

Optimal Load Balancing in Federated Locator/ID Separation Protocol (LISP) Traffic Engineering

Soumick Dasgupta¹, Raghunandan Sridhar², Stefano Secci³, Deep Medhi⁴

¹Draper, USA; ²Oracle, USA; ³CNAM & NeuroTel.AI, France; ⁴National Science Foundation, USA

Abstract—Traffic Engineering (TE) in IP networks has received attention for over a quarter century, mostly within a single autonomous system. Locator/ID Separation Protocol (LISP) has been proposed to enhance Internet routing with new features, among which to allow traffic to be sent to multiple Routing Locators (RLOCs) in LISP-enabled networks. In this paper, we consider a federated environment in which multiple autonomous systems (AS) cooperate to benefit from the LISP capability to determine a load balanced traffic distribution among their RLOCs. In particular, we assume that the ASes agree to share traffic information for federation with a traffic engineering calculator (TEC) through service level agreements (SLAs). We present two novel optimization models that can be used by the TEC for balancing traffic loads in federated LISP-enabled networks that use multiple RLOCs. We compare our models to the base case when proportioning to multiple RLOCs is not available. Based on the cases we studied, we observed that LISP-based flexible traffic balancing through our approach with multiple RLOCs provides up to 78% gain (in terms of link utilization) compared to the base case when the network is asymmetric with non-uniform traffic.

I. INTRODUCTION

Locator/Identifier Separation Protocol (LISP) is an architecture that decouples the end identifiers from core network routers [1], [3]. An advantage of LISP is that the end hosts act the same way as in a traditional IP network, sending and receiving IP traffic; at the same time, LISP allows reducing the global routing table size by a mapping process [12]. We focus here on the problem of traffic balancing in LISP-enabled networks in a federated environment.

A. LISP Overview

There are a number of entities involved in a LISP-enabled environment. Hosts are referred to as Endpoint Identifiers (EIDs) that continue to use the regular IP address space, either IPv4 or IPv6, or both. In the LISP-enabled part of the Internet (marked by dotted ovals), there are Ingress Tunnel Routers (ITRs) and Egress Tunnel Routers (ETRs) that are externally reachable through many Routing Locators (RLOCs), which have an address space separated from the end identifiers. That is, there are two address spaces involved, one for end systems in the regular IP space and the second for the LISP-enabled part of the Internet. ITRs associated with the source endpoint maintain RLOC mapping for destination EIDs. At the other end, ETRs serve to deliver traffic to the destinations.

Consider Fig. 1 where packets for a flow from EID₁ are destined for EID₂. First assume that EID₁ sends the packets to its associated ITR, say ITR₁, which uses the address space

in the LISP-enabled networks to wrap the original packets with LISP-based source address as RLOC_{I1a} and LISP-based destination address as RLOC_{E1a} to send to the associated ETR marked as ETR₁. On receiving such packets at ETR₂, the LISP-based address space for RLOCs is stripped off, and the packets are delivered using EID₂ as the destination in the regular IP address space. Similarly, packets arriving at ETR₂ with LISP-based destination address RLOC_{E2a} from LISP-based source RLOC_{I1b} would be stripped off of the LISP-based addresses to deliver to EID₂. It should be understood that EID₁ to ITR₁ is not connected by a direct link; rather, this merely depicts the flow of traffic (marked with a dashed line) from EID₁ to ITR₁ without showing additional details such as other routers that may be along this path. Similarly, we use a dashed line to show traffic flows from ETR₁ to EID₂. In this figure, ITR₁ and ITR₂ are shown to be in an autonomous system (AS) to which EID₁ and EID₂ are connected to.

The LISP architecture allows multiple RLOCs to be associated with EIDs. As for the example in Fig. 1, packets arriving from EID₃ at ITR₁ that are destined for EID₂ could be set to LISP-based source address RLOC_{I1b} at ITR₁ that could use ETR₂ reachable by RLOC_{E2b} for forwarding the traffic in the LISP-enabled networks. ETR₂ would then take responsibility of delivering the packets to EID₂. In other words, traffic arriving at an ITR may use multiple ETRs to route traffic for the destination EIDs. While the LISP provides the feature for multiple RLOCs, it is up to the network providers to determine which nodes in a network may act as LISP tunnel routers. While in the figure, EIDs are shown to be connected to a single AS, there are no restrictions; that is, originating EIDs may be connected to multiple ASes (not shown in the figure). Similarly, at the receiving end, ETR₁ and ETR₂ are shown in




Fig. 1: Relating EID, ITR, ETR and RLOC




Fig. 2: Traffic Engineering Calculator cooperating with ASes

an AS, and traffic may egress from multiples ASes to EID₂.

