Green ISP Networks via Hybrid SDN

Krishna P. Kadiyala
Telecommunications Engineering Program
The University of Texas at Dallas
Richardson, Texas 75080
Email: krishna.kadiyala@utdallas.edu

Jorge A. Cobb
Telecommunications Engineering Program
The University of Texas at Dallas
Richardson, Texas 75080
Email: cobb@utdallas.edu

***Abstract*—The development of Software Defined Networking (SDN) has introduced many benefits to legacy networks, and has become an appealing option for Internet service providers. However, doing a complete overhaul of an existing service provider network in an attempt to transform it into an SDN network is a significant economical, managerial, and technical challenge. To alleviate this, research is being performed on *hybrid* SDN networks, in which only a few routers are retrofitted to become capable of supporting SDN. In this paper, we study the impact that hybrid SDN can have on the electrical power usage of a service provider network. Most service providers have redundant routers for reliability purposes and for accommodating changes in traffic over time. These redundant routers and links are always powered on, wasting valuable energy. This waste of energy can be mitigated by identifying such routers and shutting them down. However, turning off routers reduces the number of routing paths available to the intra-domain routing protocol, which in turn has the consequence of having an unbalanced traffic load in the egress links of the service provider. This increase in link utilization experienced by some egress links leads to packet losses and long packet delays. To alleviate this, we propose retrofitting a few legacy routers to become SDN routers. By introducing just a few well-placed SDN routers, the routing flexibility increases within the network. This allows for a larger number of routers to be shutdown without exceeding a desired upper bound on link utilization. We present heuristics for choosing which routers should be augmented with SDN capabilities, and we evaluate via simulation their impact on the number of routers than can be powered down.**

Keywords–Software-defined networking; Traffic engineering; Load balancing.

I. INTRODUCTION

A current concern in society is minimizing the use of energy. This concern has reached various aspects of computing, and it is often referred to as green computing. In this paper, we focus on reducing the energy usage of enterprise networks, in particular, Internet Service Providers (ISPs).

Traditional ISP networks are over-provisioned to accommodate for unforeseen link/router failures and sudden traffic bursts. These redundant routers and links are always powered on, even though they may not be used to their full capacity at all times. Thus, the energy consumption of the network remains consistently high while the network resources remain under-utilized. By identifying such routers and shutting them down, we are able to reduce the energy consumption of the network to a certain extent.

However, an important consideration when shutting down internal routers of an ISP is that the reduced network must satisfy the traffic demand without over-provisioning the peering (i.e., external) links of the ISP. That is, the links joining the ISP to other Autonomous Systems (AS) must not reach high utilization levels. This is critical due to the fact that the peering links of an AS have been shown to be bottlenecks and are often the cause for congestion [1]. Shutting down routers may lead to high utilization of the peering links of the AS since all the traffic may have to exit through only a few reachable egress routers.

Balancing the traffic over the egress links is dependent upon the routing paradigm of the ISP. The typical routing paradigm consists of moving the transit traffic along the least-cost path. That is, each internal link has a cost, and the traffic arriving via an ingress link exits via the egress link where the total cost of the links traversed is the least. This is commonly known as hot-potato routing (HPR). Although HPR minimizes the cost of transit traffic through the ISP's network, it does not take into consideration the load on the peering links of the ISP, which, as mentioned above, are the most likely to become congested.

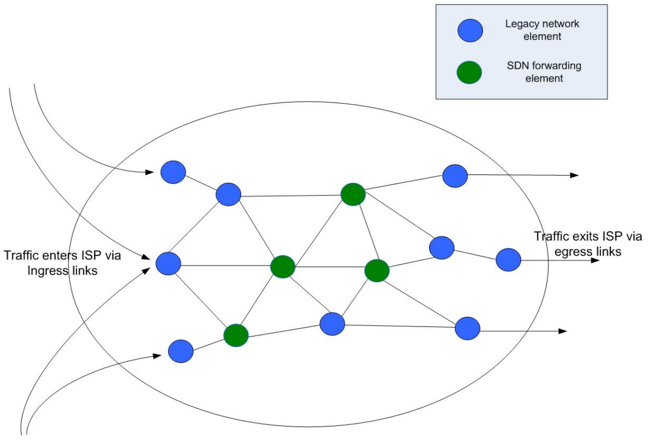

Figure 1. Hybrid Service Provider Network

The flexibility of routing inside the ISP's network is greatly limited by HPR. One possibility to free the ISP from the drawbacks of HPR is the Software-Defined Networking (SDN) paradigm. The purpose of introducing SDN routers into the legacy infrastructure is to make this internal routing less rigid, thereby enabling us to achieve greater load balancing on the egress links.

However, the complete overhaul of an existing ISP network into an SDN network is an economical, managerial, and technical challenge. To this end, ISPs may choose to transition to SDN in incremental steps, migrating from traditional networks to hybrid networks that are a combination of legacy and SDN routers. An example of such a network is shown in Figure 1. Introducing SDN incrementally reaps the potential benefits of SDN while imposing a smaller economical and managerial cost.

The increase in routing flexibility introduced by SDN will have a positive effect on the redistribution of traffic that must be performed when a legacy router is shutdown. Keeping this in mind, the problem we explore in this paper is to leverage the presence of SDN routers in the legacy network to shutdown as many routers as possible while ensuring that the maximum link utilization (MLU) at the egress links remains less than 100%. The presence of SDN routers ensures that traffic can be distributed more evenly between egress links, since it adds more flexibility to the internal routing decisions, and it allows the possibility of routing traffic over paths that would otherwise not be chosen by an intra-domain routing protocol such as OSPF.

