

Towards a Cost Model for Network Traffic

Murtaza Motiwala [†], Amogh Dhamdhere ^{*}, Nick Feamster [†], Anukool Lakhina [§]

[†] Georgia Tech ^{*} CAIDA [§] Guavus

murtaza@gatech.edu, amogh@caida.org, feamster@gatech.edu, anukool.lakhina@guavus.com

ABSTRACT

We develop a holistic cost model that operators can use to help evaluate the costs of various routing and peering decisions. Using real traffic data from a large carrier network, we show how network operators can use this cost model to significantly reduce the cost of carrying traffic in their networks. We find that adjusting the routing for a small fraction of total flows (and total traffic volume) significantly reduces cost in many cases. We also show how operators can use the cost model both to evaluate potential peering arrangements and for other network operations problems.

Categories and Subject Descriptors: C.2.3 [Network Operations]: Network management

General Terms: Algorithms, Measurement, Management

Keywords: traffic cost model, cost optimization

1. INTRODUCTION

Carrying traffic in an IP network incurs many costs, including transit fees, port costs, backhaul costs, and other operational and capital costs. How traffic is routed across a network and exchanged with neighboring ISPs can significantly affect the overall costs of forwarding traffic over the network. For example, the costs of carrying traffic over trans-oceanic or satellite links is more expensive than routing traffic over underutilized commodity backhaul links; similarly, routing traffic over transit links incurs more cost than routing over settlement-free peering or customer links. Although traffic costs may not be the dominant cost in running a network, they can play a significant role in helping operators make decisions about planning, provisioning, and traffic engineering.

Operators understand how individual elements contribute to operational costs, but they lack a holistic cost model that maps traffic flows to the costs of carrying the traffic. As a result, although business-level decisions about peering, provisioning, and interconnection may consider costs of individual elements (*e.g.*, the cost of peering or interconnection), these decisions are currently ad hoc. For example, a decision about whether an operator should peer at a particular location should not only take into account the cost of that individual peering session, but also potential costs *saved* by sending less traffic over backhaul links. The inability to attribute costs to traffic flows can result in missed opportunities for cost savings and ad hoc decisions about routing and interconnection. Previous work jointly optimized cost and performance in a multihomed stub network [7], but no similar approach exists for transit networks or networks that peer in multiple locations.

Making decisions about traffic based on cost is challenging for two reasons. First, information about traffic costs is relatively inaccessible; if this information is available, it typically comes as individual cost elements, rather than as a holistic model. Further, some aspects of traffic costs are not linear (*e.g.*, commit rates, traffic symmetry constraints, 95th percentile pricing), and these costs do not map naturally to *individual* flows. We solve this problem by developing a holistic cost model that associates a cost to each traffic flow that incorporates both interconnection and backhaul costs, as

well as non-linear cost elements (like percentile pricing) with approximate functions. Using this model, operators can input values for various aspects of cost that they are likely to know from other sources; the model outputs an overall cost for routing each traffic flow. Second, the number of traffic flows and the number of possibilities for routing each flow makes it difficult to efficiently find a solution that reduces cost. To solve this problem, we use our cost model to identify the most expensive traffic flows in the network and apply heuristics to move those flows to less expensive links. We also demonstrate how associating costs to traffic flows can help operators in reasoning about decisions, such as where and whom to peer with.

Our evaluation shows that network operators can realize significant cost savings by moving only a small fraction of overall traffic flows: For example, we find that, for three realistic cost scenarios, moving 10% to 30% of the flows that reduce traffic cost in the network can help achieve at least 65% of total possible cost savings.

Many network planning tools and techniques can build on the cost model that we present in this paper. We expect that our model might ultimately be coupled with tools that help network operators make the actual configuration changes to reassign these flows. It could also be incorporated with tools that help network operators perform forecasting, to better help make better decisions regarding network upgrades and provisioning.

2. NETWORK TRAFFIC COST: A MODEL

We develop a model for reasoning about the costs incurred by a network for carrying IP traffic. To build this model, we first need to understand the components that contribute to the cost of carrying traffic in a network. Figure 1 shows a breakdown of cost into two components: *interconnect costs* and *backhaul costs*. We further discuss what contributes to these cost components.