B. Scope of this Work

We address the traffic balancing problem of *proportioning* traffic with multiple RLOCs for the destination EIDs. The LISP-enabled part of the Internet may span multiple autonomous systems (ASes) (marked using dotted ovals in Fig. 1). Our approach assumes that multiple ASes with LISP capabilities agree to federate through service level agreements (SLAs) with a *traffic engineering calculator* (TEC).

We introduce the notion of a *group* in LISP-enabled networks for traffic proportioning. A group consists of two or more routers identified as RLOCs that may be associated with an ETR for traffic load balancing. For example, in Fig. 1, we may designate a group to consist of RLOC_{2a} and RLOC_{2b}. Typically, in practice, such a group would be RLOCs that are in geographic proximity to each other. Multiple such groups can be spread out based on geographic locations. In the first case, one of the routers (RLOCs) in a group serves as the primary destination if proportioning for multiple RLOCs was not initiated. This node is referred to as the primary RLOC or primary router and the remaining routers (one or more) in the same group as secondary routers. In the second case, all RLOCs in a group can serve as either primary or secondary routers without any distinction. In other words, traffic destined for *any router* in the group may be split among its peers within the group, without distinguishing which router is the primary RLOC. The cooperation for either of these two cases falls on the TEC that interacts with multiple ASes (Fig. 2).

For the two cases with grouping introduced above, we present optimization formulations that can be used by the TEC for traffic load balancing of LISP-enabled networks when multiple RLOCs in a group are present. The rest of the paper is organized as follows. The related work is discussed in Section II. In Section III, we present our formulations for the LISP-enabled traffic engineering model that can be used by a traffic engineering calculator. In Section IV, we present and discuss numerical results to compare LISP-enabled traffic engineering to base-case traffic engineering. We then make a few remarks on LISP-TE in Section V, followed by conclusions in Section VI.

II. RELATED WORK

A number of works discussed the benefits of LISP. In [12], Quoitin *et. al.* discussed how LISP allows a better aggregation of prefixes to reduce the routing-information-base/forwarding-information-base (RIB/FIB) sizes. The mapping reduction is studied in [6]. [7] addressed enhancing LISP with egress control. [9] presented benefits in the presence of multipath TCP. [14] presented a cooperative ISP-Driven Informed path Selection (IDIPS) system.

Reduction in latency was considered in LISP traffic engineering [2]. A federated approach for interdomain traffic engineering was presented in [10], [11].

One of the most relevant recent contributions to the implementation side of LISP is [16]: it introduces P4-LISP, a data-plane LISP implementation on programmable switches using the P4 language, aiming at line-rate operation on hardware targets such as Intel Tofino. By leveraging programmable data plane language control, the work demonstrated how the locator/identifier split paradigm can be realized with forwarding performance suitable for high-speed networks, substantially reducing control-plane bottlenecks and enabling scalable multi-homing, mobility support, and traffic engineering in hardware-accelerated environments.

It is worth noting that there is no incentive for traffic engineering in a LISP-context without a cooperative routing mechanism as observed in [15], especially since multiple autonomous systems (ASs) are involved. Finally, none of works in the current literature presents formal optimization formulations for traffic steering that we present in this work in a federated environment.

III. MATHEMATICAL PROGRAMMING FORMULATION

In our proposed federated environment, a TEC collects information from multiple ASes on source and destination with projected traffic, based on service level agreements with the ASes. We simply refer to routers exposed from LISP-enabled networks without considering the details of which ASs they are from assuming that they have service level agreements while the TEC's role is to optimize network use.

We divide all routers in LISP-enabled networks into two categories: regular routers and RLOC routers. By RLOC routers, we mean the routers that are put into one or more groups for traffic proportioning.

There are three possible cases for balancing traffic:

- (I) *Base Case*: traffic destined for a router in a group is *not* proportioned among its peers in that group. Consider again Fig. 1 where we consider RLOC_{2a} and RLOC_{2b} to be in a group. In the base case (Case-I), traffic to RLOC_{2a} from other routers is not proportioned to RLOC_{2b} and vice-versa.
- (II) *Limited Balancing*: An RLOC router in a group is identified as the *primary* router (for destination), and the others are its *secondary* routers/peers in this specific group.
- (III) *Flexible Balancing*: all RLOC routers in a group can serve as primary routers (for destination) and traffic for

any router in the group may be proportioned to its peers. That is, traffic for RLOC_{2a} may be proportioned to RLOC_{2b} and vice-versa.

