UvA-DARE (Digital Academic Repository)

Carbon-aware path provisioning for NRENs

van der Veldt, K.; Monga, I.; Dugan, J.; Grosso, P.; de Laat, C.T.A.M.

Link to publication

Citation for published version (APA):
van der Veldt, K., Monga, I., Dugan, J., Grosso, P., & de Laat, C. T. A. M. (2014). *Carbon-aware path provisioning for NRENs*. Paper presented at International Green Computing Conference IGCC14, .

General rights
It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations
If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: https://uba.uva.nl/en/contact, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (http://dare.uva.nl)

Download date: 26 Feb 2021

Carbon-aware path provisioning for NRENs

Karel van der Veldt
Universiteit van Amsterdam
The Netherlands
karel.vd.veldt@uva.nl

Inder Monga
Energy Sciences Network
California
imonga@es.net

Jon Dugan
Energy Sciences Network
California
jdugan@es.net

Cees de Laat
Universiteit van Amsterdam
The Netherlands
delaat@uva.nl

Paola Grosso
Universiteit van Amsterdam
The Netherlands
p.grosso@uva.nl

***Abstract*—National Research and Education Networks (NRENs) are becoming keener in providing information on the energy consumption of their equipment. However there are only few NRENs trying to use the available information to reduce power consumption and/or carbon footprint. We set out to study the impact that deploying energy-aware networking devices may have in terms of CO_2 emissions, taking the ESnet network as use case. We defined a model that can be used to select paths that lead to a lower impact on the CO_2 footprint of the network. We implemented a simulation of the ESnet network using our model to investigate the CO_2 footprint under different traffic conditions. Our results suggest that NRENs such as ESnet could reduce their network's environmental impact if they would deploy energy-aware hardware combined with paths setup tailored to reduction of carbon footprint. This could be achieved by modification of the current path provisioning systems used in the NREN community.**

Keywords*—*Circuit-switching networks, Energy-aware systems, Modeling techniques.

I. Introduction

A National Research and Education Network (NREN) provides connectivity to universities, research institution and laboratories in its own country.

Pushed by the demanding requirements of its scientific user base NRENs have always been a fertile territory for network innovation and cutting edge network architecture; but if we look at the current NREN landscape we see only some them being active in the energy management arena. ESnet, the network supporting the operations of DoE labs and institutions in the USA, is one of the NRENs that is active on this topic, for example they were one of the first ones to work at providing real-time visualization of their power data.

NRENs such as ESnet have path provisioning systems in place that allow for the provisioning of network circuits, as opposed to IP packet routing. Examples of path provisioning systems are ESnet's OSCARS [1] and GÉANT Bandwidth on Demand [2]. However, no NREN to date has any kind of energy-aware path creation service aimed to reduce power consumption and carbon footprint of their network operation. This is in clear stride with a lively research focus in the area of energy-efficient networking as we will show in Section II, which is driving clear implementations in for example data centers and clouds environment. In this paper we take the ESnet network as a use case (Section III).

978-1-4799-6177-1/14/$31.00 © 2014 IEEE

Given the current situation our work addresses two questions:

- Can NRENs reduce their power consumption and environmental impact, which would be the conditions under which this is possible and how much would be the improvement obtained with such an effort?
- Do we see a path toward the definition of *green paths* in NRENs?

To answer our first question we defined a realistic model for path provisioning of CO_2-aware paths using information that is already being offered by some NRENs (Section IV). Based on this model we developed a simulator and performed a number of experiments in order to determine the range of CO_2 reduction that may be achieved, depending on a number of parameters (Sections V and VI). We specifically investigate what the impact of current and future energy-aware technologies can have on reducing the environmental impact of NRENs. The results are presented in Section VII and are discussed in Section VIII. Our experiments suggest that, depending on network traffic, *green paths* could reduce the carbon footprint of ESnet by 5% to 23% under realistic circumstances.

In Section IX we provide an initial answer to our second question, namely in which way we envision that NRENs will be able to implement green paths across their networks.

II. Related work

It is estimated that ICT is responsible for 2-4% of the worldwide carbon emissions. A significant portion of which, about one-sixth, is attributed to telecommunications networks [3], [4].

To increase the energy-efficiency of networks the research community has focused along two research lines, in most cases separately from one another. Our work aims to integrate the two focuses within the context of NRENs:

- Hardware optimization: [5], [6], [7] focused for example on modeling and optimization techniques for energy aware routers.