2.1 Interconnect Costs

We refer to the cost associated with the place where traffic is exchanged with neighboring networks—including providers, peers, or customers—as *interconnect cost*. Depending on the agreement between the two networks, a network might pay for transit based on the volume of traffic exchanged, be paid by the other network, or engage in settlement-free peering. Transit fees vary depending on the geographic location of the interconnect point (*e.g.*, transit fees vary from \$2–5 per Mbps in the United States to about \$40 per Mbps in Asia [3,9]). For transit providers, which charge customers, the interconnect cost would be negative.

A network must also pay recurring *port costs*. These port costs include costs associated with buying network interface cards and paying for installation fees for buying ports. The port and installation costs depend on the exchange and the transmission medium. At a public exchange, the member network can exchange traffic with other networks present at the exchange using a single port; adding a peer at the same exchange has no incremental cost, as long as the aggregate traffic from all the peers does not exceed the port capacity. If the traffic exchanged exceeds the port capacity then




Figure 1: Classification of traffic costs for a flow f . A flow aggregates traffic from multiple individual connections.

the network can pay for additional ports at the exchange. A private peering (sometimes called a “private interconnect”) between two networks requires purchasing a separate port (and interface cards) for every neighbor network. Although private peering is more expensive than public peering, the traffic over the interconnect may be more predictable; this option may be cost-effective if two peers exchange a lot of traffic. In addition, there exist *other fixed costs* at an exchange, such as paying an annual (or monthly) fee for being a member of the exchange and a one-time installation fee.

To simplify reasoning about interconnect costs, we consider all of these cost components as *the cost per neighbor network at a given geographic location (PoP or exchange)*. For interconnect costs that depend on traffic volume, we make the above cost proportional to the rate of traffic, whereas for fixed costs, the operator can choose to amortize the cost over a period of, say, a few years. Although we may not know the current or accurate values for each of the contributors to interconnect costs, we expect that network operators can fill in the values for the different cost components.

2.2 Backhaul Costs

A network must carry traffic across its own network either to a customer network or another neighboring network; a network incurs costs from three components: the circuits themselves, the equipment for the backhaul links (e.g., routers, switches), and operational costs associated with running the network. We describe these costs below.

In practice, *circuit costs* may fall into two distinct categories: (1) *Metro-range* costs, which reflect the cost of connectivity in a metropolitan region; and (2) *Regional* costs, which depend on the geographic location and the distance of the regional circuit. Metro-range costs are often negligible for large networks but may be significant if a network leases local network connectivity from other networks. Although carrying traffic across backhaul links roughly correlates with the distance of the circuits, some regions are more expensive than others. Our model can incorporate these differences with a distance function that depends on the location of the PoPs.

Capital costs, such as the purchase, maintenance, and upgrade of routers and switches, also contribute to the backhaul costs. Depending on the network, the model might either reflect these costs as fixed or amortize them over several years. *Operational costs*, such as salaries for network operators to paying for cooling and power consumption can also be amortized. We incorporate these costs as fixed costs and include them as a component of the cost of all backhaul links in the network.

2.3 Cost Model

We now describe the cost model that we develop, based on the enumeration of costs from the previous section. The model we develop does *not* tell a network operator how to adjust the routing configuration itself to actually move a particular traffic flow from

one path to another but can help identify which traffic flows should be moved and could provide input to “what if” configuration analysis tools (e.g., [6]), or even a network designed around central routing control (e.g., [5, 8]).

We now describe a formal traffic cost model. We can write the total cost of running a network as a sum of the fixed network costs and the usage-based costs.

Fixed Costs The fixed costs (C_F) are defined by the network’s topology (backbone) and its relationships with the neighboring ASes (interconnect topology). Although the backbone and interconnect topology depend on the traffic the network is designed to carry, in the short term we assume the fixed costs are independent of traffic.

