Selective Randomized Load Balancing Outperforms VPN-Tree Routing for Disparate Demand and Link Granularities

Andrew B. Feldman[1,2], Jesse E. Simsarian[1] and Peter J. Winzer[1]
[1]Bell Labs, Alcatel-Lucent Holmdel, New Jersey 07733 Email: jesse.simsarian, peter.winzer@alcatel-lucent.com
[2]Massachusetts Institute of Technology, Cambridge, MA 02139 Email: abf149@mit.edu

***Abstract*—We find that selective randomized load balancing (SRLB) and Valiant's fully-randomized load balancing (RLB) can be lower-cost routing templates than VPN-tree routing for hose-constrained random network traffic, provided that link bandwidths have a significantly coarser granularity than individual demands, as is typically the case in today's optical networks. Our results build on the previous *VPN conjecture* and subsequent proof that VPN-tree routing provides the optimal capacity solution for the VPN design problem under the assumptions that the bandwidth demands and link capacities are both fractional or both integral with the same granularity. We consider ring and mesh network topologies and use both an analytic approach as well as a numerical exhaustive search to find the lowest-cost routing templates. In the analytic analyses, we find that using RLB on both full-mesh and ring networks can achieve lower cost than VPN-tree routed networks for fully hose-model traffic. In the exhaustive-search analyses, we calculate the cost of a general mesh network and study the effect of adding varying amounts of deterministic traffic to the hose traffic. We find that a 'background' of deterministic traffic tends to shift the optimal routing template away from VPN tree and towards SRLB and RLB solutions.**

I. Introduction

With the increasing availability of cloud-based services for storage and computing as well as the need to connect geographically-diverse sites with virtual private network (VPN) connectivity, Internet service providers become forced to design networks with a large degree of uncertainty in the bandwidth demands between network nodes. The hose model [1], [2] is a formal description of the case of maximum uncertainty in the bandwidth demands between nodes subject to bandwidth bounds from the nodes' physical interfaces. In contrast to a static demand matrix, where the demand between nodes i and j is given by the deterministic matrix $\{d_{ij}\}$, the hose model specifies the total ingress/egress capacity D_k of node k, while the demand between two nodes can take any random value subject to the hose constraint, i.e., $\sum_j d_{kj} = D_k^+$ for the ingress capacity and $\sum_i d_{ik} = D_k^-$ for the egress capacity. Frequently, again motivated by typically bandwidth-symmetric physical-layer node interfaces, the total ingress and egress capacities of each node are considered to be equal, i.e. $D_k^+ = D_k^- = D_k$, and the links connecting nodes are assumed bidirectional (undirected). We also make these assumptions in this work. While the hose model in its original form considers uniformly-random demand distributions, other traffic models have limited its pure randomness by imposing bandwidth caps [3] or a population-based correlative structure onto the random matrices [4] to introduce more realism.

Given a certain class of random demand matrices, the problem of *robust network design* [5] has aimed at finding routing templates that satisfy the hose demand under the constraint of oblivious routing, where packets follow a path from source to destination that does not depend on any network state information nor on the presence or absence of other packets in the network. Important examples of oblivious routing templates include the *shortest-path* (SP) routing template and the *hub-routing* template. For the case of unconstrained hose matrices and single-path routing, Gupta *et al.* [6] and Fingerhut *et al.* [2] found a solution to the hose-constrained robust network design problem using a shortest-path tree that is capacitated to carry all possible hose matrices (a so called *VPN tree*). It was then conjectured [7] and subsequently proven [8] that this routing template results in the minimum overall network capacity for general network graphs dimensioned to carry traffic under the hose constraint. Note also, as discussed in [9], that there is enough capacity on the VPN tree to use either *SP routing on the tree* or *single-hub routing* with the tree's root node as a routing hub. The two routing variants require the same link capacities but different (packet vs. circuit-switched) processing equipment at the nodes, and hence lend themselves to different types of network architectures with different overall cost trade-offs [9].

The *optimal* VPN tree in an N-node network is found by choosing the lowest cost of all N SP trees capacitated to route all hose matrices. One downside of routing on a tree, especially when packet routing only occurs at the root as opposed to at any node within the tree, is its susceptibility to single points of failure. This motivates the implementation of multiple trees and routing hubs, as in *selective randomized load balancing* (SRLB) [9]. SRLB is optimally accomplished by ordering the N possible SP trees in the network by a suitable cost metric, capacitating the K lowest-cost trees to accommodate $1/K$-th of the hose traffic, and during operation randomly distributing packets to these K trees for routing. In the limit of $K = N$, SRLB becomes Valiant's familiar *randomized load balancing* (RLB) scheme [10]–[14]. SRLB and RLB are robust routing templates that accommodate all possible hose matrices and generally have a lower cost than the SP routing template [9].

Previous work has studied the case where bandwidth demands and inter-node link capacities are fractional, as well as the case where both the bandwidth demands and the inter-node link capacities only come in multiples of the same

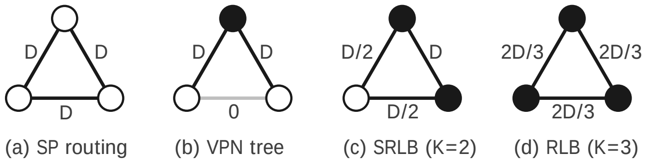

Fig. 1. Three-node network with links capacitated to carry unconstrained random hose matrices of uniform hose demands D using (a) shortest-path routing, (b) VPN-tree routing, (c) selective randomized load balancing across 2 hubs, and (d) Valiant's full randomized load balancing across the entire network.

discrete capacity units. There exist physical networks whose behavior can be approximated by each of these cases, and single-hub routing on the lowest-cost VPN tree is optimal for them under the hose constraint. Here we study the important situation frequently encountered in today's networks, where bandwidth demands and inter-node link capacities are quantized using *different* granularities. In particular, we consider infinitely fine granularities for the demand matrix elements and relatively coarse granularities for the link capacities. This model applies to, e.g., Internet protocol (IP) over wavelength-division multiplexed (WDM) optical networks and IP over an optical transport network (OTN). There, the granularities of the individual bandwidth demands, such as packet flows or lower-order OTN tributaries, are typically many orders of magnitude smaller than the unit of bandwidth quantization in the WDM network that connects the network nodes. For example, a flow in the network may be a few Mb/s, whereas the typical wavelength granularity in today's networks is 100 Gb/s [15].