- Traffic movements optimization: [8], [9], [10] have investigated data center load balancing to reduce carbon emissions. A recent proposal is the concept of Distributed Green Data Centers: interconnected small data centers powered by renewable energy sources [11]. In order to make optimal use of these systems one must decide whether it is better for the environment to execute some data- and/or compute-intensive application in a local data center or to move it to a remote location, and here the network plays an important role. This has been explored in some of our previous work [12], and in this paper we partially build upon the model discussed therein.

There is previous research in the area of reducing the carbon footprint of NRENs [13]. For example, the GreenStar Network uses the GEANT network's high speed Bandwidth-on-Demand links to create an interactive green network, which includes advanced middleware to maximize how renewable-powered resources are used. It connects data centers solely powered by renewable energy [14]. While GreenStar achieves energy-efficiency by using only renewable sources to power the network equipment, and validates this on a subset of nodes, we look at the NREN as a whole allowing for renewable and dirty energy sources.

The ESnet network has been the subject of research on energy awareness before. In a study comparable to our own [15] the authors propose the GEAR algorithm to select paths with the lowest brown power consumption. However in doing so they inadvertently consider all brown and green energy sources as producing equal amounts of CO_2, but as explained in Section III this is not the case. In our study we do not consider power consumption as such, but the carbon emissions arising from it.

III. ESNET

ESnet provides the high-bandwidth, reliable connections that link scientists at national laboratories, universities and other research institutions in the United States, enabling them to collaborate on some of the world's most important scientific challenges including energy, climate science, and the origins of the universe. Funded by the DOE Office of Science, and managed and operated by the ESnet team at Lawrence Berkeley National Laboratory, ESnet provides scientists with access to unique DOE research facilities and computing resources [16].

One service offered by ESnet is the On-Demand Secure Circuits and Advance Reservation System (OSCARS) [1], which provides virtual circuits with guaranteed end-to-end network performance. It allows network flows to be directed over a specific path, rather than routing where each packet can potentially take a different path to its destination.

ESnet also uses perfSonar [17] to publish real-time information about its network such as the amount of traffic flowing through networking elements and their power consumption.

By combining these services it is in theory possible to develop a system that creates virtual circuits in such a way that the carbon footprint of the network can be reduced. In this paper we develop a model and simulation to investigate the gains that can be made by ESnet if this CO_2-aware provisioning system would be implemented. We use a subset of the ESnet network for all our experiments (see Figure 1).

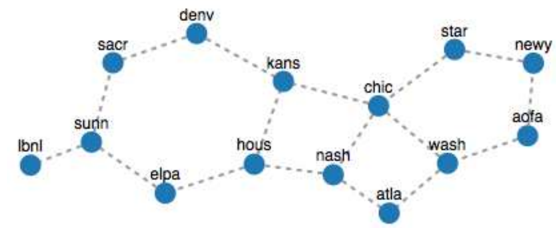

Fig. 1. Subset of ESnet topology used in our experiments [18].

We have measured the power consumption of the routers[1] of each node in the network using ESnet's monitoring system. Thus these values are realistic and reflect the real-world situation. We have observed that all routers consume a constant amount of power regardless of traffic load. That means that in the current situation it is not possible to dynamically reduce the carbon footprint. In order to dynamically reduce the carbon footprint of NRENs such as ESnet, they must make use of networking devices that implement energy-aware technologies to adjust their power consumption depending on traffic load [19], [6]. In the remainder of this paper we model and quantify what the potential benefits are of using such devices.

The minimum router power consumption we have measured is 1.5 kW, and the highest is 3.5 kW. The average is 2.1 kW.

Energy is produced from different energy sources, such as coal or solar. The amount of CO_2 emitted per kWh depends on the energy source. Table I shows a summary of the carbon footprint per energy source. In reality, different regions have different energy production mixes and therefore different carbon footprints. The Institute for Energy Research [20] collects information about energy production mixes in the United States. We combined these information sources to calculate the average carbon footprint of the States that are also present in the ESnet topology (see Table II), which corresponds to X_i in our model (see Section IV). It is important to keep in mind that this only gives an indication of the CO_2 footprint. In reality the energy sources vary over time, but this information is not readily available.

TABLE I. CO_2 OF ENERGY SOURCES [21], [22].

Energy source	gr. CO_2/kWh
Coal	950
Anthracite	870
Oil	640
Gas works gas	400
Natural gas	380
Nuclear	66
Geothermal	40
Biomass	30
Solar	22
Hydroelectric	15
Wind	10

[1] Alcatel-Lucent 7750 Service Routers

TABLE II. AVERAGE CARBON FOOTPRINT OF STATES PRESENT IN THE ESNET TOPOLOGY [20].