We abstract the backbone costs as the cost for the path between every pair of PoPs that exchange traffic. The cost component $c_{F,b}(p_1, p_2)$ is the fixed-cost component of the backbone path between PoPs p_1 and p_2 . An alternate formulation could replace the fixed backbone cost with the cost of each backbone link in the network. The fixed interconnect cost component $c_{F,i}(a, p)$ is the fixed cost for the interconnect between neighbor AS a at PoP p .

$$C_F = \sum_{p_1, p_2} c_{F,b}(p_1, p_2) + \sum_{a, p} c_{F,i}(a, p)$$

Usage-based Costs The usage-based component of the cost (C_U) depends on the volume or rate of flow f , and the route that f takes in the network. The usage-based component has three sub-components, as shown in Figure 1. A flow f enters a network at an interconnect; the cost associated with that flow is the cost at the ingress interconnect ($c_{u,i}(f)$). The flow traverses the backbone with cost $c_{u,b}(f)$ and leaves the network at an interconnect, which has an egress interconnect cost of $c_{u,e}(f)$.

$$C_U = \sum_f (c_{u,i}(f) + c_{u,b}(f) + c_{u,e}(f))$$

We will now describe how to calculate each of these cost components. For the usage-based cost components, each function and term refers to a flow f ; thus, we drop f from the notation, and the usage-based cost of a flow is simply:

$$c_u = c_{u,i} + c_{u,b} + c_{u,e} \quad (1)$$

Usage-based Interconnect Cost The equations are symmetric for the ingress and egress points, and hence both interconnect costs have the same form. For a particular interconnect, the usage-based interconnect cost is: $u_i \cdot R + s_i \cdot R^\alpha$, where R is the volume (or rate) of the total interconnect traffic that is charged, u_i is the charge per volume (or rate). We use a concave function of the form $s_i \cdot R^\alpha$ to approximate certain types of costs like port costs, which are a step function of the traffic rate. Previous work focusing on peering contracts has estimated the value of α to be between 0.4 to 0.75 [2] using market price data. The unit usage-based cost parameter, u_i , depends on the neighbor a and PoP p of the interconnect. We assume that s_i depends only on the PoP p , which yields $u_i = U_i(a, p)$ and $s_i = S_i(p)$, where a is either the ingress or egress AS and p is either ingress or egress PoP. $U_i(a, p)$ is the price per unit of exchanging traffic with AS a at PoP p . $S_i(p)$ is also in units of price per unit of traffic volume (or rate) and depends on the PoP (or exchange) p the network is present. This price reflects the port costs, which depend on the PoP. The total interconnect cost is thus: $U_i(a, p) \cdot R + S_i(p) \cdot R^\alpha$.




Figure 2: Example: Routing Decisions.

Next, we must find the contribution of flow f to the total interconnect cost. Most transit pricing on the Internet is based on the 95th percentile of traffic, where the transit provider charges for the traffic based on the 95th percentile of traffic volumes that are sampled over five-minute intervals. Because a network pays for the 95th percentile of the aggregate traffic at the interconnect, the model must reflect the contribution of any particular flow to that price. We use two techniques to approximate incorporating 95th percentile pricing of interconnect links. Let r be the volume (or rate) of the flow f .

- **Linear function** We assume that the 95th percentile is a linear function of the average or the peak traffic rate at the interconnect, as has been empirically observed for different types of networks [4]. In this case, we calculate the per flow contribution by replacing R with some constant times the volume (or rate) of the flow. Formally, $c_{u,i} = U_i(a, p) \cdot r + S_i(p) \cdot r^\alpha$.
- **Shapley value** The linear function ignores the distribution of the flow across different time intervals, which can influence the 95th percentile price at the interconnect. Stanojevic *et al.* [11] propose the use of Shapley value [10] for computing the contribution of each flow to the 95th percentile price of interconnect links. Because computing Shapley value is computationally infeasible for even a small number of flows, we instead propose way to approximate the Shapley value of each flow.