In our notation, a group may be identified through one of its member routers. If router d belongs to a group, then G^d identifies its group. For example, if router 1 and router 2 are in Group-A and router 3 and router 4 are in Group-B, then G^1 or G^2 will reflect Group-A and G^3 and G^4 will reflect Group-B. That is $G^1 = \{1, 2\} = G^2$ and $G^3 = \{3, 4\} = G^4$.

Let \mathcal{W} be the super set of the individual sets of routers in a group (G^d). That is, \mathcal{W} is the union of all the individual groups identified through any RLOC router d (i.e., G^d s). In the above example, $\mathcal{W} = \{\{1, 2\}, \{3, 4\}\}$; the set does not have duplicate elements, so there is only one occurrence of $\{1, 2\}$ and $\{3, 4\}$. Let G^* be the set of the individual elements of all such routers that are part of any group (G^d); in the example, $G^* = \{1, 2, 3, 4\}$. To sum up, in a set notation format, $\bigcup G^d = \mathcal{W}$, where $d \in G^*$. We assume that a router in one group is not present in another group. The set of primary routers belonging to different groups is labeled as \hat{R} . We label the set of source LISP routers as S and the set of destination ones as D . Notations are summarized in Table I.

With this background, we now present the problem formulations for Case-II and Case-III that can be solved by the traffic engineering calculator in a federated environment. In discussing Case-II, we will point out how this case reverts to Case-I (base-TE) when the secondary is not enabled. Without loss of generality, for ease of comparison among the three cases, we assume that there is no intra-group traffic.

A. Formulation for Case-II: Limited Balancing

Conservation of traffic flows from the source routers to the destination routers leads to three situations. First,

$$\sum_{p \in P_{sd}} x_{sd}^p = h_{sd}, \quad s \in S, d \in D \setminus \hat{R}, s \neq d \quad (1)$$

In (1), the destination nodes are the ones that are not in the set of primary routers (\hat{R}). In this case, the demand destined for the secondary routers is not proportioned with their peers in the group in Case-II. Here, P_{sd} represents a possible set of paths.

Secondly, the traffic flow to the destination primary router in a group may be proportioned to α_{sd}

$$\sum_{p \in P_{sd}} x_{sd}^p = h_{sd} \alpha_{sd}, \quad s \in S \setminus G^d, d \in \hat{R} \quad (2)$$

where α_{sd} (a variable) is the proportion of traffic from source router s and destination router d . Note that in (2), the destination router belongs to a set of primary routers; this is why demand may be proportioned to α_{sd} and the source router belongs to the routers outside the group of the primary routers.

Finally, the remaining proportion of traffic between s and d (fractioned by α_{sd}^t or $1 - \alpha_{sd}$) is sent to the peers of the primary router(s) in the same group

$$\sum_{p \in P_{st}} z_{st}^p = h_{sd} \alpha_{sd}^t, \quad s \in S \setminus G^d, t \in G^d \setminus \{d\}, d \in \hat{R} \quad (3)$$

TABLE I: List of Notations

S	= Set of source routers
D	= Set of destination routers
G^*	= Set of all routers belonging to all the groups of routers for traffic proportioning
G^d	= Set of all routers in a group where router d belongs in, usually destination router d
\mathcal{W}	= Union set of the individual set of groups G^d
L	= set of links in the network
\hat{R}	= Set of primary routers from group routers
h_{sd}	= traffic demand between s and d
C_ℓ	= capacity of link l - m
P_{sd}	= Set of paths from source router s to destination router d
$\delta_{sd,\ell}^p$	= link-path indicator: 1 if path p for demand from source s to destination d uses link ℓ ; 0, otherwise
Variables:	
x_{sd}^p	= non-negative flow variable s to d on path p
z_{st}^p	= non-negative flow variable z from s to t on path p when t is a secondary router
α_{sd}	= fraction of demand between source and primary router where d is a primary router ($0 \leq \alpha_{sd} \leq 1$)
α_{sd}^t	= fraction of demand between source s and primary router d sent to secondary router t in the same group ($0 \leq \alpha_{sd}^t \leq 1$)
y_ℓ	= link flow variable for link ℓ
r	= maximum link utilization variable

In (3), the flow variable z is the proportioned flow of the primary router to its secondary routers in a group. Thus, the demand it will carry will be for the primary router but the path it takes is to the secondary routers in its group. Here, α_{sd} is the proportion of traffic meant for the primary router directly sent to it and α_{sd}^t is the rest of the traffic proportion to its secondary routers in the group, shown in (4).