State	grams CO_2/kWh
New York	250
California	254
Massachusetts	459
Illinois	488
Texas	524
Tennessee	537
Maryland	571
Georgia	611
Kansas	698
Colorado	700

IV. MODEL

In this section we present our model to reduce the carbon footprint of a network. It is partially based on work we have previously published [12], but the focus here is on path provisioning.

Given a network topology $G = (V, E)$ where the vertices V represent the networking devices and the edges E represent the connections between them. We consider a scenario in which a user wants to transport N GBytes of data from one endpoint to another. Given that there are one or more possible paths to take, which path should be selected such that the carbon footprint (*grams* CO_2) is the smallest compared to every other possible path?

Equation 1 describes the model to calculate the carbon footprint C of a path R. Every element $i \in R$ represents a device in the path. C (*grams* CO_2) is calculated as the sum of the amount of energy dissipated TP_i (*kWh*) multiplied by the energy production mix X_i (*grams* CO_2*/kWh*), for each device $i \in R$. The constraint given in equation 2 ensures that provisioning the path does not exceed the maximum bandwidth capacity of any device in the path.

To calculate the best path among all possible paths, we first calculate all simple paths (without loops), calculate the C of each, and finally sort. This can be problematic because the algorithm to find all simple paths has a computational complexity of $O(|V|!)$ [23]. Of course it is possible to precompute all simple paths once and then look them up as needed. To reduce the computational complexity a path finding algorithm could be adapted to this model by considering a bidirectional graph where the edge weights are equal to P_iX_i, where i is the node a given edge is pointing to.

The model is as follows:

$$C = 8NT \sum_{i \in R} P_i X_i \tag{1a}$$

$$T = 3600^{-1} H^{-1} \tag{1b}$$

$$L_i = \frac{B_i^c + H}{B_i^m} \tag{1c}$$

$$P_i = P_i^S + P_i^D (1 - d_i) + P_i^D d_i U(L_i) \tag{1d}$$

Given the constraint:

$$\forall i \in R, 0 \leq L_i \leq 1 \tag{2}$$

Where:

- C is the total carbon emissions emitted to transport $8N$ Gbits over path R (*grams* CO_2).
- T is the time to transport 1 Gbit (*hours*).
- H is the throughput at which to transmit the data (*Gbit-s/s*).
- P_i is the total power consumption of the i-th device in R (*kW*).
- P_i^S is the power consumption of devices that consume a static (unchanging) amount of power, such as DWDMs[2] (*kW*). However since we only consider routers in this paper we use $\forall i, P_i^S = 0$.
- P_i^D is maximum dynamic power consumption of the i-th device (*kW*), of which the portion d_i of its total power depends on the traffic load L_i.
- d_i is the dynamic portion of P_i^D $(0 \leq d_i \leq 1)$. However in this paper $\forall i, d_i = d$.
- B_i^c is the amount of throughput currently pushing through the i-th router (*Gbit/s*).
- B_i^m is the maximum bandwidth capacity of the i-th router (*Gbit/s*).
- L_i is the load the i-th device would have after provisioning path R on it. If the constraint given in Equation 2 is violated it means the maximum capacity of one or more routers would be exceeded by provisioning this path.
- $U(L_i)$ is a function that maps the load of the dynamically powered device to actual power draw. Different functions are possible [19], but in this paper we use the linear function $U(L_i) = L_i$ as a first estimation for all routers.
- X_i is the carbon footprint of the regional energy production facilities that power the i-th node (grams CO_2 per kWh). In the case of ESnet this corresponds to the values reported in Table II.

We are also interested in the total energy consumption and emission of the whole network at any given moment, *i.e.* the consumption and emission of nodes on active paths together with the consumption and emission of node that are not transmitting data. This is explored in Section V. It is calculated according to Equation 3, where $H = 0$, and Δt is the sampling interval (*hours*):

$$C_{network} = \Delta t \sum_{i \in V} P_i X_i \tag{3}$$

V. SIMULATOR

We have built a discrete event simulator that simulates traffic flows from which we derive the total power consumption and CO_2 emissions of the network. It is written in Python using the SimPy [24] module[3]. The simulator can simulate different networks, but in this paper we only use it to simulate ESnet (see Section III).