In this paper, we re-visit the robust network design problem for single-path routing in the presence of fractional demands and quantized link capacities and find that in contrast to previous fractional-only and equal-granularity integral results, VPN tree can become *suboptimal* if the node-to-node demands exhibit a finer granularity than the link capacities. In fact, we show that SRLB and RLB can *outperform* the VPN tree under these conditions, both for networks designed to carry unconstrained hose traffic and for networks that carry a mixture of deterministic traffic and hose traffic, the latter being a more realistic model for carrier network traffic.

This paper is organized as follows: Section II gives two simple examples, illustrating our findings on a 3-node and 5-node network. Section III establishes an analytic framework to quantify the benefit of RLB and SRLB over routing on a single VPN tree for full mesh (Sec. III-A) and ring (Sec. III-B) networks. Section IV explores more complex planar mesh network topologies and includes a mix of deterministic and hose-model traffic. Finally, Section V summarizes our main findings and conclusions.

II. Two simple examples

In order to intuitively understand the behavior of fractional hose demands coupled with link capacity quantization in robust network design, consider the simple 3-node/3-link network shown in Fig. 1. Each node has an ingress/egress hose demand $D_k = D$ that is fractional. Using the SP routing template in Fig. 1a, each link needs to be dimensioned at capacity D, leading to an overall network capacity of $3D$ in the absence of link capacity quantization and of $3\lceil D \rceil$ in the presence of link quantization; $\lceil \cdot \rceil$ is the ceiling function that rounds link capacities up to the next integer (e.g., to the next number of wavelengths). In the case of VPN trees, owing to the symmetry of the problem, all three SP trees have equal cost, and one of the three VPN tree solutions is given in Fig. 1b. The filled circles in the diagrams denote routing hubs. Here, the required capacity to accommodate all possible hose matrices is $2D$ in the absence and $2\lceil D \rceil$ in the presence of link capacity quantization. As expected, VPN-tree routing outperforms SP routing for hose traffic. Figure 1c shows the SRLB solution for $K = 2$ hub nodes, where ingress traffic is randomly distributed on the two implemented trees. The links within each individual tree are dimensioned to carry $D/2$ capacity. Since the two trees share a common link, that link needs to carry a capacity of $2 \cdot D/2 = D$. The overall network capacity is hence $2D$ in the absence and $2\lceil D/2 \rceil + \lceil D \rceil$ in the presence of link capacity quantization. Similarly, full RLB ($K = N = 3$) distributes traffic among all 3 SP trees as shown in Fig. 1d, leading to a link capacity of $D/3$ on each link within each tree, and $2D/3$ on each link after the superposition of the 3 trees. The resulting network capacity is $2D$ in the absence and $3\lceil 2D/3 \rceil$ in the presence of link capacity quantization. Since all 3 trees are identical, we expect (b) through (d) to exhibit the same network capacities in the un-quantized case. In the quantized case, however, and taking a hose capacity of $D = 1.1$ as an example, we require a total network capacity of $2\lceil 1.1 \rceil = 4$ in (b), of $2\lceil 0.55 \rceil + \lceil 1.1 \rceil = 4$ in (c), and of $3\lceil 0.73 \rceil = 3$ in (d). *Hence we verify by example our conjecture that RLB can outperform VPN tree, at least for certain choices of D, on a network with link capacity quantization.*

Assuming further that the traffic matrix is composed of some deterministic portion D_D, with equal demands of $D_D/2$ between nodes, and some random hose portion D_R, and that the deterministic portion is routed SP (which is optimum for deterministic traffic), we superimpose the link capacities of Fig. 1a with those of (b) through (d) and arrive at $\lceil D_D/2 \rceil + 2\lceil D_D/2 + D_R \rceil$ for (a)+(b), $\lceil D_D/2 + D_R \rceil + 2\lceil D_D/2 + D_R/2 \rceil$ for (a)+(c), and $3\lceil D_D/2 + 2D_R/3 \rceil$ for (a)+(d). Assuming $D_D = D_R = 1$, the network requires a total link capacity of 5 for (a)+(b), of 4 for (a)+(c), and of 6 for (a)+(d). In this case, routing the random portion of the traffic using SRLB leads to the lowest overall network capacity.

Finally, we consider the simple example of a 5-node planar mesh network as shown in Fig. 2. Five SP trees on this topology are individually highlighted with darkened lines. Assuming random hose traffic with a nodal ingress/egress demand, D, the VPN-tree networks Root 1, Root 2, Root 4, and Root 5 on this topology each require a total capacity of $\lceil 2D \rceil + 3\lceil D \rceil$, and Root 3 requires a total capacity of $2\lceil 2D \rceil + 2\lceil D \rceil$. In the presence of link capacity quantization, the $K = 3$ SRLB solution based on roots 1, 2, and 5, as shown in Fig. 2, requires a capacity of $\lceil D \rceil + 6\lceil 2D/3 \rceil$. For $1 < D \leq 1.5$, each of the VPN-tree networks require a capacity of 9, while the $K = 3$ SRLB solution requires a capacity of only 8 and is hence lower cost than the VPN-tree solution.

A simple explanation for the lower cost of RLB or SRLB compared to VPN-tree routing in the above examples is that (S)RLB can spread traffic out more evenly in the network.