A. *Results and Contributions*

In this paper, we address the problem of minimizing the total power consumed by the network subject to the constraints that the traffic demand is met and that the MLU at the egress links is bounded. The goal is to minimize power consumption by shutting down routers in the network. Both legacy routers and SDN routers may be shutdown. The only restrictions we impose is that ingress routers cannot be shutdown, and a router cannot be shutdown if by doing so the network becomes partitioned.

We formulate the *green hybrid SDN problem*, and we propose a heuristic to choose which routers are to be upgraded with SDN capabilities. We show through simulations that our heuristic outperforms the random selection of routers. That is, carefully selecting which routers are upgraded with SDN increases routing flexibility in such a way that a larger number of routers can be shutdown without exceeding the desired MLU of the egress routers.

The rest of the paper is organized as follows. In Section II, we review related works that independently address the energy efficiency problem in legacy, hybrid, and SDN-only networks. In Section III, we review some background in inter-AS vs intra-AS routing, and also review our earlier work on minimizing egress link utilization in hybrid SDN networks. In Section IV, we present the green hybrid SDN problem, and our heuristics are presented in Section V. Simulation results are presented in Section VI. Concluding remarks are given in Section VII.

II. RELATED WORK

We first review work related to incremental SDN deployment and traffic engineering in hybrid networks. Hybrid SDN networks, in which legacy routers co-exist with SDN nodes, have been an interesting field of study in the SDN community starting with the ideas discussed in [2]. The problem of traffic engineering in a hybrid enterprise network has also been studied in recent years. The first paper to address network performance issues in an incrementally deployed SDN network, [3], explores how SDN can be leveraged to dynamically manage traffic in a hybrid environment.

The traffic engineering (TE) problem in an SDN/OSPF environment is studied in [4], where the goal is to optimize OSPF link weight settings to lower the MLU in the network. Optimizing TE performance over all the network links in a hybrid-SDN environment is studied in [5]. In [6], the maximum flow problem in hybrid SDN networks is explored and an FPTAS is proposed for solving it. Further, [7] focuses on ISP networks with the TE objective of minimizing the MLU over its peering links.

Energy efficiency in ISP networks has been significantly explored in various network scenarios including legacy networks, hybrid SDN networks, and pure SDN networks. We first look at the most relevant studies in the legacy network scenario.

To reduce the total power consumption in a legacy network, numerous studies propose shutting down links and/or entire routers in the network, based on the ideas discussed in [8]. The idea explored in [9] and [10] is to increase energy savings by turning off links and routers in the network subject to QoS constraints such as MLU. In the former, simple heuristics are presented to selectively turn off links and routers while in the latter, a new algorithm based on the power consumption of nodes and links is proposed. The goal in [11] is to shut down cables in bundled links while ensuring that there is enough room to satisfy the traffic demand.

In [12], the authors propose a routing algorithm that precomputes loop-free next-hops for each primary next-hop to effectively detour around links with low traffic load, allowing for traffic aggregation onto links for increased power savings. The study in [13] minimizes energy consumption by turning off unused links in cabled bundles and nodes, while ensuring the traffic demand for each session is satisfied. The authors in [14] present a technique that uses a scalable, online technique to spread the load among multiple paths so as to increase energy savings while achieving the same traffic rates as the energy-oblivious approaches.

In [15], the authors propose a framework that identifies energy critical paths and uses an online TE mechanism to deactivate and activate network elements on demand. [16] proposes a mechanism that maximizes the number of links and/or line-cards that can be put to sleep under constraints such as link utilization and packet delay. This mechanism relies on a centralized controller to make the TE decisions and disseminates the decisions to routers, which then turn on/off line cards and ports as needed. Finally, [17]–[19] are a few other green networking studies in legacy networks.

Next, we briefly go over the energy efficiency research in hybrid SDN networks.

The authors in [20] propose an SDN-based energy-aware routing and resource management model in which the SDN controller uses pre-established multi-paths and performs routing and admission control based on these paths. These paths are turned on/off based on traffic load for energy savings. In [21], the authors propose a hybrid energy-aware TE algorithm which determines the optimal setting for the OSPF link weight and the splitting ratio of SDNs to enable aggregating traffic onto partial links and turning off underutilized links to save energy. [22] focuses on finding the most appropriate percentage of legacy IP nodes to be upgraded to SDN with the goal of putting to sleep links and/or SDN nodes where applicable. The study also gives a selection criterion for selecting SDN nodes to increase energy efficiency of the network. Further, the study in [23] determines the minimum-power network subsets that can satisfy the traffic demand and shuts down unnecessary SDN switches and links.

Studies that focus on energy saving in pure SDN networks include [24]–[28]. The approaches here include modifications/extensions to the OpenFlow protocol and heuristics to aggregate traffic and/or minimize the number of active SDN elements required to satisfy traffic demands.

To the best of our knowledge, our study is the first to take into consideration inter-AS traffic engineering while shutting down nodes in hybrid SDN networks for reducing total energy consumption of the network.

III. BACKGROUND

We next review some background in inter-AS vs. intra-AS routing, and also our earlier work on minimizing egress link utilization in hybrid SDN networks.

A. Intra-AS vs. Inter-AS Routing

An *autonomous system* (AS) is a group of networks (i.e., IP prefixes) that is controlled by a single administrative entity, such as a university, a company, or an organization. Currently, the Internet has over 80,000 autonomous systems (ASms). Figure 2 shows an AS M that has four neighboring ASms. In this figure, we assume that IP prefix 210.1.0.0/16 is reachable via both ASms A and B (perhaps several AS-hops away), while IP prefixes 200.1.0.0/16 and 220.1.0.0/16 are only reachable via AS A and AS B, respectively.