Next, consider the critical set of constraints that address traffic proportioning. Specifically, (4) below reflects that the proportion of traffic to the primary router and to its associated secondary router must add up to 1

$$\alpha_{sd} + \sum_{t \in \hat{R}_{sd}} \alpha_{sd}^t = 1, \quad s \in S \setminus G^d, d \in \hat{R}, t \in G^d \setminus \{d\} \quad (4)$$

Now consider the flow on a link. The link flow on link ℓ is to account for all traffic flow variables that use the link and is given by

$$\sum_{s \in S, d \in D \setminus \hat{R}, s \neq d} \sum_{p \in P_{sd}} \delta_{sd,\ell}^p x_{sd}^p + \sum_{s \in S \setminus G^d, d \in \hat{R}} \sum_{p \in P_{sd}} \delta_{sd,\ell}^p x_{sd}^p + \sum_{s \in S \setminus G^d, d \in \hat{R}} \sum_{t \in G^d \setminus \{d\}} \sum_{p \in P_{sd}} \delta_{sd,\ell}^p z_{st}^p = y_\ell, \quad \ell \in L \quad (5)$$

Here $\delta_{sd,\ell}^p$ is the link-path indicator, which is set to 1 if the flow variable between s and d on path p exists on link ℓ and 0, otherwise.

The objective for our formulation is to obtain optimize network utilization by minimizing the maximum link utilization, which has been commonly used for IP networks [8]. Because of traffic distribution with multiple RLOCs, this is also a reasonable objective to consider for our problem. Using r as the maximum link utilization variable, we relate this to the link capacity with flow on link ℓ as follows

$$y_\ell \leq C_\ell r, \quad \ell \in L \quad (6)$$

so that we can minimize r ; i.e., the goal is to

$$\text{Minimize } r \quad (7)$$

In summary, (1)-(7) represent the formulation for the LISP-enabled traffic engineering problem. Note that in this model, α , which signifies traffic being proportioned, is also a variable.

With the above model, the base-case traffic engineering (Case-I, base-TE), where traffic is not proportioned to RLOCs in a group, becomes a special case when we set $\alpha_{sd} = 1$ in (4). This implies that $\alpha_{sd}^t = 0, s \in S \setminus G^d, d \in \hat{R}, t \in G^d \setminus \{d\}$.

B. Formulation for Case-III: Flexible Balancing

The main difference between Case-II and Case-III is that traffic may be fractioned among all the routers in a group, instead of proportioning it from the primary router to a secondary router in the group. Here, with the flow conservation equation for this case is:

$$\sum_{p \in P_{sd}} x_{sd}^p = h_{sd}, \quad s \in S, d \in D \setminus G^*, s \neq d \quad (8)$$

In (8), the destination routers are routers that do not belong to any group. Then, the traffic flow to routers in a group may be split among peers in the group and proportioned to α_{sd}

$$\sum_{p \in P_{sd}} x_{sd}^p = h_{sd} \alpha_{sd}, \quad s \in S \setminus G^d, d \in G^*, \quad (9)$$

while the remainder of the traffic (fractioned by α_{sd}^t) is sent to its peers of the same group

$$\sum_{p \in P_{st}} z_{st}^p = h_{sd} \alpha_{sd}^t, \quad s \in S \setminus G^d, t \in G^d, d \in G^*, d \neq t \quad (10)$$

Similarly, in (9), where all the routers are primary, so can the destination router, belonging to any group of routers (G^*) may be proportioned and (10) reflects the remaining portion of the traffic going to the peers.

The portion of traffic to a router and its peers in a group satisfies the following equation:

$$\alpha_{sd} + \sum_{t \in R_{sd}} \alpha_{sd}^t = 1, \quad s \in S \setminus G^d, t \in G^d, d \in G^*, d \neq t \quad (11)$$

The traffic flow on link $\ell \in L$ is

$$\begin{aligned} & \sum_{s \in S, d \in D \setminus G^*, s \neq d} \sum_{p \in P_{sd}} \delta_{sd, \ell}^p x_{sd}^p + \sum_{s \in S \setminus G^d, d \in G^*} \sum_{p \in P_{sd}} \delta_{sd, \ell}^p x_{sd}^p \\ & + \sum_{s \in S \setminus G^d, d \in G^*} \sum_{t \in G^d \setminus \hat{R}, d \neq t} \sum_{p \in P_{st}} \delta_{sd, \ell}^p z_{st}^p = y_{\ell}, \quad \ell \in L \quad (12) \end{aligned}$$

Similar to Case-II, the objective is to minimize the maximum link utilization (r), so the goal is to

$$\text{Minimize } r \quad (13)$$

By comparing the models for Case-II and Case-III, we can see that how the sets defined can be used to distinguish between the primary and other routers in a group or how to peer among all routers in a group.