It is important to note that we are not interested in accurately modeling the inner workings of networks such as packet-level traffic and specific protocols, but rather how traffic flows influence the load on networking devices, which in turn influences the power consumption of the network, and ultimately its carbon footprint.

[2]Dense Wavelength Division Multiplexer

[3]Simulator source code: https://bitbucket.org/karelvdv/greennet-sim

The network is modeled as a flow network [25] where the flows are traffic, nodes are routers, and edges are network links. We model the capacity of the nodes, which in turn determines the capacity on the links. In our ESnet simulation we use 10G routers, hence the traffic on the links cannot exceed 10G.

The simulation works as follows. Flows are set up to transport a certain amount of data from a randomly chosen source-destination pair. The inter-arrival time of the flows is exponentially distributed. A flow does not change once set up, but disappears when the data transfer is complete. The flow duration is determined by the amount of data to transport and the flow's throughput. These two parameters vary according to a beta distribution, which has a finite range to ensure that flows finish within a reasonable finite simulation time.

A path is selected from the source to destination node according to our model. The chosen path can be either the shortest or greenest or the cheapest. Flows take up a certain amount of bandwidth on every router in the path for the duration of the flow. This is what determines the load on the routers. Devices have a maximum bandwidth capacity, so it is possible that there is no path available from the source to destination. In that case the flow is dropped.

The state of the simulation only changes when a new flow arrives or finishes. At these moments, as well as at regular intervals, the total power consumption, and carbon emissions of the total network are recorded. These metrics are calculated according to Equation 3 and integrated by time passed.

We halt the simulation when a certain amount of data has been pushed through the network.

The simulator takes the following parameters:

- μ^{-1}: The mean inter-arrival time in seconds (exponential distribution).
- α, β: These parameters determine the distribution of flow lengths using the beta distribution.
- d: The percentage of router power consumption that has a dependence on traffic load. If $d = 0$ then routers are power-agnostic, if $d = 1$ they are 100% power proportional. This parameter allows us to investigate what effect energy aware routers have on the network's carbon footprint.
- $\mathit{Path\ selection\ metric}$: shortest, greenest, or cheapest are the criteria by which paths are selected. In each individual simulation run we select all paths according to the same metric.

The simulator takes as input a topology description, regional energy production description, and a configuration where all the simulation parameters are set. The output of the simulation is the total power consumption and carbon footprint of the network, as well as several statistical properties such as time elapsed, number of flows completed, etc.

VI. Experiments

We compared different network traffic scenarios in the ESnet topology with the criterium of carbon footprint. Every scenario exhibits different traffic behavior in terms of inter-arrival time, flow length, and path selection, and ends when 1TByte of data has been pushed through the network. These determine the traffic load on the network over time and thus its carbon footprint. Scenarios are permutations of the following parameters:

- Inter-arrival time: $\mu^{-1} \in \{0.1, 1, 10\}$ seconds.
- Dynamic power: $d \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$.
- Path selection metric: shortest, cheapest, or greenest path.
- Flow type: long, short, or uniform. This simulates users with certain bandwidth usages (see Table III).

TABLE III. Flow types used in our experiments.

Flow	Throughput (Gbit/s)	Data (GBytes)
Long-lived	0.1	[1, 10]
Short-lived	1.0	[1, 10]
Uniform	[0.1, 1.0]	[1, 10]

The amount of data transmitted per flow type is determined by the beta distribution, which is identical for every scenario. We chose the following values for α and β (see Figure 2 for the PDF and CDF):

$$\alpha = 0.5,\ \beta = 1.0$$

For uniform flows we ensure that all of them always have the same duration: $1/0.1 = 10/1 = 10$. The amount of data and throughput are scaled from the beta distribution as such:

$$x \sim B(\alpha, \beta)$$
$$Data = 1 + 9x \text{ GBytes}$$
$$Throughput = 0.1 + 0.9x \text{ Gbit/s}$$

This distribution resembles the user classification according

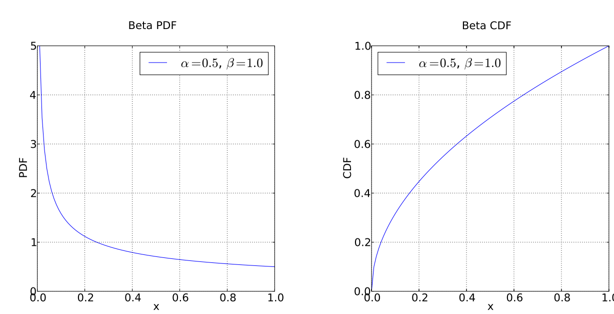

(a) (b)

Fig. 2. PDF and CDF of the beta distribution.

to bandwidth usage and type of connectivity used in [26]. The authors observe that in NRENs there are many users with small bandwidth requirements and a couple with high bandwidth requirements. This is also what can be expected in the ESnet network, where a few scientists and scientific collaborations move large amount data while the majority of communication is short lived.

