Resource Provisioning for Dynamic Multi-domain WDM Networks: Effectiveness and Fairness

Xiaolan J. Zhang
Coordinated Science Laboratory
Department of ECE
Univ. of Illinois at Urbana-Champaign
Email: xzhang29@illinois.edu

Sun-il Kim
Info. and Technology Systems Center
Department of Computer Science
Univ. of Alabama in Huntsville
Email: sunilkim@cs.uah.edu

Steven S. Lumetta
Coordinated Science Laboratory
Department of ECE
Univ. of Illinois at Urbana-Champaign
Email: lumetta@illinois.edu

Abstract—**In this paper, we approach new problems of resource dimensioning and routing for multi-domain networks supporting future user-driven dynamic traffic. Three basic dimensioning approaches–independent, global, and normalized–are studied in conjunction with the appropriate routing schemes that match each dimensioning policy based on how much information is shared across domains. We quantify the performance of inter-domain and intra-domain traffic while considering changes in traffic loads as well as network rescaling. We introduce the notion of fairness, which raises questions about how network domains with separate ownership and possibly conflicting economic interests should decide on co-providing bandwidth-on-demand services. The results motivate further research in fair, multi-domain resource provisioning, and highlight the need for new cost-models that are critical in designing multi-domain provisioning and peering policies.**

I. INTRODUCTION

The ability to dynamically establish high-bandwidth optical network connections is increasingly important to support applications. Given the diverse geographic locations of end-users and service providers, many of these requests will require dynamic provisioning across multiple network domains with varying protocols and standards (as well as ownership and management policies). Interoperability between domains has received substantial attention in recent years. In contrast, few studies have examined the problem of dimensioning network resources such as transceivers to support a mix of inter-domain and intra-domain traffic routed dynamically through the domain. In this paper, we examine the relationship between the amount of network information shared between two domains and metrics of cost, such as blocking probability and capacity, demonstrating that more complete information can provide substantial benefits. The degree to which today's pre-negotiated agreements, possibly based on inferences about the peer networks, can match these benefits, we leave as the subject of future work.

In order to illustrate the main idea that generally applies on many existing network architectures and cost models, we assume in this paper a network model with full wavelength conversion at each node and uniform link cost. The dimensioning problem is to allocate the number of wavelength channels (or capacity) on each link to support dynamically routed node-to-node requests (here "provisioning" is used interchangeably with "dimensioning"). We define path cost as the number of hops. Online state-based routing algorithms are used to route connection requests. Our study is a first step toward the determination of equipments, such as switches, amplifiers and transceivers, with specific defined costs/weights on practical multi-domain architectures, considering reachability, wavelength-continuity, and survivability.

Dimensioning in a single domain [1] has shown that an unbalanced network can artificially obscure the differences in performance between various routing algorithms and that a dynamically routed network must be dimensioned properly in order to take full advantage of the overall capacity set in place. When considering a multi-domain network, each domain should not only be dimensioned for its own intra-domain (internal or local) traffic, but must also provide support for inter-domain (external) traffic. Resource usage in each domain becomes more complicated. As resources internal to each network are shared between intra- and inter-domain connections, poorly dimensioned/managed networks can adversely impact the resources available for routing intra-domain traffic in other domains. Although such effects can be (and currently are) limited to some extent through peering agreements, adding such constraints cannot match the performance available in a singly-managed domain, and their negative impact on performance in the presence of dynamic inter-domain traffic is unclear. Due to the interdependence between domains and resources, simply dimensioning each domain independently produces poor results. Improving these results requires a better understanding of the interaction between the amount and type of information shared between domains (for dimensioning and routing) and the impact of this sharing on overall performance.

Routing through multiple domains, especially when the domains are bound by conflicting economic interests, raises questions about fairness issues that require careful attention. A simple example of a conflict between two domains is shown here. Figure 1 is a two-domain network connecting NSFNET and ARPANET (these topologies are only used for illustration. Real domain networks can have arbitrary topolo-

This work was made possible with the generous donations from Intel and support from the Information Trust Institute of the University of Illinois at Urbana-Champaign and the Hewlett-Packard Company through its Adaptive Enterprise Grid Program.