Fig. 3: Topology A4

TABLE II: Uniform Demand, Uniform Capacity

(a) base-TE		(b) LISP-TE/II			
Capacity	r	Capacity	r	α_{13}	α_{23}
300	0.63	300	0.63	1	1

TABLE III: Uniform Demand, Capacity adjusted on link 1-3

(a) base-TE		(b) LISP-TE/II		
C_{13}	r	C_{13}	r	α_{13}
200	0.760	200	0.713	0.875
100	0.950	100	0.814	0.714
80	1.000	80	0.838	0.676

C. Illustrative Example

Consider the 4-node topology A4 shown in Fig. 3. It has 4 routers where routers 3 and 4 are considered in a group by considering 3 as the primary router and 4 as the secondary router. Here, $G^3 = \{3, 4\} = G^4, \mathcal{W} = \{\{3, 4\}\}$ and $G^* = \{3, 4\}$. Then, equation (4) reduces to the following two equations:

$$\alpha_{13} + \alpha_{13}^4 = 1 \quad (14)$$

$$\alpha_{23} + \alpha_{23}^4 = 1 \quad (15)$$

We first set the traffic demand to 190 for all pairs (except for no traffic due to our assumption between node 3 and node 4, which form a group) and we set link capacities to 300 for all links for Topology-A4.

In Table II, we show the results for base-TE and LISP-TE/II. We observe that the maximum link utilization (r) is 63% for both cases. For LISP-TE/II, $\alpha_{13} = \alpha_{23} = 1$. This means that we do not see any traffic between routers 1 and 3 proportioned to router 4 and similarly, for traffic between routers 2 and 3. Thus, in this symmetric traffic/capacity example, there is no benefit of LISP-TE/II over base-TE.

We next change the capacity of link 1-3. The results for base-TE and LISP-TE/II with a number of different capacity values for link 1-3 are shown in Table III. When the capacity of link 1-3 is reduced, the optimal value of r increases, which is as expected. The optimal value of α_{13} is strictly less than 1 when the capacity of link 1-3 is reduced, thus implying that some traffic destined for router 3 is diverted to the secondary router 4 as link capacity 1-3 decreases. In other words, this illustrates the benefit of LISP-TE/II over base-TE with this non-symmetric network case.

TABLE IV: Maximum link utilization results for Topology A12.

Scenario	Capacity	Traffic ($s-d$)	Total Traffic	RLOC Traffic	r (Case-I) (base)	r (Case-II) (limited)	r (Case-III) (flexible)
A12:1	10,000	1,000 (uniform)	124,000	47,000	0.583	0.583	0.583
A12:2		1.5 times (RLOCs)	+ 23,500	+ 23,500	0.792	0.792	0.604
A12:3		2.0 times (RLOCs)	+ 47,000	+ 47,000	1.000	1.000	0.750
A12:4	-5,000 (RLOCs)	1,000 (uniform)	124,000	47,000	1.000	1.000	0.786
A12:5		1.5 times (RLOCs)	+ 23,500	+ 23,500	1.357	1.357	1.036
A12:6		2.0 times (RLOCs)	+ 47,000	+ 47,000	1.750	1.750	1.286
A12:7	10,000	Generated/Non-Uniform	44,888	17,386	0.219	0.219	0.219
A12:8		1.5 times (RLOCs)	+ 8,678	+ 8,678	0.304	0.304	0.227
A12:9		2 times (RLOCs)	+ 17,386	+ 17,386	0.389	0.389	0.286
A12:10	-5,000 (RLOCs)	Generated/Non-uniform	44,888	17,386	0.376	0.376	0.287
A12:11		1.5 times (RLOCs)	+ 8,678	+ 8,678	0.521	0.521	0.389
A12:12		2 times (RLOCs)	+ 17,386	+ 17,386	0.668	0.668	0.490

TABLE V: Maximum link utilization results for Topology A26.

Scenario	Capacity	Traffic ($s-d$)	Total Traffic	RLOC Traffic	r (Case-I) (base)	r (Case-II) (limited)	r (Case-III) (flexible)
A26:1	10,000	100 (uniform)	63,200	20,700	0.568	0.568	0.503
A26:2		1.5 times (RLOCs)	+ 10,350	+ 10,350	0.700	0.700	0.535
A26:3		2.0 times (RLOCs)	+ 20,700	+ 20,700	0.850	0.850	0.568
A26:4	-5,000 (RLOCs)	100 (uniform)	63,200	20,700	0.843	0.825	0.660
A26:5		1.5 times (RLOCs)	+ 10,350	+ 10,350	1.163	1.163	0.690
A26:6		2.0 times (RLOCs)	+ 20,700	+ 20,700	1.500	1.500	0.825
A26:7	10,000	Generated/Non-Uniform	65,218	22,215	0.591	0.591	0.517
A26:8		1.5 times (RLOCs)	+ 11,055	+ 11,055	0.732	0.732	0.553
A26:9		2 times (RLOCs)	+ 22,215	+ 22,215	0.890	0.890	0.591
A26:10	-5,000 (RLOCs)	Generated/Non-Uniform	65,218	22,215	0.868	0.847	0.681
A26:11		1.5 times (RLOCs)	+ 11,055	+ 11,055	1.177	1.177	0.719
A26:12		2 times (RLOCs)	+ 22,215	+ 22,215	1.526	1.526	0.855