Border routers exchange prefix reachability information with each other via the BGP protocol. E.g., border router r_1 in AS M learns about IP prefix 210.1.0.0/16 via its neighboring router r_A in AS A, and border router r_B learns about this same prefix via its neighboring router r_B in AS B. It is possible that these IP prefixes are located many AS-hops away from ASms A and B. For the purposes of this paper, we only consider the fact that IP prefixes are being advertised by border routers, and ignore the number of AS-hops to reach them.

We assume that interior routers (i.e., routers not located at the border of the AS) do not speak BGP. This is commonly the case for a medium-sized AS. Thus, interior routers are not aware of the existence of other ASms. They do, however, run an Internal Gateway Protocol (IGP), such as RIP or OPSF, to find a path to every IP prefix available within its own AS.

To allow interior routers to find a path to the external prefixes, such as 210.1.0.0/16, the border routers employ *route redistribution*. That is, border routers advertise the external prefixes over the IGP as if these prefixes belonged to a link directly attached to them. In this way, each internal router can reach an external prefix by following the shortest path to any border router that advertises the prefix. In Figure 2, routers r_1, r_3, r_5 and r_6 will reach prefix 210.1.0.0/16 via neighboring router r_A, while routers r_2 and r_4 will reach this same prefix via neighboring router r_B.

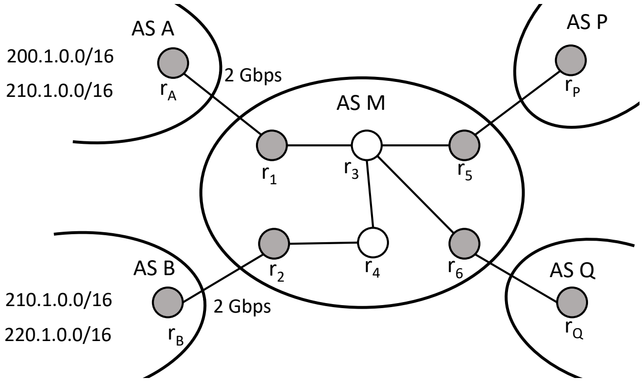

Figure 2. Autonomous systems.

B. Optimizing Egress Link Utilization via Hybrid SDN

As discussed in the next section, shutting down routers will increase the utilization of the egress links. Therefore, finding a method to distribute the traffic evenly over the egress links is beneficial to this problem. Since we assume that the network has SDN routers, we choose to use *SDN Egress Selection* (SES), which we introduced in earlier work [7], to minimize the utilization of egress links. We briefly overview this technique below.

Consider again Figure 2. The links between AS M and its neighbors A and B are labeled with their respective bandwidth, and assume that the IGP uses minimum-hop routing.

Let AS P send 1.5 Gbps to prefix 200.1.0.0/16. This traffic is only advertised by r_1 and thus it must exit via r_1. Let AS Q send 1.5 Gbps to prefix 220.1.0.0/16. Similarly, this must exit via r_2 because only r_2 advertises it. Finally, let each of P and Q send 0.3 Gbps to prefix 210.1.0.0/16. Due to minimum-hop routing, this 0.6 Gbps will exit via r_1, causing this egress link to overflow. Consider now turning a single router into an SDN router, in particular, r_3. This allows r_3 to divert traffic in any way we choose. Let r_3 forward the traffic for prefix 210.1.0.0/16 towards r_1 if it originates from AS P, and towards r_2 (via r_4) if it originates from Q. In this way, both links receive only 1.8 Gbps each, reaching a utilization of 90%.

The SES problem has similarities with the NP-hard problem of minimizing the makespan in unrelated parallel machines [29], [30]. This scheduling problem consists of *m* parallel machines and *n* independent jobs, such that, processing job *j* on machine *i* requires time $p_{i,j}$. The makespan of a schedule is the maximum total time used by any machine. The objective is to find a schedule that minimizes the makespan. Note that we can map the SES problem to the above scheduling problem by considering each egress link to be a machine, and each traffic flow from an ingress router to a destination prefix to be a job. The processing time of a flow at an egress is set to either the bandwidth of the flow or infinity, depending on whether or not the routing of the flow via minimum-hop routing plus SDN re-routing reaches that egress router.

A 2-approximation solution to the makespan problem is given in [29], [30]. Given the specific nature of the SES problem, we have shown in [7] that the rounding obtains a solution that is very close to optimal, and thus, much smaller than the theoretical bound of twice the optimal.

IV. THE GREEN HYBRID SDN PROBLEM

We consider an AS where each router is either a border router or an interior router. Border routers are divided into two sets: *ingress routers* and *egress routers*. Interior routers are also divided into two sets: *SDN routers* and *legacy routers*. Each egress router r has an egress link of capacity $C(r)$.

A traffic *flow* $f(i,p)$ corresponds to the traffic from ingress router i destined for IP prefix p. Each flow $f(i,p)$ has a *demand*, $D(i,p)$, that corresponds to the amount of traffic of the flow. A traffic flow exits the AS via a single egress router, i.e., we assume that a flow cannot be split among multiple egress routers. For each IP prefix p and egress router e, $avail(e,p)$ is true if and only if e received an advertisement for p from its neighboring AS. Thus, flow $f(i,p)$ can only exit the AS via some egress e where $avail(e,p)$ is true.

A sequence of routers, $r_0, r_1, \ldots, r_n$, is said to be a *hybrid routing path* iff r_0 is an ingress router, r_n is an egress router, and for each r_i, $0 \leq i < n$, r_{i+1} is the next-hop router along the IGP path from r_i to r_n, or r_i is an SDN router whose neighbors include r_{i+1}.

We assume that there exists an SDN controller node that determines the forwarding tables of the SDN routers. The SDN controller is assumed to be aware of the network topology and the paths chosen by the IGP. E.g., the IGP could be OSPF, and the SDN routers forward to it a copy of the link-state advertisements that they receive. The controller is also aware of the traffic matrix either directly from the network operators or via some interaction with the ingress routers.