Usage-based Backhaul Cost The backhaul cost $c_{u,b}$ is: $c_{u,b} = r \cdot u_b$ and has two factors: the geographic location of the ingress PoP, p_i , and egress PoP, p_e , and the distance between them. The unit usage-based cost u_b depends on properties of the ingress and egress PoPs, and the distance between the ingress and egress. We represent this cost as $u_b = U_b(p_i, p_e) = R(p_i, p_e) \cdot D(p_i, p_e)$, where $U_b(p_i, p_e)$ is the unit cost per traffic rate unit, $D(p_i, p_e)$ is the distance between p_i and p_e , and $R(p_i, p_e)$ accounts for the dependence of the usage-based backhaul cost on the ingress and egress PoP. For a flow f with rate r , the usage-based backhaul cost is: $c_{u,b} = r \cdot R(p_i, p_e) \cdot D(p_i, p_e)$.

3. APPLICATIONS OF THE COST MODEL

In this section, we present examples to motivate and demonstrate the utility of understanding the traffic cost of flows in a network. These examples are empirical and we have developed them after talking with network managers at different ISPs around the world. Attribution of costs to traffic flows can help in determining how to route traffic in a network to reduce cost, and also in making long term planning decisions like which networks to select for peering. We classify these applications into two categories, based on whether the decision can be implemented with changes to existing routing configurations, or whether the changes require more fundamental modifications to existing peering relations. We present these from the perspective of a network that we refer to as “network X”.

3.1 Routing Decisions - Cost Optimization

In the example in Figure 2(a), network X can forward a flow arriving at PoP S via either PoP A or B . This situation could arise if network X peers with a particular neighbor at two locations, A and B for instance, and can choose to route traffic via

either PoP. Further, as shown in Figure 2(b), network X may be able to route traffic to a particular destination via multiple neighbor ASes. The operator of network X thus has choices for the egress AS and egress PoP over which to route a given flow, from which X would prefer to use the least expensive (*egress AS, egress PoP*) pair for routing a particular flow. Network X can use the traffic costs, current routing, and topology information to select the least expensive (*egress AS, egress PoP*) pair for each flow, thus minimizing the total cost. For each flow, the operator must account for the total cost (interconnect and backhaul) for routing that flow via each of the (*egress AS, egress PoP*) pairs. It is not necessarily sufficient to use the egress PoP that incurs the lowest backhaul cost for a flow (if, for instance, that egress PoP is closest to the ingress PoP), because the interconnect costs of exchanging traffic with ASes at that egress PoP may be high. The operator can also introduce capacity and performance constraints to avoid rerouting traffic in ways that might create congestion or introduce high performance penalty. We now describe this example and its formulation in more detail and evaluate simple greedy heuristics to solve the problem.

Formulation Our formulation uses the traffic cost model from Section 2 as an input, along with additional routing information from the network. Given a network topology, routing information, and the set of s, d flows, our goal is to reduce the total cost of routing the flows while satisfying constraints on backhaul and interconnect links. Note that the optimization assumes that the network topology and neighbor AS relationships are fixed, hence the optimization only deals with optimizing the usage-based cost (C_U) of carrying traffic flows in the network.

Inputs The input to the optimization problem is the complete set of s, d flows routed on the network, and the fully parameterized cost model that determines the usage-based interconnect and backhaul cost for routing each flow (as defined in Section 2). In addition, the optimization requires information about the capacity of the interconnect links and backhaul paths in the network. The optimization also takes as input information about availability of a neighbor at the different PoPs in the network. We obtain the (*egress PoP, egress AS*) pair for each flow f at PoP p based on the destination d of the flow from the routing table dumps at each PoP.

Output The desired output is the routing configuration that minimizes the total cost of routing every flow. This takes the form of a *mapping*, which defines the ingress AS, ingress PoP, egress AS and egress PoP for every flow f . The realization of the routing decisions may be complicated, depending on how the network is configured, but, fortunately we find that that most of the cost benefits can be achieved by routing only a small fraction of the flows.