Fig. 4: Topology A12 (groups are shown with shadowed ovals)




Fig. 5: Topology A26 (groups are shown with shadowed ovals)

IV. STUDY CASES

For our study, we considered two different moderate-size topologies. All computational works were performed using IBM ILOG CPLEX as the optimization tool. The problem formulations presented in Section III-A and Section III-B are linear programming formulations. Thus, a standard package is extremely efficient to solve our formulations in just a few seconds; no specialized algorithms are needed. Furthermore, since such balancing acts would need to be performed only periodically (such as a few times a day), using a standard tool is viable.

The topologies have 12 and 26 nodes (see Fig. 4 and Fig. 5). For both the topologies, we consider uniform as

well as non-uniform traffic, where the non-uniform traffic was generated using the demand general model given in [5]. We chose RLOC groups for A12 to be $\{6, 8\}$ and $\{9, 10, 11\}$. For A26, RLOC groups were $\{16, 17, 18\}$, $\{20, 21, 22\}$, and $\{24, 25, 26\}$. Routers in each group are selected based on their proximity to each other.

The results are summarized in Table IV and Table V for A12 and A26, respectively. We first clarify various information contained in the tables. Under the *Capacity* column, 10,000 refers to capacity uniformly fixed at 10,000 on all links in the network; $-5,000$ (RLOCs) means that the link capacity from any router directly connected to the RLOC routers in all the

groups (G^*) is reduced to 5,000 while all the remaining link capacities are kept fixed at 10,000. Under the *Traffic* column: (i) *uniform* means that the traffic volume between all routers in the network are uniformly fixed with the value specified being the traffic for each source-destination (s-d) pair; (ii) *1.5 times (RLOCs)* means that *only* the traffic volume to the destination RLOC routers in groups (i.e., G^*) are increased by 1.5 times from the original traffic to those routers; similarly, *2 times (RLOCs)* means that only the traffic is increased 2 times the same way (i.e., to G^*); (iii) *Generated/Non-Uniform* means that non-uniform traffic is generated for origin-destination pairs, using the traffic generation model from [5]. Under the *Total traffic* column, “+ n ” means (n) traffic units have been added to the original total traffic in the network due to the scenario considered. Similarly, under the *RLOC Traffic* column, “+ n ” means (n) traffic units have been added to the original traffic to the RLOCs in the network due to the scenario considered. Columns r (*Case-I*), r (*Case-II*) and r (*Case-III*), show the optimal value of the TE objective, r , for the different cases. Optimal utilizations are shown for the most drastic gains for Case-III (flexible balancing) compared to base case and Case-II for non-uniform traffic with capacity asymmetry (i.e., links to RLOCs reduced by 5,000 units) in Fig. 6.

Link capacities are reduced to RLOCs in order to observe how the optimal link utilization is affected when we impose a tighter bound to those links to force the traffic to be further proportioned among the RLOC routers in a group. Usually, increasing the capacities would allow more traffic to flow and we might not get the advantage that we observe with minimizing the maximum link utilization in a tighter bound scenario. Similarly, we increased the traffic going to those routers that belong to all the group routers (G^*) to observe the response of link utilization, when the overall traffic to RLOCs is relatively more dominant compared to the total network traffic. In other words, we want to observe how much better our LISP-TE model is when there are high traffic demands to those routers. If LISP-TE is better than the base case under such scenarios, then this may also serve as an indicator that routers that have a high inflow of demands may be grouped in a LISP scenario for a better network utilization.

Optimality for A12:: From Table IV, we see that for the uniform load case with uniform capacity throughout the network, there are no differences between base-TE, LISP-TE with Case-II (LISP-TE/II), and LISP-TE with Case-III (LISP-TE/III) since for these situations, the optimal value of r remains the same. This is the same observation we noted earlier with topology A4.