Our *Green Hybrid SDN problem* is as follows. Consider a legacy AS network, and let R be the set of legacy routers in this network. We are given two upper bounds. The first is an upper bound U on the link utilization on egress links. The second is an upper bound k on the number of legacy routers that will be replaced by SDN routers. The output consists of finding two sets, S and R', such that:

- $R' \subset R$. This set contains the routers that are to be powered down.
- $S \subset R$ and $|S| \leq k$. Each router in S will be replaced by an SDN router.
- For every IP prefix p, every ingress i, and every egress e, the reduced network (i.e., the network consisting of routers in $R - R'$) contains a hybrid routing path from ingress i to an egress e such that $avail(e,p)$ is true.

  Notice that, assuming that every ingress has traffic from at least one prefix, then no ingress router can be part of R'. Egress routers, along with interior routers, may belong to R'.
- The egress MLU is at most U. That is,

  $$\forall e,\ \frac{Load(e)}{C(e)} \leq U \leq 1$$

  Above, $Load(e)$ is the sum of all the traffic demands of flows that are assigned to egress router e.
- The cardinality of R' is the largest possible, i.e., the energy savings is maximized by powering down the largest possible number of routers.

Note that S and R' are not defined to be mutually exclusive. However, if a router r belongs to both sets, then r will be powered down. In this case, the SDN functionality of r is of no use, and its inclusion in S is meaningless. Thus, effectively, S and R' can be thought of as being disjoint.

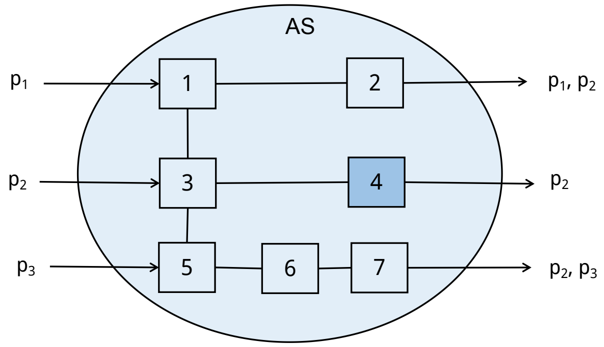

Figure 3. SDN selection example.

As routers are turned off and network paths become unavailable, it is paramount to distribute the traffic evenly among the remaining egress routers. Otherwise, the egress MLU increases beyond the desired bound. As mentioned above, we use the (SES) method that we introduced in [7] to perform this load balancing.

In [7], we assumed that the location of the SDN routers was given as input to the problem. In this paper, however, we assume that that we are free to choose which routers will be upgraded with SDN.

Different choices for set S will yield significantly different cardinalities for set R', and hence, different savings in power. As an example, consider Figure 3, where there is a single AS with three ingress routers, three egress routers, and three prefixes. Note that in light of the IGP, prefixes p_1, p_2, p_3, will exit via egresses 2, 4, and 7, respectively.

Note that 4 is the only router that can be shutdown. This is because the ingress routers cannot be shutdown (otherwise incoming traffic is dropped), and furthermore, each of p_1 and p_3 is available at a single egress, so egress routers 2 and 7 cannot be shutdown either. Finally, 6 cannot be shutdown since otherwise p_3 would not reach its egress. Thus, assume router 4 is shutdown. The IGP would then route traffic for p_2 via egress 2 since it is closer than egress 7.

Assume next that the combined traffic of p_1 and p_2 exceeds the desired utilization of egress 2. However, assume egress 7 can easily handle the combined traffic of p_2 and p_3. Let $k = 1$, i.e., we are only allowed to transform a single router with SDN. The only sensible choice is router 3, which can divert the traffic of p_2 towards 5, and hence, towards egress 7. Any other choice for the SDN router would be unable to affect the traffic, leading to an over-utilization of egress 2, and thus, router 4 would not be allowed to shutdown in this case.

The above example illustrates the importance of the heuristic to select the routers in S. Here, there will be energy savings if and only if the heuristic chooses router 3.

V. HEURISTICS

We next present several heuristics for selecting set S. Before doing so, we present the overall steps of the method.

1) First, k routers are chosen by one of the heuristics below to be transformed into SDN routers.
2) A router is chosen at random (not including ingresses) and is shutdown.
3) The network is checked to ensure that it is not partitioned, and that there is a hybrid routing path for each flow to an egress router advertising the flow's prefix.
4) The method in [7] is used to see if the SDN nodes can help in routing the traffic in such a way that the egress MLU is at most U.
5) If bound U is not violated, the router is permanently removed from the network.
6) We return to step 2 above. We end when no router can be removed from the network without violating U.

A. Diverting Traffic

For this heuristic, we pick routers that have the ability to "bump" traffic towards any egress other than the IGP-chosen egress, provided the egress is advertising the prefix. To elaborate, let p be an IP prefix, and let $\mathit{distance}(r, e)$ be the IGP distance or cost from router r to egress router e. Also, let $\mathit{exit}(r, p)$ be the egress router through which the traffic from r to p exists the AS. That is,

$$\forall e,\ \mathit{avail}(e, p) \Rightarrow \mathit{distance}(r, \mathit{exit}(r, p)) \leq \mathit{distance}(r, e)$$

Finally, let $path(r, e)$ be the IGP path from router r to egress router e. Then, we say that r is an *SDN candidate* if there is a router s and a prefix p such that: s is a neighbor of r, $exit(r, p) \neq exit(s, p)$, and $r \notin path(s, exit(s, p))$.