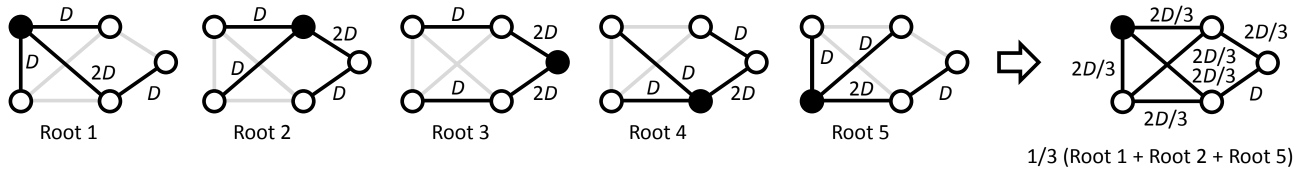

Fig. 2. Five-node network with five VPN trees (Roots 1 through 5) and a $K = 3$ SRLB solution, capacitated to carry unconstrained hose matrices with uniform hose demands D.

When link bandwidth is quantized, it may happen that for some choices of D, the quantization step is not reached in the (S)RLB case on some links where it is reached for VPN-tree routing. Therefore the (S)RLB template may allow for more efficient packing of the quantized bandwidth in the network.

III. Analytic results

Analytic expressions can be derived that quantify the benefit of RLB and SRLB over VPN-tree routing in the presence of link capacity quantization for full meshes as well as for ring network topologies. We assume random hose traffic with fractional and equal ingress/egress hose demands of $D_k = D$ at each node. In what follows, we will use the term 'cost' to denote the aggregate (bidirectional) link capacities required to route all hose matrices of marginals D.

A. Full mesh topologies

For a fully-connected N-node mesh network, we calculate the following costs of networks subject to different routing templates that are robust to any traffic variation satisfying the hose constraint:

SP routing: Owing to the full-mesh connectivity, each shortest path between any pair of nodes contains only a single link. As a result, a SP network design forces a required capacity of D onto each link, resulting in a total network cost of

$$\frac{N(N-1)}{2} D \tag{1}$$

for no link capacity quantization, and in the case of link capacity quantization a network cost of

$$\frac{N(N-1)}{2} \lceil D \rceil \ . \tag{2}$$

VPN tree: Due to the symmetry of the problem, all SP trees within the network capacitated to accommodate all hose matrices are identical VPN trees and require the hose capacity D on each direct link to the root. This results in a network cost of

$$(N-1)D \tag{3}$$

for no link capacity quantization, and in the case of link capacity quantization a network cost of

$$(N-1) \lceil D \rceil \ . \tag{4}$$

SRLB: Distributing the traffic randomly among K SP trees within the network requires a capacity of $2D/K$ for each link connecting two root nodes, and D/K for each link connecting a non-root node to a root. Since there are $K(K-1)/2$ links between K roots, and $K(N-K)$ links from non-root nodes to root nodes, we have a total network cost of

$$\frac{K(K-1)}{2} \frac{2D}{K} + K(N-K) \frac{D}{K} = (N-1)D \tag{5}$$

for no link capacity quantization, which due to the symmetry of the network (i.e., the equal cost for each root tree) also equals the cost of the optimum VPN-tree solution. In the case of link capacity quantization, however, we have a network cost of

$$\frac{K(K-1)}{2} \left\lceil \frac{2D}{K} \right\rceil + K(N-K) \left\lceil \frac{D}{K} \right\rceil \ . \tag{6}$$

RLB: Distributing all traffic randomly among all N SP trees, i.e., letting all nodes act as routing hubs ($K = N$), the above equations for K hubs simplify to

$$(N-1)D \tag{7}$$

for no link capacity quantization, and in the case of link capacity quantization to

$$\frac{N(N-1)}{2} \left\lceil \frac{2D}{N} \right\rceil \ . \tag{8}$$

We observe from the above equations that the cost of VPN-tree routing is always a factor of $N/2$ less than that for SP routing, independent of whether link capacities are quantized or not. Further, owing to the symmetry of the problem, all N possible SP trees are equal-cost VPN trees, which leads to any combination of $1 \leq K \leq N$ hubs in an (S)RLB architecture having equal cost in the un-quantized case. However, with link-capacity quantization present, this fact can change considerably. Figure 3a (symbols) shows the cost ratio of link-capacity quantized RLB to link-capacity quantized VPN-tree routing as a function of N, evaluated at the respective hose demands, $D_{\min}$ that minimize the cost ratio for each N. The figure clearly shows that in the presence of link capacity quantization, RLB with an odd number of nodes can outperform VPN-tree routing in terms of network cost. For networks with an even number of nodes, RLB has the same cost as VPN-tree routing. To understand this behavior, we use Eqns. (4) and (8) to write the network cost ratio $r_{\min}$ as

$$\begin{aligned} r_{\min} &= \min_D \left\{ \frac{N(N-1)/2 \ \lceil 2D/N \rceil}{(N-1) \lceil D \rceil} \right\} \\ &= \frac{N(N-1)/2 \ \lceil 2D_{\min}/N \rceil}{(N-1) \lceil D_{\min} \rceil} = \frac{x \lceil D_{\min}/x \rceil}{\lceil D_{\min} \rceil} \end{aligned} \tag{9}$$

with $x = N/2$ and $\min_D\{\cdot\}$ denoting minimization with respect to D, leading to $D_{\min}$. Examining the stepwise functions