When we increase the demands to the RLOCs, LISP-TE gives us a better (minimum) r with the flexible balancing case (Case-III) compared to the base case (Case-I) and Case-II. When we decrease the link capacities of the links to the RLOCs, under both generated or fixed uniform demands or non-uniform demands (by increasing the demands to the RLOCs), we see a better value of r for LISP-TE for Case-III than for Case-I and Case-II. Note that for one scenario we considered, all three cases have an infeasible value of

r (i.e., optimal $r > 1$) suggesting that if we increase the traffic for RLOCs too high with capacities for links to RLOCs relatively low, then we may get infeasible solutions. While practical networks would not allow infeasibility, this scenario was studied to see the implications.

For non-uniform traffic with most constrained asymmetric capacity (i.e., links to RLOCs reduced by 5,000 units of capacity) marked as Scenarios A12:10, A12:11, and A12:12 in Table IV, optimal utilizations are shown for Case-III (flexible balancing) compared to base case and Case-II in Fig. 6(a). We see that the gain for Case-III compared to the other case is from 31% to 36%.

Optimality for A26:: Unlike in A12, for non-uniform (Generated) and uniform demands with fixed capacities, we see noticeable gain in the optimality value for Case-III suggesting that it is beneficial for LISP-TE/III for A26. This also gives us an insight that although the network topology does influence optimal r , LISP-TE with Case-III has a better gain than the base case. Similar to the discussion for A12, under fixed capacities when we increase the demand we get a better (minimal) value of r with Case-III compared to Case-I. In most cases, Case-II is the same as Case-I. When we decrease the link capacities of the links to the RLOCs, under both generated or fixed uniform demands or non-uniform demands (by increasing the demands to the RLOCs), we see a better value of r for LISP-TE for Case-II and Case-III than Case-I.

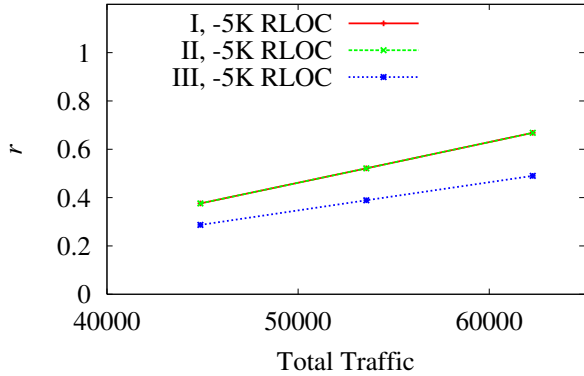
For non-uniform traffic with most constrained asymmetric capacity (i.e., links to RLOCs reduced by 5,000 units of capacity) marked as Scenarios A26:10, A26:11, and A26:12 in Table V, optimal utilizations are shown for Case-III (flexible balancing) compared to base case and Case-II in Fig. 6(b). We see that the gain for Case-III compared to the other case is from 27% to 78%.

V. LISP-TE: REMARKS

We now make a few remarks on LISP-enabled TE.


- Our TE model uses continuous variables since the scope is to determine the ensemble percentage of traffic steering. Secondly, in actual practice, traffic steering could depend on which IP prefixes are to be used based on the ensemble percentage from the the LISP-TE model output.
- Our TE model does not include time-based traffic changes, not necessary since no temporal coupling is needed. That is, our model can be rerun if traffic matrices change significantly to re-determine steering percentage.
- The use of the TE calculator does not mean the actual end-to-end traffic flows need to be fully defined. Note that EIDs are also in different ASes (not shown in Fig. 1). Based on our approach, conventional Internet TE techniques such as path prepending, AS egress routing optimization, multi-exit discriminator (MED) could still be used along with IP prefixes for traffic steering.
- LISP TE is to connect cooperating ASes, which would in general result in smaller topologies. The main benefit of LISP-TE was observed when traffic and network capacity

A12, non-uniform traffic / Asymmetric Capacity



(a) Scenarios: A12:10, A12:11, A12:12

A26, non-uniform traffic / Asymmetric Capacity



(b) Scenarios: A26:10, A26:11, A26:12

Fig. 6: Optimal r : non-uniform traffic, asymmetric capacity

were nonuniform. Thus, we posit that studying large topologies would not change this main benefit.

- Partnerships between ASes are not uncommon through SLAs resulting in sharing many details. Furthermore, the TE calculators could be placed at an Internet Exchange Point (IXP), where ASes who are connected at this IXP agree on SLAs about traffic etc.
- LISP-TE traffic proportioning can be better controlled and ensured also by using Re-encapsulating Tunnel Routers (RTRs), for relay routing across RLOCs through LISP Canonical Address Format (LCAF) extension [4], [13]. Indeed, as traffic fluctuates over times, disposing of programmable RTRs can ease adjusting the traffic proportioning in real-time.
- Programmable networking could also be used for LISP-TE. This implies an external control-plane module configuring programmable device as for instance a P4 switch [16], e.g. using P4 runtime, NETCONF or other forms of south bound Interface using the load-balancing information computed by our model.