Consider the example in Figure 4. Let prefix p be announced by egress routers 9 and 10. Then, for routers 1, 3, 5, and 7, egress 9 is the exit router for prefix p. Similarly, for routers 2, 4, 6, 8, and 11, their exit router is 10. Thus, one SDN candidate is router 3, because $exit(3, p) \neq exit(4, p)$, and 3 is not contained in $path(4, 10)$. Similarly, routers 4, 7, and 8 are also SDN candidates.

On occasions, the number of SDN candidates can be greater than the desired number of SDN routers. If so, we simply choose randomly within the set of candidates.

B. Most Visited

This heuristic is based on the diverting-traffic heuristic. The steps are as follows:

- Apply the diverting traffic heuristic to find the SDN candidate routers.
- Calculate all the IGP paths from each input flow to its egress router.
- Rank the SDN candidate routers according to the number of these IGP paths that cross the candidate router.

For Figure 4, the most-visited heuristic results in choosing routers 4 and 8 first since they both appear in two paths: (2,10) and (11,10). These are followed by 3 and 7 since they are visited by only one path: (1,9).

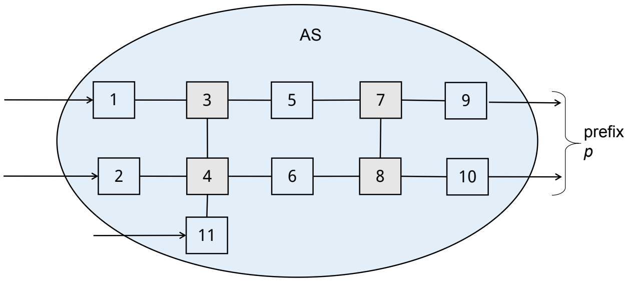

Figure 4. Network example.

C. Degree

This heuristic first identifies SDN candidates as in the diverting traffic heuristic, and the degree of the router is then used to rank the SDN candidates. We arrange routers in decreasing order of their degree. Thus, in Figure 4, router 4 is picked first. This is followed by routers 3, 7, and 8.

D. Most Traffic

For the final heuristic, we assume that we know the pattern of the traffic demand in the network. For each flow, we first calculate the shortest path from its ingress to the closest egress that advertises its prefix, i.e., the path taken by the legacy IGP. We prioritize routers by adding the traffic demands of all the flows that the IGP routes through them. Higher priority is given to those routers that handle the largest amount of traffic.

VI. SIMULATION RESULTS

We evaluate our heuristics by performing simulations on a Rocketfuel [31] ISP topology with 53 routers and 84 Intra-AS links in the network. Using the Rocketfuel topology information, we conclude that the routers that are not acting as backbone routers are acting as border routers, and we separate these border routers into 14 ingress routers and nine egress routers. We assume that the intra-domain protocol is OSPF, and that the distances are hop-based. Therefore, the shortest path to a destination is the path with the least number of hops.

We generate synthetic traffic flows from each ingress router, where the number of traffic flows through the network is the number of ingress routers times the number of prefixes. We also assume that each ingress router has incoming traffic destined to all the prefix advertised in the network. For example, in the case of 40 prefixes, this would give us 14 $\times$ 40 $=$ 560 traffic flows.

We consider a traffic scenario with 40 prefix advertisements. Typically in an ISP, the number of prefixes in the routing table can scale to large numbers. However, it has been shown that only a small fraction of these prefixes are actually responsible for a major portion of the traffic traversing the ISP network [32]. Generally, a prefix advertisement may be received and advertised by multiple egress routers in the AS. In this paper, we simply choose to advertise each prefix at all the egress border routers.

Each of the egress routers is assumed to connect with the neighboring AS with a single peering link. The capacity of the egress links is set to 1000 scaled units. The total amount of traffic generated is a fraction of the total capacity of the egress links. E.g., with nine egress links, where each link has a capacity of 1000, the total traffic generated is $f \times 1000 \times 9$, where $0 < f < 1$. In the case of six egress links, the total traffic generated is $f \times 1000 \times 6$. This total traffic is distributed randomly across the input flows, ensuring that the total traffic is exactly this amount. We have chosen f to be 0.2. Although relatively small, this amount allows us to shut down many routers in the network and observe the impact of adding SDN routers.

In each of the scenarios, we start with zero SDN routers and increment up to twenty SDN routers. Each point in our plots represents the number of routers that can be turned off averaged over ten simulation runs. Of the heuristics discussed above, the diverting traffic heuristic performed the best. No improvement was seen by adding the most visited refinement nor the degree refinement. For lack of space, we focus on comparing the diverting traffic heuristic against randomly selecting SDN routers.

We begin by presenting the diverting-traffic heuristic in Figures 5, 6 and 7. The number of available egress routers varies from three up to nine. For each of these cases, the bound U on the MLU is varied from 0.7 up to 0.9. This is followed by Figures 8, 9, and 10 with a similar configuration except that the heuristic is just random selection of the SDN routers.

An interesting phenomenon occurs in Figures 5 and 8. Even without SDN routers, having only three egress routers allows us to turn off a total of 29 routers. It appears to suggest that a lower number of egress routers is best. However, this is just a side-effect of how we chose to generate traffic. Recall that the total input traffic generated is $f \times 1000 \times$ (number of egress routers). Thus, the input traffic in the case of only three egress routers is only a third of the input traffic in the case when nine egress routers are available. Thus, as routers are shut down, the nine egress routers case has to squeeze a larger amount of traffic through a smaller number of egress links, and thus, requires a large number of SDN routers to turn off the same number of routers as the case of three egress routers.

Figures 6, 7, 9, and 10, clearly show that as the number of SDN routers in the network is increased, a larger number of routers can be shutdown without violating the link utilization bound. The number of SDN routers need not be large. For example, from Figure 7, with only six SDN routers we are able to shutdown about 23 routers of the maximum 29 possible. No simulation point, regardless of its parameters, was able to shutdown more than 29 routers.