gies using different control planes, such as GMPLS/ASON, SONET/Ethernet over WDM, etc). Each network assumes three border nodes (marked in solid color) connected to the borders of the other domain. Various sub-domains may exist (in different format depending on the whether the overall network has a hierarchical structure), but for illustration purposes we use a simple two-domain scenario. Between the node pair (A, B), we have two choices of global shortest paths denoted path 1 and path 2. Assume uniform link cost, path 1 favors ARPANET because less resources are used on the ARPANET side. However, path 2 favors NSFNET. When both paths are available, which network should bear the extra cost to establish a connection from A to B? This question not only requires a closer look at performance tradeoffs, but motivates cost models for dimensioning and routing policies that take into consideration the notion of fairness. At the same time, fairness can subjective, but it is dependent on the economic model that each domain must agree on. The question of building an mutually attractive cost model between peer domains cannot be solved without understanding the impact of various routing and dimensioning choices made. As networks become fully autonomous at the WDM layer, these problems will become more important.

In this paper, we quantify the performance of three network dimensioning approaches paired with appropriate routing algorithms. We propose a fairness measure to capture the penalty of inter-domain traffic on each network. We utilize the two basic routing schemes presented by [2], End-to-End global shortest path routing and source initiated concatenated shortest path routing (equivalent to single node for two domain case), to illustrate potential fairness issues in peer-viewed domains. We also introduce a normalized (and more fair) routing scheme while considering limited information sharing, and show the impact of network scaling as well as traffic load deviations. Our study of multi-domain provisioning for dynamic traffic motivates a new direction in network design where fairness in network operation cost and benefit is considered.

II. BACKGROUND

Recent progress in new standards development and multi-domain control plane design improves support for multi-domain services at the optical layer and introduces new problems that arise in designing algorithms and evaluating performance of multi-domain provisioning, routing and management [3], [4]. Network domains are divided based on different reasons from management and geographic locations to vendor-specific component technologies. Multi-domain networks can be defined by joining of peer individual networks (eg. single routing areas (RA)) through a set of connected border nodes (external network-network interface (E-NNI) as defined in OIF or edge routers (ER) in GMPLS networks). Domains can also be defined hierarchically corresponding to multiple layers of the optical network that have been historically implemented in many carrier networks [5]. Many studies have been devoted to designing control and signaling platforms to enable inter-domain path computation. [6], [7] proposed path computation schemes with various QoS constraints in a multi-domain context. A network service plane has been developed to integrate diverse transport network systems [8]. Path computation element (PCE) framework has been extended to support cross-domain shortest path selection [9]. [10] proposed a framework for wavelength path establishment using a ranking database.

Performance study with different levels of information sharing has been the main research focus. Sharing complete network state information across multi-domain can be impractical for two reasons: the scale of the aggregated domain exceeds the signaling capacity to flood all link states and information such as the complete network topology is often considered proprietary/private. With regards to the scalability issue, [11] proposed an efficient information aggregation and updating system for large scale networks. [12] is a theoretical study on minimal information needed between network domains to route within a tolerable error rate. Topology aggregation techniques have been used to reduce the amount of information shared between domains, often utilizing virtual topologies to provide more compact and abstract information. [2] compared the performance of global shortest path, concatenated SPF (similar to source routing in other studies), and single node aggregation/hierarchical routing (equivalent to stitched path routing). They also proposed three different border selection criteria for source routing: random selection, closest border, and least loaded path to border. [13] further proposed a quantitative study on hierarchical routing performance by using simple node (same as single node), full mesh and symmetric star aggregation schemes.

The study of routing algorithms generally falls into three categories with different levels of information sharing: i. End-to-End global shortest path routing (complete information sharing, though not necessarily needed as discussed in later sections). ii. Source initiated concatenated shortest path routing that searches for a path domain by domain starting from the source (minimum information sharing). Crankback signaling may be used to re-attempt a failed search in a segment. iii. Hierarchical/hybrid routing with various topology aggregation techniques (and thus varying amount of information shared depending on topology exposure). Generally, the performance improves as more information is shared, but the scalability problem increases (in addition to inter-domain privacy issues).

