://wiki.internet2.edu/confluence/download/attachments/17383/Internet2+Headroom+Practice+8-14-08.pdf*

[2] M. Chamania, A. Jukan, O. de Dios, J. Jimenez, *"Offloading Excess IP Traffic with Optical Bypass: A Simple Capacity Upgrade, or More?"* , ONTC Prism Newsletters, *www.ontc-prism.org/*

[3] K. H. Liu et al. , *"Experimental Study of Dynamic IP Topology Reconfiguration in IP/WDM Networks,"* IEEE GLOBECOM, 2001.

[4] M. Chamania, M. Caria, A. Jukan, "Effective Usage of Dynamic Circuits for IP Routing," IEEE ICC'2010.

[5] L. Cheng, et al.*"Network Engineering - Control of Dynamic Link Topology in User Networks,"* Bell Labs Technical Journal (BLTJ), July 2004.

[6] Policy Based Routing - Cisco White Paper, *http://www.cisco.com/en/US/products/ps6599/products_white_paper09186a00800a4409.shtml*

[7] The Lambda Station Project: *http://www.lambdastation.org/*

[8] C. Xin, C. Qiao, Y. Ye, S. Dixit, *"A Hybrid Optical Switching Approach,"* IEEE GLOBECOM 2003.

[9] A. Shaikh, A. Greenberg, *"OSPF Monitoring: Architecture, Design and Deployment Experience,"* NSDI 2004.

[10] P. Casas et al., *"Efficient methods for traffic matrix modeling and on-line estimation in large-scale IP networks,"* ITC 2009.

[11] B. Mukherjee, *Optical WDM Networks*, Springer, 2006, pp. 422.

[12] Gurobi Optimizer 3.0, *http://www.gurobi.com/*