A direct comparison of the diverting traffic vs. random is given in Figure 11. It clearly shows the superiority of the diverting traffic approach over the random approach regardless of the utilization bound chosen.

Finally, Figure 12 shows the diverting-traffic heuristic with $U = 0.75$, and a curve for each of 3, 6, and 9 routers. The figure clearly shows that, for each of these cases, as the number of SDN routers increases, the number of routers that can be shutdown also increases, as desired.

VII. CONCLUDING REMARKS

We have introduced the green SDN hybrid problem and evaluated several heuristics for it. Our goal was to ensure

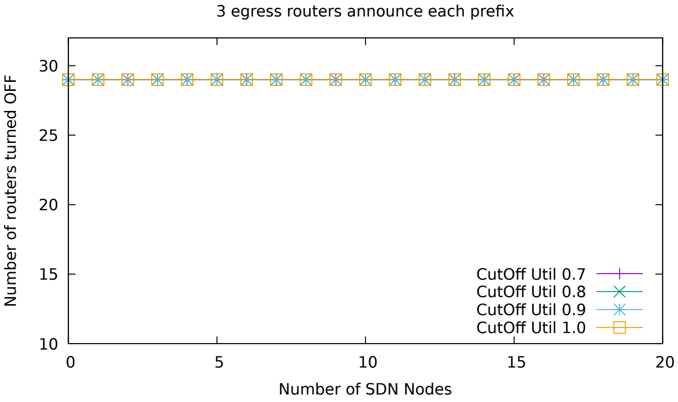

Figure 5. Diverting traffic for three egress routers

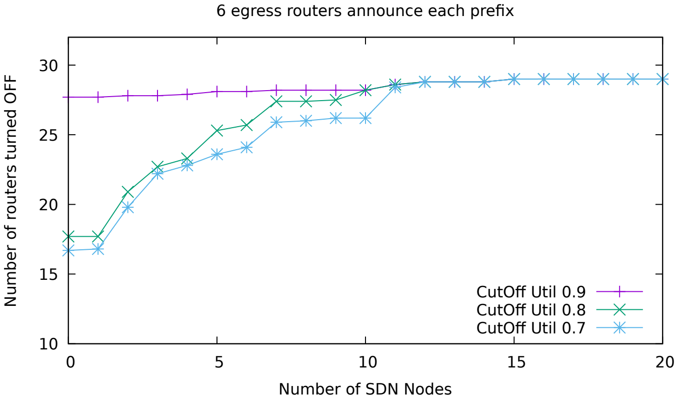

Figure 6. Diverting traffic for six egress routers

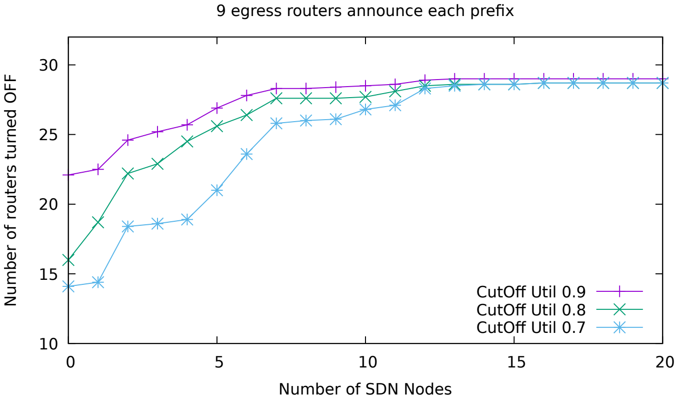

Figure 7. Diverting traffic for nine egress routers

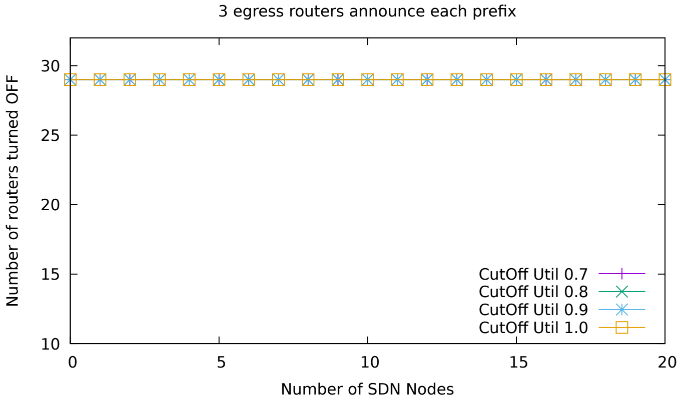

Figure 8. Random placement for three egress routers

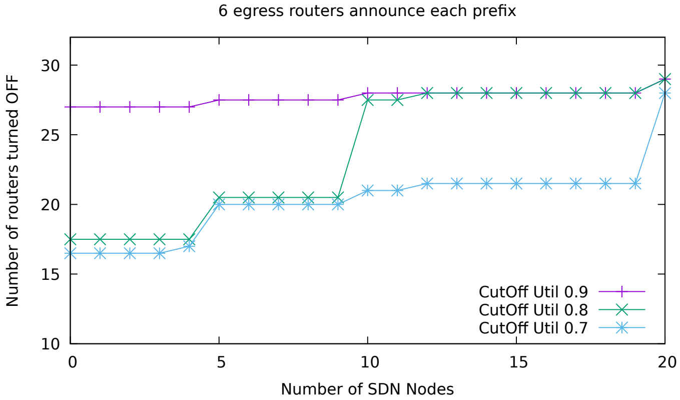

Figure 9. Random placement for six egress routers

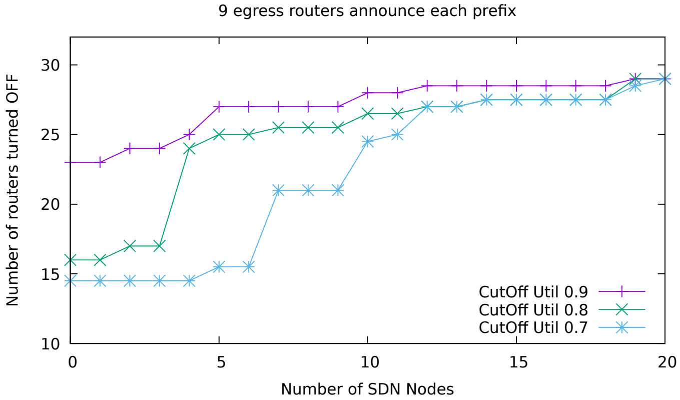

Figure 10. Random placement for nine egress routers

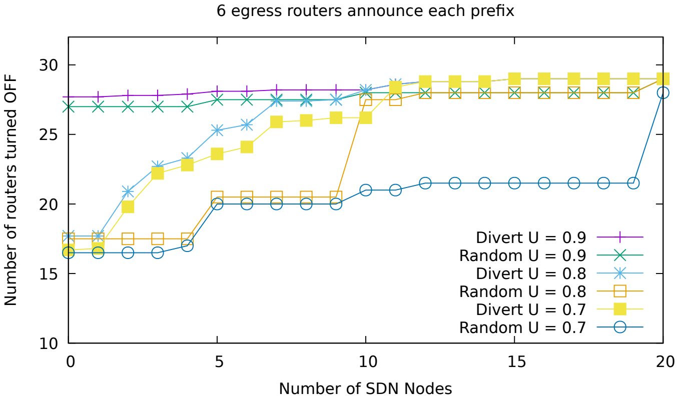

Figure 11. Diverting traffic vs. random placement

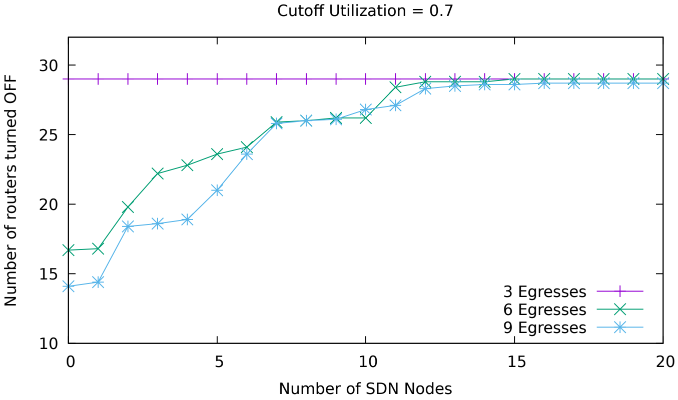

Figure 12. Diverting traffic from three up to nine egress routers

egress links are not over utilized due to the diverting of all the traffic onto a few egress links. We show through our simulations that it is possible to achieve significant energy savings and maintain a bounded link utilization with only a few SDN routers. We also show that in such a hybrid network, the location of the SDN routers play an important role in maximizing energy savings.