III. MULTI-DOMAIN NETWORK DIMENSIONING

A dynamic network domain must provision enough capacity (or wavelength channels) on each link (or hop) to support both inter-domain and intra-domain traffic. Inter-domain traffic is routed through its border nodes (selection of border nodes critical to the performance and fairness for both domains), introducing additional loads between internal nodes to border nodes in each network domain. In this section, we discuss general multi-domain dimensioning algorithms. For illustration purposes, all discussions are presented for two-domain scenarios, but our results can be readily extended to an arbitrary number of domains using standard information sharing/hiding

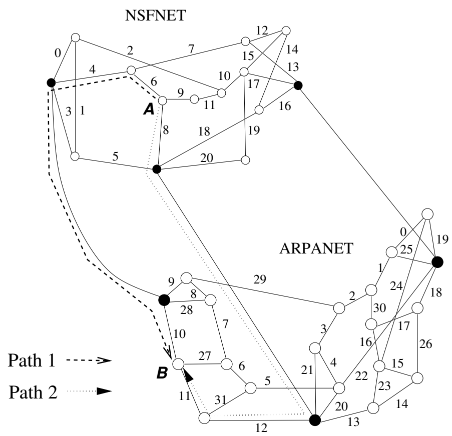

Fig. 1: NSFNET-ARPANET joint topology (link # is shown on each link).

TABLE I: Network notation

N	Set of nodes in the local network.
N'	Set of border nodes in the local network. $N' \subset N$
S	Set of nodes in the foreign network.
S'	Set of foreign border nodes. $S' \subset S, \lvert S'\rvert = \lvert N'\rvert$
E	Set of links in the local network.
C_e	Total capacity of link $e \in E$ provisioned for internal traffic.
C'_e	Total capacity of link $e \in E$ provisioned for for external traffic.
C^I_b	Total capacity of inter-domain link connecting border nodes $b \in N', b' \in S'$ provisioned for external traffic.
R	Set of all intra-domain connection request pairs. $R = N \times N$.
λ	Mean arrival rate for internal traffic.
λ'	Mean arrival rate for external traffic.
μ	Mean departure rate for internal traffic.
μ'	Mean departure rate for external traffic.
κ	Mean capacity request for internal traffic.
κ'	Mean capacity request for external traffic.
TSL_i	Topological shortest path length (TSL) for pair $i \in R$
T	Internal traffic matrix
T'_L	Equivalent external traffic matrix on the local network
$p_{n,s,b}$	The probability that border b is used to connect a local node $n \in N$ to a foreign node $s \in S$. $\sum_{b \in N'} p_{n,s,b} = 1$.

techniques such as topology aggregation or virtual topology abstraction.

In our study, we assume that every node in each domain will initiate a connection request to all other nodes in the same domain and to all other nodes in the other domains (equivalent to a global full mesh demand). Domains are joined through their border nodes, where each border node is connected to a border in in the other domain. In our model, we assume that external and internal traffic share the wavelength resources provisioned on each link. In the two-domain case, the numbers of border nodes in both domains are the same. There is a link connecting each pair of border nodes (as shown in Figure 1). In practice, it is possible for some of the border node pairs to be situated in the same building/site. In such cases, they may be viewed logically as a single node with finite capacity (equal to the inter-domain link capacity). The mapping between border nodes between two domains may also not be one-to-one. However, such variations do not affect how dimensioning and routing techniques are designed.

Dimensioning for external traffic for each domain is equivalent to dimensioning for additional traffic between each pair of internal nodes to border nodes. The network parameters are defined in Table I. We separate the capacity for external and internal traffic load for the purpose of analysis, but , the resources are still shared in our study model. We also define the capacity C^I_b for the inter-domain links that connect border node pairs. Each dynamic connection pair (internal or external) has arrival/departure/capacity rates defined.

The intra-domain traffic matrix is defined the same way as in a single domain ($T = \{(\lambda_i, \mu_i, \kappa_i) | i \in R\}$). Here, λ_i is the arrival rate drawn from a distribution of averaged value λ for request i in a particular traffic matrix. μ_i and κ_i mean the same. The external traffic generated from each node n (including border nodes themselves) is an aggregated arrival rate $\sum_{i \in \{n\} \times S} \lambda'_i$ that is effectively augmented on the rate from n to all border nodes split based on probability $p_{n,s,b}$. The equivalent external traffic matrix (on the local network) is defined in Equation 1. The entire traffic matrix including both is then the sum of T and T'_L for the same pair of nodes.