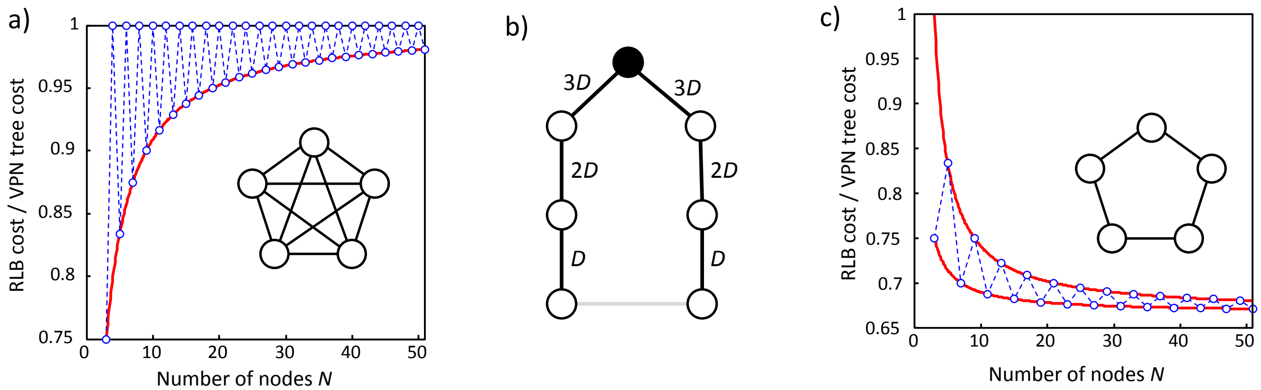

Fig. 3. (a) Analytic results for the cost ratio of RLB to VPN-tree routing for fully-connected meshes with N nodes and link-capacity quantization. (b) Diagram of a capacitated VPN tree for a 7-node ring. (c) Analytic results for the cost ratio of RLB to VPN-tree routing for ring networks with N nodes and link-capacity quantization. The inset diagrams of (a) and (c) show the mesh and ring network topologies for the case of 5 nodes.

represented by the numerator and denominator of the above expression as a function of D, we find that the ratio cannot be lower than 1 for any D when N is even, and that for odd N the lowest ratio as a function of D is $N/(N+1)$, such that

$$r_{\min} = \begin{cases} 1 & , N \text{ even} \\ \dfrac{N}{N+1} & , N \text{ odd} . \end{cases} \tag{10}$$

The solid line in Fig. 3a gives $r_{\min}$ for odd N.

Note that for full meshes and in the presence of link capacity quantization, we never found SRLB to be lower-cost than RLB at $D_{\min}$. However, SRLB *has* been found to be lower-cost than RLB in such networks at other values of D. It is also important to stress at this point that RLB and SRLB can both be *higher-cost* than VPN tree routed networks within certain ranges of D, but in the context of this study we are mostly interested in those values of D that yield the lowest-cost networks.

B. Ring topologies

We now describe the results of an analytical analysis for the case of ring networks. Restricting ourselves for simplicity to ring networks with an odd number of nodes (to avoid path selection ambiguities in minimum-hop routing), we find the following analytical solutions:

SP routing: The hose constraint forces a required capacity of $(N-1)D/2$ onto each of the links in the ring, resulting in a total network cost of

$$\frac{N(N-1)}{2}D \tag{11}$$

for no link capacity quantization, and in the case of link capacity quantization a network cost of

$$N\left\lceil \frac{N-1}{2}\,D \right\rceil . \tag{12}$$

VPN tree: As shown in Fig. 3b for the case $N = 7$, any of the N possible trees require a capacity of $(N-1)D/2$ on the two links directly emerging from the root, and D less on each subsequent link down to the two leaves of the tree, with a capacity of D on the final edges reaching the two leaf nodes. This results in a network cost of

$$\frac{D}{4}\left(N^2-1\right) \tag{13}$$

for no link quantization, and in the case of link quantization a network cost of

$$2\lceil D\rceil + 2\lceil 2D\rceil + 2\lceil 3D\rceil + \ldots + 2\left\lceil \frac{N-1}{2}\,D\right\rceil . \tag{14}$$

SRLB: Since all SP trees on the ring have equal cost owing to the symmetry of the problem, the cost for SRLB across K network nodes is the same for all choices of $K = 1$ through N in the un-quantized case, as was found in the full mesh topology studied above. A cost calculation for the link-capacity quantized case requires separately accounting for all possible placement combinations of the root nodes, leading to the results summarized in Tab. I for a 3-node, 5-node, and 7-node ring network.

RLB: If all nodes act as routing hubs, a closed form expression for the aggregate network capacity can be given both for the un-quantized case as

$$\frac{D}{4}\left(N^2-1\right) , \tag{15}$$

and for the link-capacity quantized case as

$$N\left\lceil \frac{D}{4N}\left(N^2-1\right)\right\rceil . \tag{16}$$

Comparing the above expressions, we find that in the un-quantized case, the cost of VPN-tree routing (as well as the cost of SRLB with arbitrary K) is a factor of $2N/(N+1)$ less than that for SP routing. In the quantized case, the cost ratios take a more complex form, especially for SRLB. Analytical expressions for the minimum cost ratio $r_{\min}$ between RLB