There are several directions possible for future work. We plan to continue to investigate various heuristics and apply them to a wide variety of topologies to study their effectiveness. Also, we have assumed that the traffic load is static. If the traffic load changes over time, the SDN controller must recalculate routes and propagate them to the SDN routers. It would be beneficial to come up with a scheme that would allow a smooth transition between the old and new set of routing tables. Finally, we have assumed that all routers consume the same amount of energy, and we have not considered shutting down individual links rather than entire routers. Thus, we also plan to investigate more complex energy models.

REFERENCES

[1] T. Bressoud, R. Rastogi, and M. A. Smith, “Optimal configuration for BGP route selection,” in *In Proc. IEEE INFOCOM*, 2003.

[2] D. Levin, M. Canini, S. Schmid, a. feldmann, and f. schaffert, “Panopticon: Reaping the benefits of incremental sdn deployment in enterprise networks,” in *USENIX Annual Technical Conference*, 2014.

[3] S. Agarwal, M. Kodialam, and T. V. Lakshman, “Traffic engineering in software defined networks,” in *2013 Proceedings IEEE INFOCOM*, 2013, pp. 2211–2219.

[4] Y. Guo, Z. Wang, X. Yin, X. Shi, J. Wu, and H. Zhang, “Incremental deployment for traffic engineering in hybrid sdn network,” in *2015 IEEE 34th International Performance Computing and Communications Conference (IPCCC)*, Dec 2015, pp. 1–8.

[5] J. He and W. Song, “Achieving near-optimal traffic engineering in hybrid software defined networks,” in *2015 IFIP Networking Conference (IFIP Networking)*, May 2015, pp. 1–9.

[6] Y. Hu, W. Wang, X. Gong, X. Que, Y. Ma, and S. Cheng, “Maximizing network utilization in hybrid software-defined networks,” in *2015 IEEE Global Communications Conference (GLOBECOM)*, Dec 2015, pp. 1–6.

[7] K. P. Kadiyala and J. A. Cobb, “Inter-as traffic engineering with sdn,” in *2017 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, Nov 2017, pp. 1–7.

[8] M. Gupta and S. Singh, “Greening of the internet,” in *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, ser. SIGCOMM ’03. New York, NY, USA: ACM, 2003, pp. 19–26. [Online]. Available: http://doi.acm.org/10.1145/863955.863959

[9] L. Chiaraviglio, M. Mellia, and F. Neri, “Reducing power consumption in backbone networks,” in *2009 IEEE International Conference on Communications*, June 2009, pp. 1–6.