The *projected load* of a network is the amount of traffic that the total given network capacity can support (without being overloaded). Equation 2 defines the projected load for the internal traffic. It is the ratio of average traffic load (stochastic arrival/departure rate times the topological shortest path lengths) to the total available network capacity. E is an operator that computes the expected value for a set. One can use single call rate rate parameters ($\lambda_i, \mu_i, \kappa_i$) to substitute the averaged rate if one particular traffic matrix is considered. *Topological Shortest Length (TSL)* is the minimum number of hops for a connection in an empty network (Obviously, TSL may be less than the length of the actual path selected in a dynamic network using available shortest paths). For external traffic, the average amount of resources used for each connection is the average shortest path lengths to all border nodes. Equation 3 defines the external traffic load. Same load computation metric will be used on all dimensioning techniques to provide a fair comparison.

Note that dimensioning each network domain to support traffic to the other domain can be the same as dimensioning a single network with an estimate of external traffic distribution on borders. Once we have the traffic matrices, the basic dimensioning algorithm proposed in [1] can be extended (Algorithm 1) to dimension each network separately. The dimensioning algorithm basically provisions capacity for each simulation request over a random selected topological shortest path (more than one shortest path exist for most requests) and uses the steady-state result to provision the network. Using a similar load estimation technique, the capacity of inter-domain links is determined according to Equation 4. By symmetry, the computed inter-domain link capacity from the other domain is the same.

$$T'_L = \Big\{ \Big(\sum_{n,s \in S} p_{n,s,b} \lambda'_{(n,s)}, \sum_{n,s \in S} p_{n,s,b} \mu'_{(n,s)}, \sum_{n,s \in S} p_{n,s,b} \kappa'_{(n,s)} \Big) \Big| \sum_{b \in N'} p_{n,s,b} = 1, n \in N, b \in N' \Big\} \tag{1}$$

$$proj\_load_{int} = \frac{\mathsf{E}_{i \in R} \left(\frac{\lambda \kappa}{\mu} TSL_i \right)}{\sum_{e \in E} C_e} \tag{2}$$

```
1: Set infinite available capacity on each link $C_e \leftarrow \infty \ \forall e \in E$
2: while System has not reached steady state do
3:     Generate a new request $i$ from one or more traffic matrices $T + T'_L$.
4:     Route $i$ by shortest path first algorithm (SPF).
5: end while        ▷ $\widetilde{C}_e$ is the actually used network capacity of each link.
6: Repeat from Line 1 to get the distribution of $\widetilde{C}_e$, mean $\mathsf{E}_l(\widetilde{C}_e)$
```

Algorithm 1: Dimensioning with internal and external traffics.

$$proj\_load_{ext} = \frac{\mathsf{E}_{n \in R, s \in S}\left(\frac{\lambda'\kappa'}{\mu'}\sum_{b \in N'}\frac{TSL_{(n,b)}}{|N'|}\right)}{\sum_{e \in E} C'_e} \tag{3}$$

$$C_b^I = \sum_{n \in N}\sum_{s \in S} p_{n,s,b}\frac{\lambda'\kappa'}{\mu'} \tag{4}$$

$$l_{b,b'}^{E2E} = TSL_{(n,b)} + TSL_{(s,b')} \tag{5}$$

$$l_{b,b'}^{nE2E} = \frac{|R|TSL_{(n,b)}}{\sum_{i \in R} TSL_i} + \frac{|S|TSL_{(s,b')}}{\sum_{i \in S \times S} TSL_i} \tag{6}$$

$$\begin{aligned} penalty = & \sum_{n \in N \setminus N', s \in S}\frac{\lambda'\kappa'}{\mu'}\sum_{b \in N'} p_{n,s,b}\frac{TSL_{(n,b)} - \min_{n,k \in N'} TSL_{(n,k)}}{\min_{n,k \in N'} TSL_{(n,k)}} \\ & + \sum_{n \in N', s \in S}\frac{\lambda'\kappa'}{\mu'}\sum_{b \in N'} p_{n,s,b} TSL_{(n,b)} \end{aligned} \tag{7}$$

IV. SEVERAL ROUTING AND DIMENSIONING TECHNIQUES

In this section, we discuss three dimensioning techniques that estimates the amount of traffic loads to each border node (i.e. $p_{n,s,p}$) depending on the amount of information shared crossing domain borders and inter-domain routing agreement.