K	3 nodes	5 nodes	7 nodes
1	$2\lceil D\rceil$	$2\lceil 2D\rceil + 2\lceil D\rceil$	$2\lceil 3D\rceil + 2\lceil 2D\rceil + 2\lceil D\rceil$
2	$\lceil D\rceil + 2\lceil D/2\rceil$	$\lceil 2D\rceil + 2\lceil 3D/2\rceil + 2\lceil D/2\rceil$ $2\lceil 3D/2\rceil + 3\lceil D\rceil$	$\lceil 3D\rceil + 2\lceil 5D/2\rceil + 2\lceil 3D/2\rceil + 2\lceil D/2\rceil$ $2\lceil 5D/2\rceil + 2\lceil 2D\rceil + 3\lceil D\rceil$ $3\lceil 2D\rceil + 4\lceil 3D/2\rceil$
3	$3\lceil 2D/3\rceil$	$2\lceil 5D/3\rceil + 2\lceil 3D/3\rceil + \lceil 2D/3\rceil$ $3\lceil 4D/3\rceil + 2\lceil D\rceil$	$2\lceil 8D/3\rceil + 2\lceil 2D\rceil + 2\lceil D\rceil + \lceil 2D/3\rceil$ $2\lceil 7D/3\rceil + \lceil 2D\rceil + \lceil 5D/3\rceil + 2\lceil 4D/3\rceil + \lceil D\rceil$ $3\lceil 2D\rceil + 2\lceil 5D/3\rceil + 2\lceil 4D/3\rceil$ $2\lceil 2D\rceil + 4\lceil 5D/3\rceil + \lceil 4D/3\rceil$
4	n/a	$\lceil 6D/4\rceil + 2\lceil 5D/4\rceil + 2\lceil D\rceil$	$\lceil 10D/4\rceil + 2\lceil 9D/4\rceil + 2\lceil 6D/4\rceil + 2\lceil D\rceil$ $\lceil 9D/4\rceil + 2\lceil 2D\rceil + \lceil 7D/4\rceil + \lceil 6D/4\rceil + 2\lceil 5D/4\rceil$ $2\lceil 2D\rceil + 2\lceil 7D/4\rceil + 3\lceil 6D/4\rceil$ $\lceil 2D\rceil + 4\lceil 7D/4\rceil + 2\lceil 6D/4\rceil$
5	n/a	$5\lceil 6D/5\rceil$	$2\lceil 11D/5\rceil + 2\lceil 9D/5\rceil + 2\lceil 7D/5\rceil + \lceil 6D/5\rceil$ $3\lceil 2D\rceil + 2\lceil 8D/5\rceil + 2\lceil 7D/5\rceil$ $4\lceil 9D/5\rceil + 3\lceil 8D/5\rceil$
6	n/a	n/a	$\lceil 2D\rceil + 2\lceil 11D/6\rceil + 2\lceil 10D/6\rceil + 2\lceil 9D/6\rceil$
7	n/a	n/a	$7\lceil 12D/7\rceil$

TABLE I. EXPRESSIONS FOR THE COST OF 3, 5, AND 7-NODE RING NETWORKS FOR SRLB WITH DIFFERENT K IN THE CASE OF LINK BANDWIDTH QUANTIZATION. EACH SUBROW IN A TABLE ENTRY CORRESPONDS TO A DIFFERENT COMBINATION OF THE K ROOT NODES.

and VPN-tree routing can be obtained as follows:

$$r_{\min} = \min_D \left\{ \frac{N\lceil D(N^2-1)/4N\rceil}{2\lceil D\rceil + 2\lceil 2D\rceil + \ldots + 2\lceil (N-1)/2)D\rceil} \right\}$$

$$= \frac{N\lceil D_{\min}(N^2-1)/4N\rceil}{2\lceil D_{\min}\rceil + 2\lceil 2D_{\min}\rceil + \ldots + 2\lceil ((N-1)/2)D_{\min}\rceil}$$

$$= \frac{N}{2}\left\{ \frac{\lceil x_1 D_{\min}\rceil}{\lceil D_{\min}\rceil + \lceil 2D_{\min}\rceil + \ldots + \lceil x_2 D_{\min}\rceil} \right\}, \tag{17}$$

with $x_1 = (N^2-1)/(4N)$ and $x_2 = (N+1)/2$. Examining the resulting stepwise functions represented by the numerator and denominator of the cost ratio as a function of D, we find that the minimum value that the ratio in brackets can attain is $2/(3x_2+1)$ for odd x_2 and $2/(3x_2)$ for even x_2, such that

$$r_{\min} = \begin{cases} \dfrac{2N}{3N-3} & ,(N-1)/2 \text{ even} \\ \dfrac{2N}{3N-1} & ,(N-1)/2 \text{ odd}\,. \end{cases} \tag{18}$$

Figure 3c shows a numeric minimization of the cost ratio to arrive at $r_{\min}$ (symbols) as well the two analytic expressions given by (18) (solid lines). Interestingly, in a ring network with link capacity quantization, the cost ratio of RLB versus VPN-tree routing approaches $2/3$ for large N, *making the benefit of RLB a sustained feature on ring topologies*, although $D_{\min}$ tends to decrease as N increases. We stress again that Fig. 3c shows the relative cost evaluated at $D_{\min}$. RLB and SRLB on ring topologies can cost more than than VPN-tree routing for some other choices of the hose demand D.

C. Lower bounds for general mesh topologies

To arrive at a lower bound for the cost of general mesh network topologies, we begin with the provably lowest-cost routing template in the un-quantized (fractional) capacity case as well as in the uniform-granularity (integral) capacity case given by the VPN tree [2-6]. Denoting the capacities of the L links making up the lowest-cost VPN tree in a general mesh network by c_i ($i = 1, \ldots L$), this lower bound becomes

$$\sum_{i=1}^{L} c_i\,. \tag{19}$$

A slightly tighter bound is given by taking into account link-capacity quantization as

$$\left\lceil \sum_{i=1}^{L} c_i \right\rceil\,. \tag{20}$$

Note that while tighter than the former, even the latter bound may not be achievable on all network topologies.

For the example mesh network of Fig. 2, these two lower bounds would be at $2D+3D = 5D$ and at $\lceil 5D\rceil$, respectively. Varying D from 1 to 1.5 in that example lets this lower bound grow from 5 to 8. Hence, the discussed example SRLB network with $K = 3$ actually achieves its lower bound in the hose demand region $1.4 < D \leq 1.5$.

IV. NUMERICAL ANALYSIS

A. Method

To solve for general network topologies, and to explore various combinations of deterministic and hose traffic, we use numerical analyses with exhaustive search to compare the cost of different routing templates. The exhaustive search calculates the network cost for every possible combination of root nodes with different values of deterministic demands, D_D, and hose demands, D_R. For a network topology with N nodes, the number of root configurations explored by the exhaustive search algorithm is $2^N - 1$. The exponential growth in computational complexity constrains us to test networks with relatively small N.