Some of the above remarks are regarding LISP-TE implementation. While this is not in the scope of this paper, these

are mentioned to show how it is possible to enable LISP-TE.

VI. CONCLUSION

In this paper, we present two formulations of the federated LISP-enabled network traffic engineering problem, one with limited balancing and another with flexible balancing. To our knowledge, this is the first work on optimization formulations for LISP-TE with a study on traffic balancing to multiple routers for federated LISP-enabled networks.

Our study shows that when capacity is most constrained with non-uniform traffic, the flexible balancing model (Case-III) provides the most benefit, ranging from 27% gain to 78% gain. Other objectives than the minimizing the maximum link utilization can also be studied; while the actual solution would change for any different objective for the LISP-TE problem, the primary observation will remain the same since there are no changes in the set of constraints.

REFERENCES

- [1] A. Cabellos and D. Saucez, "An architectural introduction to the locator/ID separation protocol (LISP)," *IETF, RFC 9299*, 2022.
- [2] J. Chavula, M. Densmore, and H. Suleman, "Reducing latency in African NRENs using performance-based LISP/SDN traffic engineering," *University of Cape Town*, 2015.
- [3] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "Locator/ID separation protocol (LISP)," *IETF, RFC 6830*, 2013.
- [4] D. Farinacci, D. Meyer, and J. Snijders, "RFC 8060: LISP canonical address format (LCAF)," *Internet Engineering Task Force*, 2017.
- [5] B. Fortz and M. Thorup, "Internet traffic engineering by optimizing OSPF weights," in *Proc. of IEEE INFOCOM 2000*, 2000, pp. 519–528.
- [6] Y. Li, S. Guo, Q. Song, Y. Wang, X. Wei, and J. Ma, "Mapping prediction with recurrent neural networks for future LISP enabled networks," *J. of Information and Intelligence*, vol. 1, pp. 134–147, 2023.
- [7] H. D. D. Nguyen and S. Secci, "LISP-EC: Enhancing LISP with egress control," in *2016 IEEE Conference on Standards for Communications and Networking (CSCN)*. IEEE, 2016, pp. 1–7.
- [8] A. Nucci, B. Schroeder, S. Bhattacharyya, N. Taft, and C. Diot, "IGP link weight assignment for transient link failures," in *Proc. 18th International Teletraffic Congress (ITC18)*, Berlin, Germany, Sep. 2003, pp. 321–330.
- [9] C.-D. Phung, M. Coudron, and S. Secci, "Internet acceleration with LISP traffic engineering and multipath TCP," in *21st Conf. Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, 2018.
- [10] B. Quoitin and O. Bonaventure, "A cooperative approach to interdomain traffic engineering," in *Next Generation Internet Networks, 2005*, april 2005, pp. 450 – 457.
- [11] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig, "Interdomain traffic engineering with BGP," *IEEE Communications Magazine*, vol. 41, no. 5, pp. 122 – 128, may 2003.
- [12] B. Quoitin, L. Iannone, C. de Launois, and O. Bonaventure, "Evaluating the benefits of the locator/identifier separation," in *Proceedings of 2nd ACM/IEEE international workshop on Mobility in the evolving internet architecture*. New York, NY, USA: ACM, 2007, pp. 5:1–5:6.
- [13] A. Rodriguez-Natal, V. Ermagan, A. Smirnov, V. Ashtaputre, and D. Farinacci, "RFC 9306: Vendor-specific LISP canonical address format (LCAF)," *Internet Engineering Task Force*, 2022.
- [14] D. Saucez, B. Donnet, L. Iannone, and O. Bonaventure, "Interdomain traffic engineering in a locator/identifier separation context," in *Internet Network Management Workshop, 2008.*, Oct. 2008, pp. 1–6.
- [15] S. Secci, K. Liu, G. K. Rao, and B. Jabbari, "Resilient traffic engineering in a transit-edge separated internet routing," in *Proc. of 2011 IEEE Int. Conference on Communications (ICC 2011)*, June 2011.
- [16] B. Steinert, M. Hberle, J.-O. Nick, D. Farinacci, and M. Menth, "P4-LISP: A P4-based high-performance router for the locator/identifier separation protocol," in *2023 IEEE 9th International Conference on Network Softwarization (NetSoft)*, 2023, pp. 89–97.