Although the formulation we have described can determine both the ingress interconnect and egress interconnect, there are important differences between the ingress and egress mappings. Changing the ingress AS or PoP for a flow depends on neighboring and remote networks. For example, attempting to change the ingress AS for a flow f using AS path prepending assumes that remote ASes prefer short AS paths and do not override this behavior with local preference. Changing the ingress PoP for a flow involves negotiating arrangements (*e.g.*, hot-potato routing) with neighboring ASes. On the other hand, given a destination prefix, the network has complete flexibility in choosing to route traffic towards that prefix via any neighbor AS that advertises that prefix. Similarly, the network can choose to send traffic to multiple PoPs where a particular neighbor AS may be peering. Given that the network cannot deterministically control the ingress mapping for a flow, our model retains the ingress mapping. The model assumes that the network can only control the egress mapping for a flow f ; in other words,




Figure 3: Example: Planning Decisions.

the network can route the traffic internally and choose the appropriate egress AS or PoP to reduce the cost of routing that traffic.

Objective Function The objective is to minimize the total cost of network traffic. From the cost model developed in Section 2, there are two types of costs associated with each flow: interconnect and backhaul costs. Thus, the *objective* is to minimize the total usage-based cost over all the flows in the network.

Constraints The routing configuration that minimizes costs must also satisfy capacity constraints in the network. We consider two categories of capacity constraints: interconnect link capacities and backhaul link capacities. Operators can add other constraints, e.g., related to performance. These can also be modelled as linear constraints but would require additional information about performance of various paths. Adding constraints will restrict the set of alternate paths for each flow, thereby reducing potential cost saving. We show here how we model interconnect link capacity constraints. If $C_i(a, p)$ is the capacity of an interconnect link with AS a at PoP p , then the total rate of flows that map to the AS a and PoP p should be less than the capacity $C_i(a, p)$. Because the optimization does not change the ingress mappings of the flows, the constraint only applies to the egress interconnect links. Formally:

$$\sum_{f: a_e, p_e = a, p} r \leq C_i(a, p) \quad \forall a, p$$

Solving the Optimization Solving the optimization involves determining the egress mappings for every flow so as to minimize the total cost of all the flows, subject to the various capacity constraints. The capacity constraints restrict the amount of traffic that can be routed on a particular interconnect link or backhaul path in the network. This is similar to the bin packing problem where objects (flows in this case) are assigned to bins (links in this case) and the bin has a fixed capacity and each object has a fixed size (rate of the flow), which is NP-hard. Simple greedy approaches yield good approximate solutions for the bin-packing problem (e.g. first-fit, best fit decreasing and first fit decreasing). We present a simple greedy assignment that respects the interconnect and backhaul capacity constraints, and assigns a flow to the lowest cost path on which it can be routed while respecting the capacity constraints.

The traffic flow assignment is not a direct mapping to bin packing, so we use the following variation of the first fit decreasing strategy. We consider flows in decreasing order of their cost and assign each flow to a path that has enough backhaul and interconnect capacity, and has the least cost among all such paths. We show the results from our greedy assignment, using two different methods of attributing interconnect costs to flows, in Section 4.3.

3.2 Planning Decisions

Mapping traffic flows to their associated costs can also help to identify potential opportunities for reducing cost or increasing revenue by re-evaluating existing interconnections. A network operator may wish to evaluate the locations where the network is peering with a particular AS, or he may wish to evaluate the profitability of peering with that AS at all. We present two examples here:

Determining Peering Locations Network X can use the available cost information and, based on its current traffic demands, estimate how peering with a neighboring network at additional locations

might affect the overall cost of carrying traffic. Such a decision will depend on the costs of transporting traffic over various backhaul links, as well as the costs of various interconnection and peering points along the path. For example, as shown in Figure 3(a), network X might have a significant amount of ingress traffic near a certain location, A , that is also destined for locations near A in a neighboring network. Depending on the cost of interconnecting at A relative to backhaul costs (i.e., if interconnection is less expensive than carrying the traffic to B via backhaul links), it may make sense for network X to also peer with this neighboring AS at location A . If, on the other hand, a second peering location B offers more attractive pricing (e.g., port costs and exchange fees could be lower at B), it may be more profitable to simply send all of the traffic to the neighbor through a peering location at B .