Three dynamic routing algorithms are used here. The first two are same as the algorithms found in [2]. In **source initiated concatenated shortest path routing (CSR)** , the requesting node chooses the closest border node and uses the shortest path. Starting from this border node, the path in the other domain is selected. No information (other than the destination) is shared between the border nodes. If the downstream domain cannot find a path to the destination through the selected border or the inter-border link is full, the call is rejected. [1] **End-to-End global shortest path routing (E2E)** chooses the border nodes that result in global available shortest paths. The PCE-based BRPC approach [9] implemented such global optimal routing without domain topology disclosure. Equation 5 shows the computation of shortest paths between connected border pairs in two-domain networks. The border pair is picked with minimum $l_{b,b'}^{E2E}$. **Normalized global shortest path routing (nE2E)** is similar to E2E, but each reported distance is normalized by the average TSL of each respective domain before being advertised and summed up. Equation 6 shows the computation of nE2E path length. The pair of border nodes is picked with minimum normalized distance $l_{b,b'}^{nE2E}$. nE2E is important because E2E favors the border node selection for the larger network when two domains are of different sizes. The larger domain is more likely to have longer paths to the border, their shortest path length can dominate the length of the global shortest path and force the smaller domain to pick unfavorable paths. For example, the available SPF length from a source node to three borders are 5, 7, and 11 hops in the larger domain A (average TSL=7). In the smaller domain B (average TSL=2), three corresponding borders to the destination node can be 3, 2, and 1 hops away. In E2E, the path with 5 hops in A and 3 hops in B is chosen. In nE2E, the path with 7 hops in A and 2 hops in B is chosen. Although the path picked by nE2E is longer, it is a fairer path for B because the hops used in each domain is comparable to their sizes. Fairness will be defined in the next section.

```
for each inter-domain connection pair $(n, s)$ do
    for each border node pair $(b, b')$ from each domain do
        Compute the total path length $l_{(b,b')}^{E2E}$ (or $l_{b,b'}^{nE2E}$ for NS).
    end for
    pick up the set of borders $B \in N'$ of minimal total path length.
    for each shortest path border nodes $b \in N'$ do
        $p_{n,s,b} = \frac{1}{|B|}$ for $b \in B$, $p_{n,s,b} = 0$ for $b \notin B$
    end for
end for
```

Algorithm 2: Computation of $p_{n,s,b}$ for GS/NS.

Next, we discuss the computation of $p_{n,s,b}$ in three dimensioning algorithms. Once the distribution is computed, we can use the general dimensioning algorithm (Algorithm 1) to dimension each network separately. **Independent stitched path dimensioning (IS)** is used for two networks that only share node information (minimum amount of shared information to enable end-to-end connection services). In this case, an internal node has no idea which border node(s) external calls will come from. Therefore, external traffic is splits uniformly (random) across all border nodes. In this case, $p_{n,s,b} = \frac{1}{|N'|}$. **Global shortest path dimensioning (GS)** provisions the network using least cost routes crossing two domains (assuming E2E). This is used when two networks are willing to share some form of cost information (in this paper we use path lengths) at the border node (at the appropriate E-NNI interface) to pick the global shortest path. The traffic rate to each border node is then weighted by the likelihood each border node is chosen by using E2E. Algorithm 2 shows the computation of $p_{n,s,b}$ for each inter-domain connection pair (n, s). **Normalized shortest path dimensioning (NS)** is similar to GS, except the metric used to compute the concatenated shortest path length. Algorithm 2 also shows the procedure to compute $p_{n,s,b}$ for NS assuming the use of nE2E.

V. FAIRNESS MEASUREMENT ON SHARED MECHANISMS

We propose *penalty ratio* to measure the fairness of resource usage on each network domain for shared routing and dimensioning. The penalty incurred by each network is computed by Equation 7. It basically sums the normalized number of extra hops used for each inter-domain routing. For example, if the TSLs of an internal node to three borders are 1, 4, 10 and the selected path uses the border that is 4 hops away, the penalty for that connection is 3. The total penalty for the dimensioned network is the weighted sum of penalties to all border nodes

---

[1] We omit crankback path selection because paths found by other shared routing algorithms present similar property to those found by fail and retry mechanisms in two-domain cases.

according to the traffic load distribution. Then, penalty ratio of two networks is the ratio of the penalty of each domain. This inherently assumes the use of a global fairness notion (for illustration purposes) without considering billing models. The metric can be weighted to conform to the needs of different peering agreements. Therefore, it is important to note that our presentation motivates the need for cost-performance studies given a fairness model that the domains have agreed upon.