Figure 4a shows the model six-node mesh network that we analyzed, with the links labeled by distance in km. Our previous analytical results assumed equal distance between nodes so that routing decisions were based solely on the

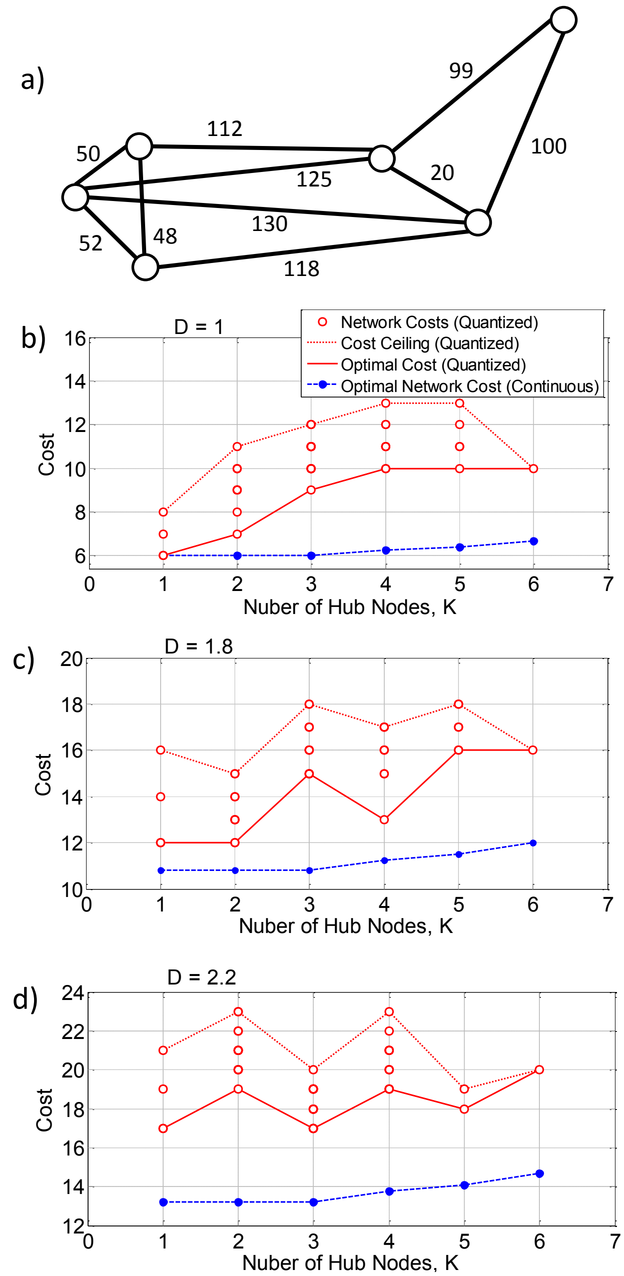

Fig. 4. Results of the numerical cost analysis on (a) the model 6-node mesh network with distances in km and pure hose traffic. Network cost as a function of SRLB hub count K for (b) $D = 1$ (c) $D = 1.8$, and (d) $D = 2.2$. Open circles are the network costs for all possible combinations of K routing hubs with link capacity quantization. $K = 1$ and $K = 6$ represent VPN tree and RLB, respectively. The dotted red line traces the highest-cost networks and the solid red line traces the lowest-cost networks. The filled circles are the network costs of Eq. (19) for continuous link capacities.

minimum number of hops. In the numerical analysis reported in this section, we include link lengths so that shortest-path routing becomes shortest-distance routing. For the random (hose) portion of the traffic, D_R, the exhaustive search routine begins by using Dijkstra's algorithm to efficiently construct N SP trees rooted at every node. For the case of single-tree routing, $K = 1$, the link capacities are calculated according to the VPN tree capacity allocation algorithm [1] [2] for each of the N trees.

For SRLB or RLB, i.e., $2 \leq K \leq N$, the algorithm creates routing templates for every combination of SP trees rooted at all possible K-node combinations. For a given routing template, the link capacities are calculated by distributing the node traffic uniformly across the K SP trees, each tree carrying $1/K$-th of the hose traffic.

As above, we take the network 'cost' to be the aggregate link capacity in the network, $\sum_i c_i$. The task of finding the optimal routing then reduces to finding the set of K hub nodes that minimizes network cost. For the case of mixed traffic, we assume for each node an equal amount of random hose traffic D_R and deterministic node traffic D_D, with uniform deterministic node-to-node demands of $D_D/(N-1)$. In order to capacitate the network for mixed traffic, we first calculate the (un-quantized) link capacities $c_{i,R}$ required for all possible combinations of cost of VPN trees for the random hose traffic of marginal D_R. Next, we calculate the required (un-quantized) link capacities $c_{i,D}$ for the deterministic portion of the traffic, assuming SP routing on the original network graph, which for deterministic traffic represents the optimum routing template. The un-quantized overall link capacities are then given by $c_i = c_{i,R} + c_{i,D}$. While the optimum choice of VPN tree(s) for the random portion of the hose traffic is independent of the deterministic 'background' traffic in the case of un-quantized link capacities, this situation changes for link capacity quantization. Here, the overall link capacities are given by $c_i = \lceil c_{i,R} + c_{i,D} \rceil$, and the *presence of deterministic traffic does affect the optimal choice of VPN tree(s)* as well as the overall network capacity requirements and hence the overall network cost. In fact, the presence of a deterministic demand 'background' inside the ceiling function leads to more cases where SRLB is lower-cost than simple VPN-tree routing, as discussed in Subsec. IV-C below.