Evaluating Existing Peering Contracts The peering relationships of network X are beneficial to X when they are created. Over time, network X may connect to additional peers, or the traffic flow and interconnection costs may change sufficiently for the peering link to no longer be beneficial to network X . An operator at network X may want to periodically re-evaluate the value of peering with a certain AS. Figure 3(b) shows network X and its peer P . When X created a peering relationship with P , it may have been less expensive to route traffic destined to D via P , as opposed to using a transit provider, T . Over time, however, transit provider T might offer a better price, or the backhaul cost of routing traffic to T might decrease; X might add another peer Q that can route traffic to the same destination D . The operator of network X must continually re-evaluate whether there is value in continuing to peer with P . For example, the operator may wish to compute the cost for routing traffic towards a customer AS, D , if it depeered P and instead routed this traffic over either T and/or Q . In Section 4.4, we describe how to evaluate the value of an existing peering contract.

4. EVALUATION

We now evaluate the different applications of the traffic cost model, as described in Section 3. We evaluate the greedy algorithm to reduce cost of routing traffic in the network and the two planning decision examples described in the previous section.

4.1 Setup

We use traffic flow statistics, routing data, and topology data from a large access provider in the UK. The traffic statistics consist of NetFlow data from a weekday in July 2009; statistics are based on a 1-in-1000 packet-sampling rate. The routing data consists of full BGP routing table dumps from the edge routers and the complete IS-IS topology for the network.

We extract flow-level statistics from the NetFlow data that gives us the traffic (in bytes) between every s, d pair, where s is the source prefix and d is the destination prefix. The s, d pair defines a flow f ; we compute its rate, r , by dividing the total bytes transferred by the duration of our measurement. Combining this flow-level data with available BGP and IGP routing data, we obtain the path in the network for each flow f . We set the unit traffic costs to be between \$1 and \$10 per Mbps, which corresponds to publicly available data on the current prices for Internet transit. We evaluate three different scenarios for different relative prices of backhaul and interconnect cost:

- **Backhaul \approx Interconnect:** We scale the unit backhaul cost $U_b(p_i, p_e)$, to be in the same range as the unit interconnect cost.
- **Backhaul \gg Interconnect:** Represents the case where transit prices and peering costs are very low, due, possibly, to competition in the transit market or the presence of IXPs.




Figure 4: Cumulative fraction of savings for the fraction of reassigned flows using greedy heuristic with capacity constraints.

- **Backhaul \ll Interconnect:** Represents the case where transit prices and peering costs are much higher than backhaul costs. This could be in regions where certain ISPs have monopolies in the transit market, and peering opportunities are limited.

These three scenarios can represent the cost structure for links in different types of networks. For example, transit providers may have high backhaul but lower interconnect costs; on the other hand, content providers might have relatively higher interconnect costs.

4.2 Shapley Value Computation

We estimate the Shapley values for a subset of flows at every interconnect link in the network. The computation of Shapley values quickly becomes computationally infeasible even for a small number of flows, so we use the approximation technique described in previous work [11]. The complexity of estimating Shapley values for a given interconnect is $O(|flows|^2 * K)$, where K is the number of permutations used. For a fixed number of flows, the smaller the value of K , the faster the computation, but the higher deviation from the true Shapley values. We computed the coefficient of variation (CV) for the Shapley values for a fixed set of flows at a particular interconnect for $K = 10, 100, 1000$. We found that the CV is $> 100\%$ for most flows for $K = 10$, between 50% and 100% for most flows for $K = 100$ and $< 30\%$ for all flows for $K = 1000$, with a median of 11%. Although it is computationally infeasible to calculate the ground truth, these results show that for $K = 1000$, the permutation does not significantly affect the Shapley value estimate. We use this value of K in our evaluation.

4.3 Greedy Cost Reduction

We evaluate the greedy cost optimization described in Section 3.1. We aim to demonstrate the benefits of using a simple greedy strategy. We assume that the network operator has a target utilization of 30%. For more than 73% of flows, an alternate path is available, but a small fraction of flows have more than two alternate paths.