VI. SIMULATION RESULTS AND DISCUSSION

Simulations were performed on the joint NSFNET-ARPANET (Figure 1) with uniform arrival rate, departure rate and requested capacity for Poisson traffic. The 95% confidence interval falls within $\pm 5\%$ of the results. The dimensioned capacity on each link (link # matches the number in Figure 1) using IS, GS and NS dimensioning algorithms are shown in Table II-III. As a result, the computed *penalty ratio* for NSFNET to ARPANET using GS is 2.46 and the ratio using NS is 0.99. If global fairness policy is enforced (where both networks should have comparable penalty), NS-nE2E is more fair compared to GS-E2E since its ratio is closer to one.

We first compare the performance of different schemes with varying loads. For each experiment, blocking probabilities of inter-domain traffic, NSFNET internal traffic and ARPANET internal traffic are separately shown to illustrate the impact of varying conditions on one or both domains. Blocking rates with varying inter-domain traffic loads are shown in Figure 2. Both domains gain significant reduction in blocking using shared information provisioning (GS and NS) compared to IS. At load 0.95 in Figure 2(b), over 80% reduction of internal blocking in NSFNET can be seen between IS and GS. NS-nE2E further improves blocking by 90% over GS-E2E. However, IS-E2E does not show much improvement over IS-CSR. This, in particular, shows the importance for domains to share information at the resource dimensioning stage. Joining two existing domains without re-provisioning may can result poor performance, even when a joint routing approach is used.

We next show how scaling of one network domain may affect the performance of other domains. Figure 3 shows the blocking performance when the size of NSFNET is scaled (under- or over-provisioning each link by the same ratio). Figure 3(c) shows that the internal blocking of ARPANET can increase with IS when NSFNET increases capacity. It is because an over-provisioned network can sustain more traffic (both internally and externally) and push additional inter-domain traffic to the other networks to compete with their own internal traffic. When one network continues to over-provision, with limited capacity provisioned for inter-domain links, the blocking probability of inter-domain traffic and internal traffic of the other domain will converge. This trend occurs as inter-domain links become saturated and limits the volume of the external traffic to the other domain. When one network is under-provisioned, the inter-domain traffic will suffer more blocking due to reduction in resources. ARPANET scaling can influence NSFNET in the same way, but the figures are omitted for brevity. Figure 4 shows the case where there is no capacity limit on inter-domain links (it is possible to have a significantly more capacity allocated for long reach cross-domain links or intra-office/site links). Similar trends are shown, but much lower blocking can be achieved for shared provision algorithms. Sufficiently large capacity available in inter-domain links prevent external traffic from being routing through unfavorable paths (much longer than TSL) via non-optimal border nodes.

TABLE II: NSFNET capacity (total 1491)

link #	IS	GS	NS
0	60	65	74
1	31	38	37
2	75	68	70
3	79	62	61
4	108	94	91
5	126	114	112
6	65	57	52
7	104	88	79
8	80	105	119
9	43	57	66
10	92	96	97
11	44	45	42
12	36	38	38
13	84	91	94
14	34	45	48
15	21	31	34
16	82	62	48
17	90	113	117
18	122	97	86
19	49	49	47
20	66	76	79

TABLE III: ARPANET capacity (total 2496)

link #	IS	GS	NS
0	26	26	26
1	71	67	64
2	144	116	118
3	93	93	96
4	48	48	51
5	123	114	115
6	89	78	81
7	55	53	51
8	29	40	37
9	116	95	92
10	118	80	75
11	127	81	90
12	145	126	125
13	135	166	165
14	49	59	60
15	10	19	17
16	41	52	51
17	42	56	51
18	67	111	106
19	50	82	75
20	100	91	101
21	68	99	96
22	205	153	172
23	63	80	80
24	38	48	48
25	65	91	84
26	34	45	46
27	36	39	37
28	31	60	55
29	167	119	124
30	65	62	61
31	46	47	46

In terms of resource allocation, NSFNET has a higher external/internal capacity ratio (2.82) compared to ARPANET (1.38) because NSFNET is a smaller network. It is important to note that, in general, a smaller network, such as NSFNET, has to pay more for in capacity to sustain the inter-domain traffic in proportion to its own size when it is joined with a larger domain. If equal cost-profit model is assumed, fair routing (as discussed earlier) can greatly improve internal blocking on small networks. As shown in Figure 2(b) and 3(b), nE2E improved the performance of NSFNET without affecting the performance of the other domain.