B. Results for pure hose traffic on a mesh network

Figures 4b – d show the results of applying the above numerical procedure to the 6-node network of Fig. 4a with different values of pure hose traffic, $D = D_R$, which is assumed to be the same at each node. No deterministic traffic 'background' is assumed here. The x-axis of the graphs is the number of SRLB routing trees, K, and the multiple data points for each K are the network costs for each of the VPN tree combinations. The open symbols are the costs including the effect of link quantization and the closed data points are the network costs for continuous link capacities. The cost ceiling and optimal cost bounds for quantized link capacity are also shown in dotted and solid lines, respectively. In these results, the single VPN tree cost is always lowest cost (or at least tied for lowest). In contrast to the un-quantized capacity case, Figs. 4(c) and (d) show that the lowest network cost under link capacity quantization as a function of K is non-monotonic. The results of Sec. III-A prove that the network cost in the case of continuous link capacities is independent of K for a fully-connected mesh. However, for a non-fully connected mesh with continuous link capacities, as considered here and in previous work [9], we observe a monotonically-increasing relationship between cost and hub-count, which implies that

the added redundancy of SRLB comes at an additional cost. The non-monotonic relationships seen in Fig. 4 suggest that extra hub nodes can be added for minimal or even for no extra cost for some D.

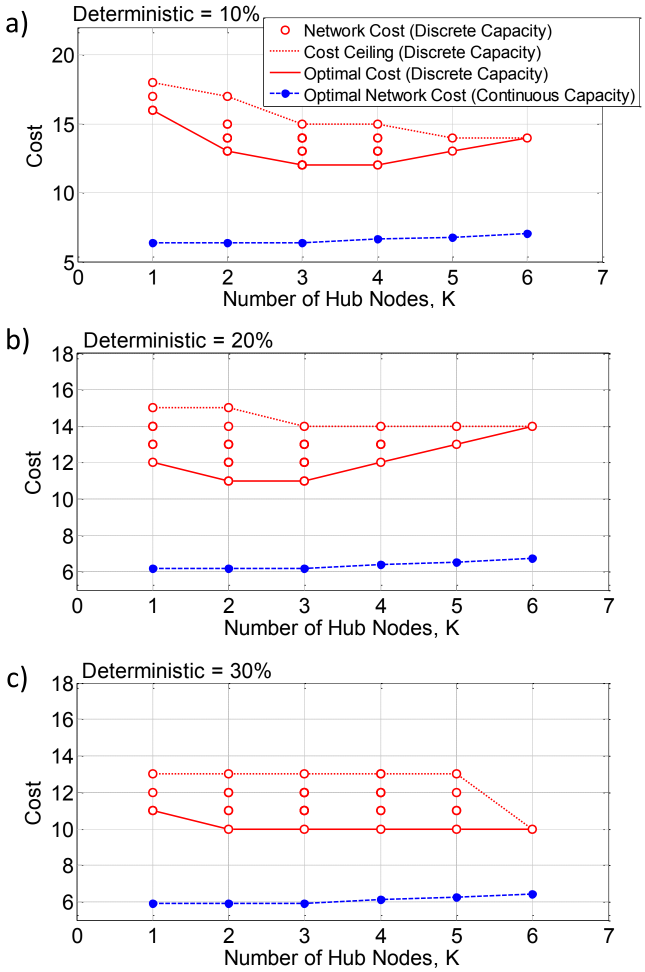

Fig. 5. Results of the numerical analysis for the 6-node mesh network of Fig. 4a with fractions of deterministic traffic equal to (a) 10%, (b) 20%, and (c) 30% and total hose traffic of $D = 1.1$.

C. Results for mixed traffic a mesh network

Fig. 5 shows the results of our numerical cost analysis with $D = D_D + D_R = 1.1$ but different fractions of random and deterministic traffic. For the deterministic traffic, we assume uniform node-to-node demands of $D_D/(N-1)$. Fig. 5a shows the case with 10% deterministic traffic and 90% hose traffic. We observe that the added deterministic traffic makes the single VPN tree the most expensive option, and the minimum network cost occurs for SRLB with $K = 3$ and 4 nodes. The achievable lower cost of SRLB compared to VPN-tree routing persists for the 20% and 30% deterministic traffic cases in Figs. 5(b) and (c). We can explain this behavior by noting that a small amount of static traffic on its own forces the network links to carry entire wavelengths that are only partially filled. The added hose traffic then makes the final network cost dependent on how efficiently the hose traffic can be packed into