We evaluate cost savings using two different techniques for calculating the flow contribution to the interconnect costs. Figure 4 shows the cost savings for the three different scenarios of backhaul and interconnect costs. When using *linear function* (Figure 4(a)), moving only the most expensive 10% of flows that have alternate paths that are less expensive achieves 68% of the maximum possible saving (in the case of Backhaul \approx Interconnect). This result is significant, because the network operator may not wish to reassign many flows, since doing so might require large changes in routing configuration or disrupt a large fraction of traffic. When using *Shapley values* (Figure 4(b)), moving the most expensive 30% of the flows achieves 65% of the maximum possible saving (in the case of Backhaul \approx Interconnect). Because the greedy strategy as-

signs flows in the order of their original cost and also obeys the capacity constraints, some flows may traverse routes that are more expensive than the original route, but significant cost savings are possible regardless.

Breakdown of Cost Savings We evaluate the relative contribution of interconnect and backhaul cost savings to the total cost savings, for each cost scenario. We find that the relative contributions depend on the particular cost scenario. In the Backhaul \ll Interconnect case, almost all the cost savings are because of reduction in the interconnect cost for the reassigned flows. For the Backhaul \gg Interconnect case, there are a number of reassigned flows for which interconnect cost actually *increases*. For these flows, however, the backhaul cost savings are sufficient to give a positive total cost saving. This finding highlights the importance of optimizing *both interconnect and backhaul costs* for flows. Considering interconnect costs in isolation might obscure certain cost saving opportunities.


4.4 Peering Decisions

Now, we describe the evaluation of two “what-if” scenarios (described in Section 3.2). We use the linear function for calculating interconnect costs for these examples.


Peering Location Evaluation For an existing peer A , we consider each PoP where the network does not currently peer with A , and try to route existing flows (which use A as the egress AS) via the new PoP. We calculate the total cost of routing flows after adding the new PoP and pick the additional PoP which gives the maximum cost savings for the peer A . For our analysis, we ignore any capacity constraints while reassigning flows and assume that A is available for peering at each additional location. It is easy to extend our method to include capacity constraints and the availability of peer A at the new peering location.

Figure 5(a) shows the CDF of savings by selecting one additional peering location for each existing peer. We find that when Backhaul \gg Interconnect, the benefit of adding a peering location with an existing peer depends on the peer. For about 35% of existing peers, there is no benefit from adding an additional peering location, perhaps because the network already connects to certain peers at the best possible PoP. On the other hand, for some peers, adding an additional peering location saves $> 80\%$ of the current cost of routing traffic via that peer. This could happen, for example, if most of the traffic that X routes via A enters X at a certain PoP p_i . If there is an egress PoP p_e close to p_i , then adding a peering location with A at p_e will yield significant backhaul cost savings.

Existing Peer Evaluation A network X may wish to periodically re-evaluate the value of a peering link with an existing peer A . We describe a method using which X can estimate the value of a peering link with a neighbor. For a neighbor AS A , we try to reassign



(a) New PoP selection for a neighbor



(b) Depeering with an existing neighbor

Figure 5: Peering Decisions: CDF of benefits (% savings) for different scenarios of the cost function.

each flow that was routed via A to some other AS. If a flow cannot be routed via any other AS, then we assume that that flow must be routed via a transit provider, and charge it by the maximum rate. We then calculate the difference in total cost by reassigning the flows which used A as egress. This is the net saving for network X by depeering network A . If the net saving is negative, then it makes sense for network X to keep peering with A , while if the net saving is positive, then network X would benefit from depeering A .

Figure 5(b) shows the CDF of the net savings from depeering each neighbor. The CDF is skewed, showing that some networks are extremely beneficial. Some peers for which the net saving is positive would benefit by depeering these peers. When Backhaul \gg Interconnect, a smaller fraction of peers gives a net saving. The intuition is that moving traffic from an existing peer to another peer (which may offer cheaper interconnect) may lead to a large increase in backhaul cost, and no net saving. But, when Backhaul \ll Interconnect, we see a larger number of peers which X can benefit by depeering. This is because moving flows from an existing peer to other peers offering less expensive interconnect is beneficial, even if it involves carrying the traffic for longer distances on less expensive backhaul paths.