VII. CONCLUSION

We quantified the performance and fairness of various routing and dimensioning schemes in a multi-domain environment with a simple abstract levels of information sharing between the domains. We proposed the metrics to measure the loads of multi-domain traffic and the fairness of dimensioning and routing schemes. Our results showed that joint provisioning (in both dimensioning and routing) and playing fair is crucial to the performance of both networks for both internal and external traffic. We also showed that scaling of one network domain can affect the performance of the other network, especially if

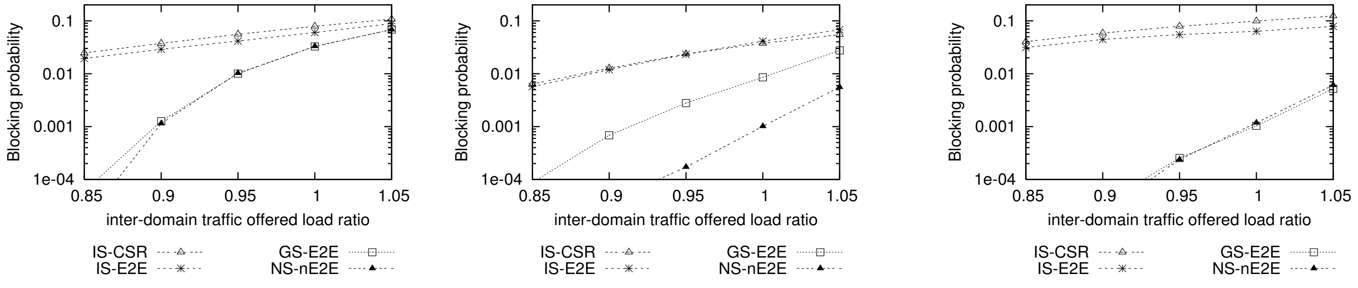

(a) Inter-domain traffic blocking rate (b) NSFNET internal traffic blocking rate (c) ARPANET internal traffic blocking rate

Fig. 2: Blocking probabilities in dimensioned networks with varying external offered loads. Fixed internal offered load 1.0.

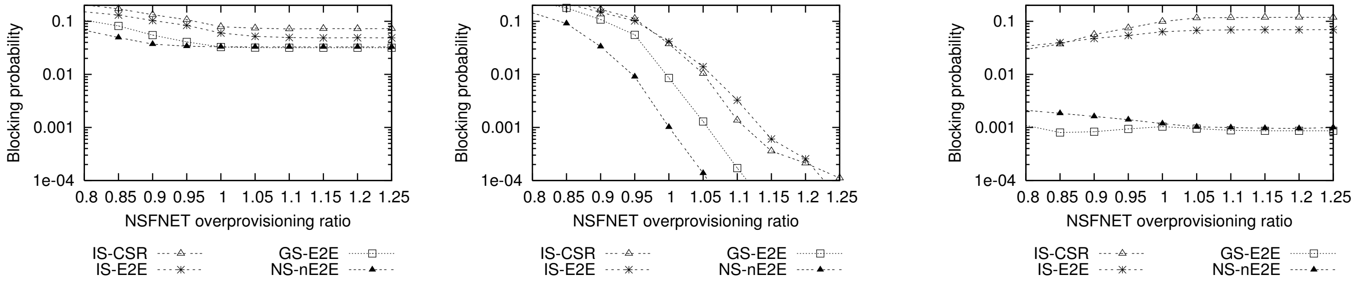

(a) Inter-domain traffic blocking rate (b) NSFNET internal traffic blocking rate (c) ARPANET internal traffic blocking rate

Fig. 3: Blocking probabilities in dimensioned networks with varying NSFNET sizes. Fixed all offered load 1.0.