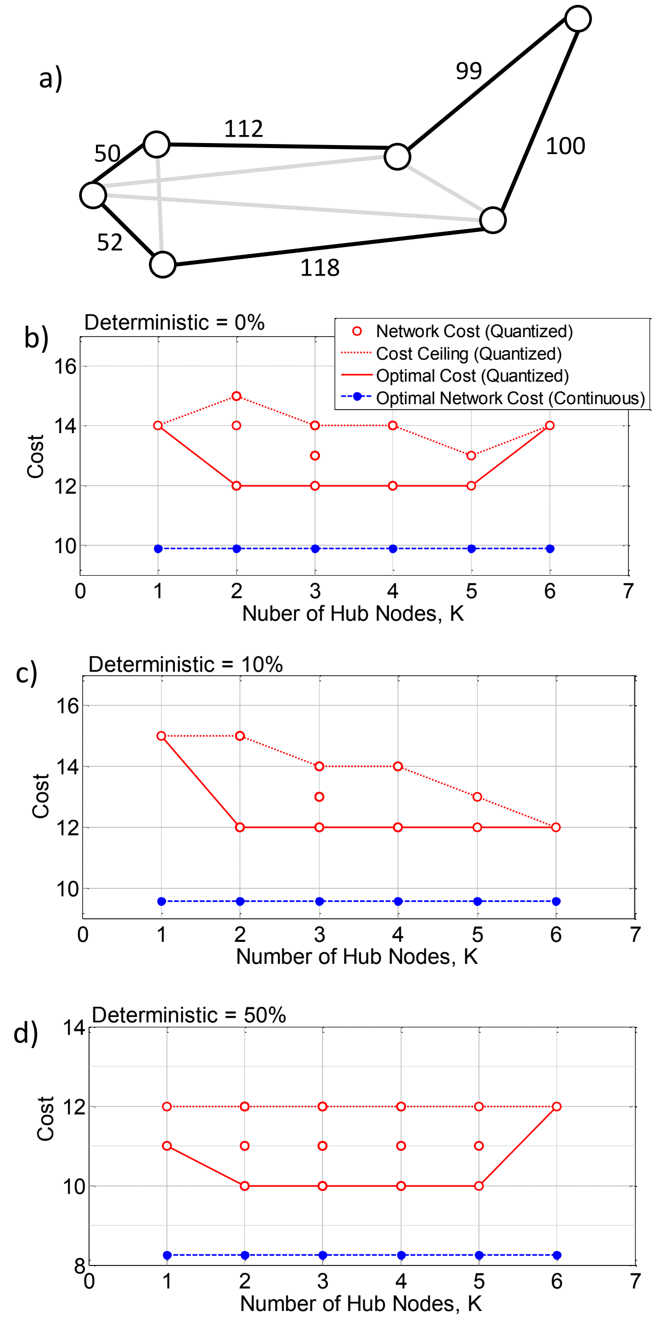

Fig. 6. (a) Diagram of the 6-node ring with the ring links drawn in dark lines, and the result of the numerical cost analysis for the 6-node ring with $D = 1.1$ and deterministic traffic equal to (b) 0%, (c) 10%, and (d) 50%.

the already-provisioned wavelengths. Comparing VPN tree to SRLB with $K > 1$, we note that SRLB spreads traffic over a greater number of links (because it is the superposition of multiple VPN trees) and therefore achieves more efficient packing onto the wavelengths provisioned for the deterministic traffic 'background'. It is also worth noting that the the network cost tends to decrease for both quantized and un-quantized link capacities as the fraction of deterministic traffic increases. This reflects the fact that deterministic traffic requires less network resources to route than random hose traffic and that there is

always a *'robustness premium'* to be paid when designing a network for random traffic matrices [9].

D. Results for hose and mixed traffic on a ring network

We also performed the numerical cost analysis on the ring network shown in Fig. 6a, implemented on the mesh topology of Fig. 4a. The addition of link distance defines shortest-path routes without ambiguities on a ring with an even number of nodes. The cost of the various routing strategies on this ring network versus the number of trees, K, is shown in Fig. 6b for full hose traffic. Here we observe that for $D = 1.1$, the cost for $K = 2$ to 5 SRLB nodes can be lower than the VPN tree. The lower cost of SRLB persists for 10% and 50% deterministic traffic.

V. Conclusion

Through simple examples, analytical, and numerical analyses, we have shown that *VPN-tree routing can be sub-optimal* in networks with fractional demands and link capacity quantization that is characteristic of today's optical networks. Under these conditions, selective randomized load balancing (SRLB) or Valiant's fully-randomized load balancing (RLB) can outperform VPN-tree routing. The effect becomes even more pronounced on networks with a mix of random hose and deterministic point-to-point traffic, which is a more representative model for traffic in many carrier networks.

Acknowledgment

The authors acknowledge valuable discussions with M. Kodialam, T. V. Lakshman, and T. Klein of Bell Labs.

References

[1] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. E. van der Merive, "A flexible model for resource management in virtual private networks," In *Proceedings of ACM SIGCOMM*, pp. 95-108, 1999.

[2] A. J. Fingerhut, S. Suri, and J. S. Turner, "Designing least-cost nonblocking broadband networks," *Journal of Algorithms*, vol. 24, no. 2, pp. 287-309, August 1997.

[3] A. Frechette, F. B. Shepherd, M. K Thottan, P. J. Winzer, "Shortest path versus multi-hub routing in networks with uncertain demand," In *Proceedings IEEE INFOCOM*, pp. 710-718, 2013.

[4] Franz Fidler, Peter J. Winzer, Marina K. Thottan, and Keren Bergman, "Impairment-Aware Optical Networking Using Cross-Layer Communication," *J. Opt. Commun. Netw.*, vol. 5, no. 2, Feb. 2013.

[5] C. Chekuri, "Routing and network design with robustness to changing or uncertain traffic demands," *SIGACT News*, vol. 38, no. 3, pp. 106-128, Sept. 2007.

[6] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener, "Provisioning a virtual private network: a network design problem for multicommodity flow," In *Proceedings of ACM STOC*, pp. 389-398, 2001.

[7] G. Italiano, S. Leonardi, and G. Oriolo, "Design of trees in the hose model: the balanced case," *Operations Research Letters*, vol. 34, no. 6, pp. 601-606, 2006. A preliminary version appeared in *Proceedings of ARACNE*, 65-76, 2002.

[8] N. Goyal, N. Olver, and F. P. Shepherd, "The VPN conjecture is true," In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pp. 443-450, 2008.

[9] F. B. Shepherd and P. J. Winzer, "Selective randomized load balancing and mesh networks with changing demands," *Journal of Optical Networking*, vol. 5, no. 5, pp. 320-339, 2006.

[10] L. G. Valiant, "A Scheme for Fast Parallel Communication," *SIAM J. Comput.*, vol. 11, no. 2, May 1982.

[11] C. S. Chang, D. S. Lee, Y. S. Jou, "Load balanced Birkoff Von Neuman switches, part I: one stage buffering," *Computer Communications*, vol. 25, pp. 611-622, 2002.

[12] Rui Zhang-Shen and N. McKeown, "Designing a Fault-Tolerant Network Using Valiant Load-Balancing," in *Proceedings of INFOCOM'08*, pp. 2360 - 2368, april 2008.

[13] Albert Greenberg *et al.*, "VL2: A Scalable and Flexible Data Center Network," in *ACM SIGCOMM Computer Communication Review - SIGCOMM '09* vol. 39, no. 4, pp. 51-62, Oct. 2009.

[14] Murali Kodialam, T. V. Lakshman, and Sudipta Sengupta, "Traffic-oblivious routing in the hose model," *IEEE/ACM Transactions on Networking (TON) archive*, vol. 19, no. 3, pp. 774-787, June 2011.

[15] P. J. Winzer, "Beyond 100G Ethernet," *IEEE Communications Magazine*, vol. 48,no. 7, pp. 26-30, July 2010.