5. SUMMARY AND DISCUSSION

We have developed a holistic traffic cost model for associating costs for individual traffic flows in the network. We have shown how a network operator could use the cost model to reduce the cost of routing traffic both across backhaul links and at interconnection points in a network. Although network operators currently apply some heuristics to control the cost of network traffic, they lack a holistic cost model that incorporates all contributors to the cost of forwarding individual traffic flows. This paper presents the first such cost model for network traffic, which we believe could serve as the foundation for many tools to help network operators control network costs and rationalize planning decisions.

Operators could also use our cost model to jointly optimize cost and performance, as previous work has done for stub networks [7] and for datacenters [12]. The cost model could also be integrated with a configuration tool that helps an operator determine a set of configuration changes that could achieve the appropriate re-mapping of traffic flows; alternatively, a controller (as in networks based on RCP [1] or 4D [8]) could directly map flows onto the appropriate paths. A central controller could collect all inputs to the optimization, compute the optimal mapping of traffic flows to respective paths, and implement these decisions by pushing them to the routers directly.

There are a number of directions that require additional work. Instead of random interconnect costs we used in our evaluation, it would be useful to work with real cost data. We used Shapley val-

ues for distributing interconnect costs across flows; if the number of flows is large, computing these values is expensive. Another important avenue for future work is to incorporate feedback in this model. When a network reroutes flows to reduce cost, the change can affect incoming traffic patterns, making the resulting cost sub-optimal. This change might trigger more routing changes, which could again affect traffic patterns. We plan to explore the conditions under which our cost-based routing optimization converges to a stable routing configuration. We are also exploring the feasibility of a tool that continuously monitors traffic patterns and cost information and re-optimizes the routing to reduce traffic costs in the network.

6. REFERENCES

- [1] M. Caesar, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe. Design and implementation of a routing control platform. In *Proc. 2nd USENIX NSDI*, Boston, MA, May 2005.
- [2] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In *Proc. IEEE INFOCOM*, Barcelona, Spain, Mar. 2006.
- [3] Why care about Transit Pricing? http://drpeering.net/a/Peering_vs_Transit_The_Business_Case_for_Peering.html, 2009.
- [4] Internet Video Traffic. <http://tinyurl.com/2777mhl>.
- [5] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and K. van der Merwe. The case for separating routing from routers. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Portland, OR, Sept. 2004.
- [6] N. Feamster and J. Rexford. Network-Wide Prediction of BGP Routes. *IEEE/ACM Transactions on Networking*, pages 253–266, Apr. 2007.
- [7] D. K. Goldenberg, L. Qiu, H. Xie, Y. R. Yang, and Y. Zhang. Optimizing cost and performance for multihoming. In *Proc. ACM SIGCOMM*, pages 79–92, Portland, OR, Aug. 2004.
- [8] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang. A clean slate 4D approach to network control and management. *ACM Computer Communications Review*, 35(5):41–54, 2005.
- [9] Wholesale Internet Bandwidth Prices. http://www.circleid.com/posts/wholesale_internet_bandwidth_prices/, 2008.
- [10] L. Shapley. A Value for n-Person Games. *Classics in Game Theory*, page 69, 1997.
- [11] R. Stanojevic, N. Laotaris, and P. Rodriguez. On economic heavy hitters: Shapley value analysis of the 95th-percentile pricing. In *Proc. Internet Measurement Conference*, Melbourne, Australia, Nov. 2010.
- [12] Z. Zhang, M. Zhang, A. Greenberg, Y. C. Hu, R. Mahajan, and B. Christian. Optimizing Cost and Performance in Online Service Provider Networks. In *Proc. 7th USENIX NSDI*, San Jose, CA, Apr. 2010.