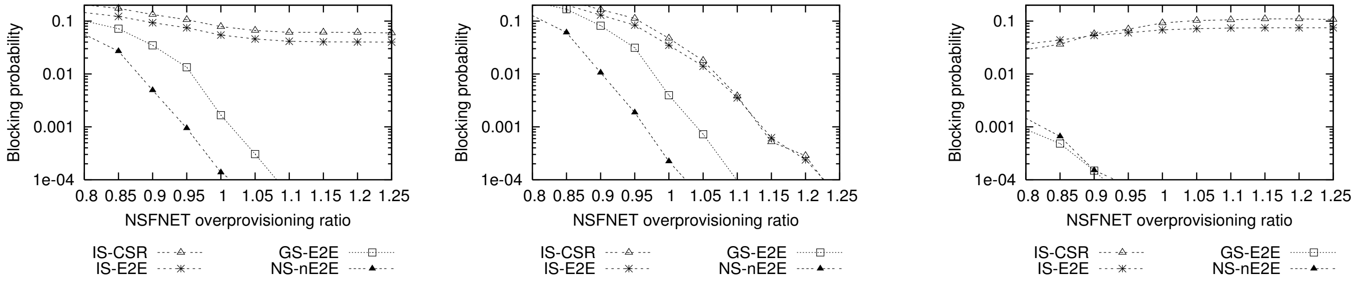

(a) Inter-domain traffic blocking rate (b) NSFNET internal traffic blocking rate (c) ARPANET internal traffic blocking rate

Fig. 4: Blocking probabilities in dimensioned networks with varying NSFNET sizes. Fixed all offered load 1.0. $C_b^I = \infty$

they are independently dimensioned. Our results motivate the need to gain a deeper understanding of the interaction between network domains of varying sizes, especially in maximizing the overall performance as well as fairness across domains that are independently owned. Being able to quantify the impact of decisions made in routing and dimensioning is critical in designing sound billing and control models needed for peering agreements. We must be able to reason about and design better solutions to support the ability to automatically and dynamically establish lightpaths.

REFERENCES

[1] X. J. Zhang, S. Kim, and S. S. Lumetta, “Resource dimensioning in WDM networks under state-based routing schemes,” in *IEEE/BroadNets’07*.

[2] Y. Zhu, A. Jukan, and M. Ammar, “Multi-segment wavelength routing in large-scale optical networks,” *IEEE/ICC ’03*.

[3] N. Ghani, Q. Liu, A. Gumaste, D. Benhaddou, N. Rao, and T. Lehman, “Multidomain optical networks: issues and challenges - control plane design in multidomain/multilayer optical networks,” *Communications Magazine, IEEE*, vol. 46, no. 6, pp. 78–87, June 2008.

[4] D. King, Y. Lee, H. Xu, and A. Farrel, “Path computation architectures overview in multi-domain optical networks based on itu-t ason and ietf pce,” *NOMS Workshops 2008*.

[5] R. Doverspike and P. Magill, *Optical Fiber Telecommunications*, I. P. Kaminow, T. Li, and A. E. Willner, Eds. London, UK: Elsevier Inc., 2008, vol. B: Systems and Networks.

[6] X. Yang and B. Ramamurthy, “Inter-domain dynamic routing in multi-layer optical transport networks,” *GLOBECOM ’03*.

[7] T. Saad, H. Mouftah, and A. Nouroozifar, “Constraint-based routing across multi-domain optical wdm networks,” *Electrical and Computer Engineering, 2004. Canadian Conference on*.

[8] S. Figuerola, J. Garcia, A. Sanchez, C. de Waal, and A. Willner, “The network service plane: An approach for inter-domain network reservations,” *ICTON 2008.* .

[9] J. Vasseur, R. Zhang, N. Bitar, and J. Roux, “A Backward Recursive PCE-based Computation (BRPC) procedure to compute shortest inter-domain Traffic Engineering Switched paths,” 2007.

[10] T. Tachibana and H. Harai, “Lightpath establishment without wavelength conversion based on aggressive rank accounting in multi-domain wdm networks,” *Lightwave Technology, Journal of*, vol. 26, no. 12, pp. 1577–1585, June15, 2008.

[11] J. Szigeti, I. Ballok, and T. Cinkler, “Efficiency of information update strategies for automatically switched multi-domain optical networks,” *Transparent Optical Networks, 2005*.

[12] G. Liu, C. Ji, and V. Chan, “On the scalability of network management information for inter-domain light-path assessment,” *Networking, IEEE/ACM Transactions on*, vol. 13, no. 1, pp. 160–172, Feb. 2005.

[13] G. Maier, C. Busca, and A. Pattavina, “Multi-domain routing techniques with topology aggregation in ason networks,” *ONDM 2008*.