VII. Results

The carbon footprint of the scenarios we have simulated are shown in Figures 3, 4, and 5. Each graph presents the carbon footprint of the network as a function of router dynamic power

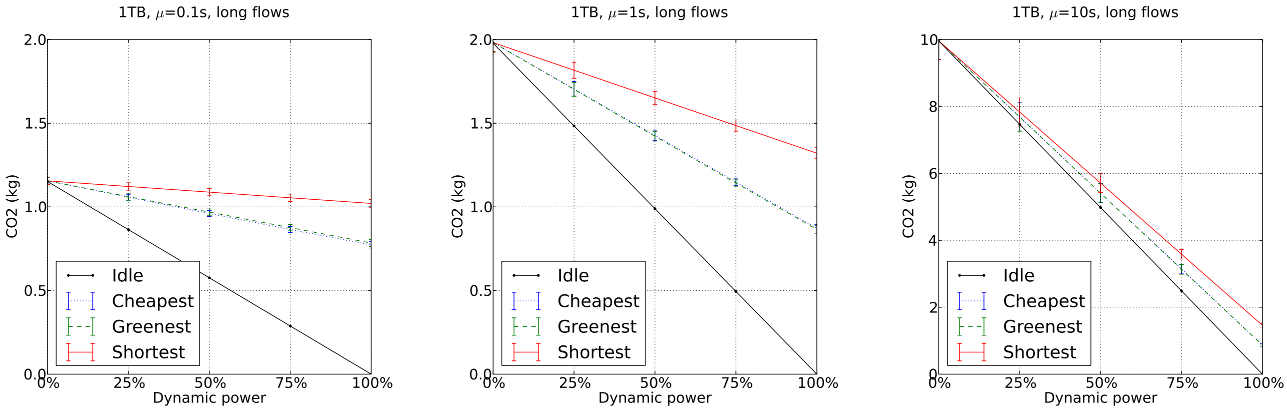

(a) (b) (c)

Fig. 3. Carbon footprint of 1TByte of traffic: Long flows.

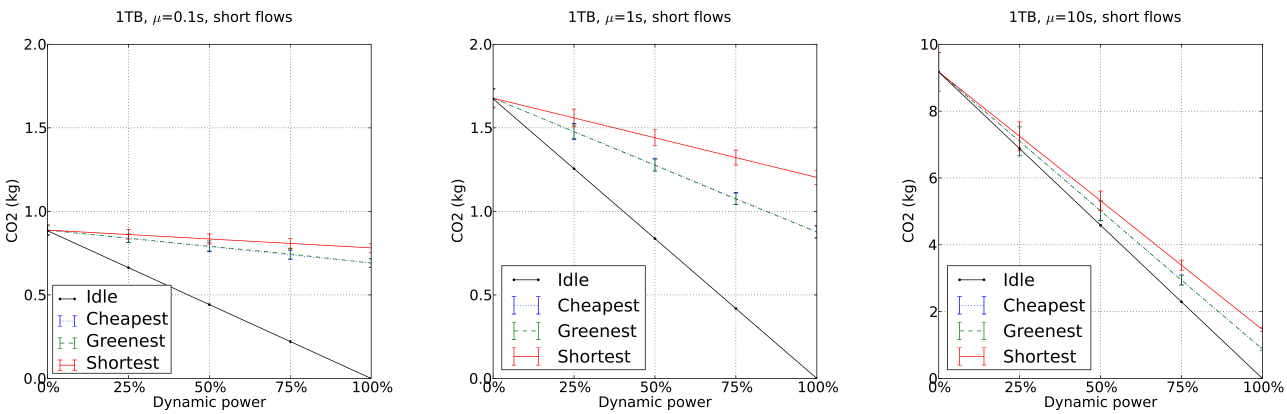

(a) (b) (c)

Fig. 4. Carbon footprint of 1TByte of traffic: Short flows.