[10] ——, “Energy-aware backbone networks: A case study,” in *2009 IEEE International Conference on Communications Workshops*, June 2009, pp. 1–5.

[11] W. Fisher, M. Suchara, and J. Rexford, “Greening backbone networks: Reducing energy consumption by shutting off cables in bundled links,” in *Proceedings of the First ACM SIGCOMM Workshop on Green Networking*, ser. Green Networking ’10. New York, NY, USA: ACM, 2010, pp. 29–34. [Online]. Available: http://doi.acm.org/10.1145/1851290.1851297

[12] Q. Li, M. Xu, Y. Yang, L. Gao, Y. Cui, and J. Wu, “Safe and practical energy-efficient detour routing in ip networks,” *IEEE/ACM Transactions on Networking*, vol. 22, no. 6, pp. 1925–1937, Dec 2014.

[13] B. Mumey, J. Tang, and S. Hashimoto, “Enabling green networking with a power down approach,” in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 2867–2871.

[14] N. Vasić and D. Kostić, “Energy-aware traffic engineering,” in *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*, ser. e-Energy ’10. New York, NY, USA: ACM, 2010, pp. 169–178. [Online]. Available: http://doi.acm.org/10.1145/1791314.1791341

[15] N. Vasić, P. Bhurat, D. Novaković, M. Canini, S. Shekhar, and D. Kostić, “Identifying and using energy-critical paths,” in *Proceedings of the Seventh COnference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT ’11. New York, NY, USA: ACM, 2011, pp. 18:1–18:12. [Online]. Available: http://doi.acm.org/10.1145/2079296.2079314

[16] M. Zhang, C. Yi, B. Liu, and B. Zhang, “Greente: Power-aware traffic engineering,” in *Proceedings of the The 18th IEEE International Conference on Network Protocols*, ser. ICNP ’10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 21–30. [Online]. Available: http://dx.doi.org/10.1109/ICNP.2010.5762751

[17] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, “An energy saving routing algorithm for a green ospf protocol,” in *2010 INFOCOM IEEE Conference on Computer Communications Workshops*, March 2010, pp. 1–5.

[18] O. Okonor, N. Wang, Z. Sun, and S. Georgoulas, “Link sleeping and wake-up optimization for energy aware isp networks,” in *2014 IEEE Symposium on Computers and Communications (ISCC)*, June 2014, pp. 1–7.

[19] R. Bolla, R. Bruschi, A. Cianfrani, and M. Listanti, “Enabling backbone networks to sleep,” *Netwrk. Mag. of Global Internetwkg.*, vol. 25, no. 2, pp. 26–31, Mar. 2011. [Online]. Available: http://dx.doi.org/10.1109/MNET.2011.5730525

[20] M. R. Celenlioglu, S. B. Goger, and H. A. Mantar, “An sdn-based energy-aware routing model for intra-domain networks,” in *22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Sept 2014, pp. 61–66.

[21] Y. Wei, X. Zhang, L. Xie, and S. Leng, “Energy-aware traffic engineering in hybrid sdn/ip backbone networks,” *Journal of Communications and Networks*, vol. 18, no. 4, pp. 559–566, Aug 2016.

[22] J. Galn-Jimnez, “Legacy ip-upgraded sdn nodes tradeoff in energy-efficient hybrid ip/sdn networks,” *Comput. Commun.*, vol. 114, no. C, pp. 106–123, dec 2017. [Online]. Available: https://doi.org/10.1016/j.comcom.2017.10.010

[23] H. Wang, Y. Li, D. Jin, P. Hui, and J. Wu, “Saving energy in partially deployed software defined networks,” *IEEE Transactions on Computers*, vol. 65, no. 5, pp. 1578–1592, May 2016.

[24] R. Bolla, R. Bruschi, F. Davoli, and C. Lombardo, “Fine-grained energy-efficient consolidation in sdn networks and devices,” *IEEE Transactions on Network and Service Management*, vol. 12, pp. 132–145, 2015.

[25] A. Markiewicz, P. N. Tran, and A. Timm-Giel, “Energy consumption optimization for software defined networks considering dynamic traffic,” *IEEE 3rd International Conference on Cloud Networking (CloudNet)*, pp. 155–160, 2014.

[26] S. S. Tadesse, C. Casetti, C.-F. Chiasserini, and G. Landi, “Energy-efficient traffic allocation in sdn-basec backhaul networks: Theory and implementation,” *14th IEEE Annual Consumer Communications Networking Conference (CCNC)*, pp. 209–215, 2017.

[27] S. Oda, D. Nobayashi, Y. Fukuda, and T. Ikenaga, “Flow-based routing schemes for minimizing network energy consumption using openflow,” in *The Fourth International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY 2014)*, April 2014.

[28] H. Ying, L. Tao, B. N. C., and W. Wenjie, “An initial load-based green software defined network,” *Applied Sciences*, vol. 7, no. 5, 2017. [Online]. Available: http://www.mdpi.com/2076-3417/7/5/459

[29] J. K. Lenstra, D. B. Shmoys, and E. Tardos, “Approximation algorithms for scheduling unrelated parallel machines,” in *Foundations of Computer Science, 1987., 28th Annual Symposium on*. IEEE, 1987, pp. 217–224.

[30] D. Du, K.-I. Ko, and X. Hu, *Design and Analysis of Approximation Algorithms*. New York, NY: Springer, 2012.

[31] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, “Measuring isp topologies with rocketfuel,” *IEEE/ACM Trans. Netw.*, pp. 2–16, feb 2004.

[32] N. Feamster, J. Borkenhagen, and J. Rexford, “Guidelines for interdomain traffic engineering,” in *ACM SIGCOMM Computer Communication Review*, October 2003, pp. 19–30.