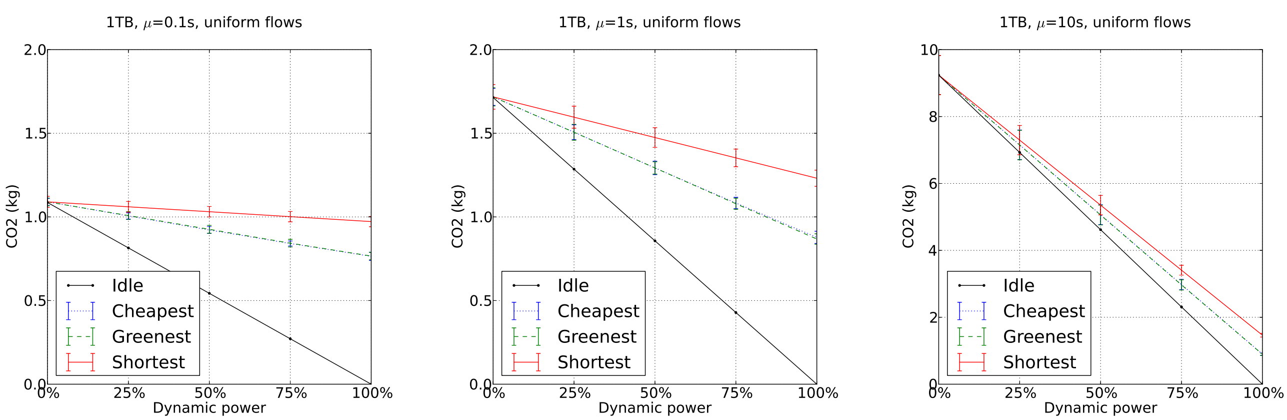

(a) (b) (c)

Fig. 5. Carbon footprint of 1TByte of traffic: Uniform flows.

and path selection metric. The error bars show the standard deviation. The graphs contain also the carbon footprint of the network in the idle state (zero traffic); this represents the lower bound. The upper bound is the carbon footprint at 0% dynamic power. Although 100% dynamic power is an unrealistic scenario it provides useful insight into the actual range of improvement possible. Table IV shows the actual CO_2 reduction compared to 0% dynamic power.

TABLE IV. CARBON FOOTPRINT REDUCTION FACTOR AS THE RATIO OF CO_2 AT X% AND 0% DYNAMIC POWER.

μ^{-1}	Flow	Path Selection Metric	CO_2 reduction per Dynamic power 25%	50%	75%	100%
0.1	Long	Shortest	0.97	0.94	0.91	0.88
0.1	Long	Greenest	0.92	0.84	0.76	0.68
0.1	Short	Shortest	0.97	0.94	0.91	0.88
0.1	Short	Greenest	0.95	0.89	0.84	0.78
0.1	Uniform	Shortest	0.97	0.95	0.92	0.89
0.1	Uniform	Greenest	0.92	0.85	0.77	0.70
1.0	Long	Shortest	0.92	0.83	0.75	0.67
1.0	Long	Greenest	0.86	0.72	0.58	0.44
1.0	Short	Shortest	0.93	0.86	0.79	0.72
1.0	Short	Greenest	0.88	0.76	0.64	0.52
1.0	Uniform	Shortest	0.93	0.86	0.79	0.72
1.0	Uniform	Greenest	0.88	0.75	0.63	0.51
10	Long	Shortest	0.79	0.57	0.36	0.15
10	Long	Greenest	0.77	0.54	0.31	0.09
10	Short	Shortest	0.79	0.58	0.37	0.16
10	Short	Greenest	0.77	0.55	0.32	0.10
10	Uniform	Shortest	0.79	0.58	0.37	0.16
10	Uniform	Greenest	0.77	0.55	0.32	0.10

VIII. DISCUSSION

The first observation we can draw from the simulation results is that the carbon footprint is reduced in every scenario compared to the 0% dynamic power scenario. The carbon footprint decreases linearly with the dynamic power percentage which corresponds to the linear load function we used (see Section IV).

The inter-arrival time has a large impact on the carbon footprint. This makes sense, as a short inter-arrival time means that many flows are active at the same time, resulting in a high average load on the network, which in turn corresponds to a high average power consumption and finally also to a high carbon footprint. In contrast, a long inter-arrival time means that there is only sporadic load on the network and thus the carbon footprint is mostly determined by the idle time. This can clearly be seen by comparing the graphs where the mean inter-arrival times are $\mu^{-1} = 0.1$ versus $\mu^{-1} = 10$ seconds.

The carbon footprint is further reduced by green path selection. As seen in Table IV, the CO_2 reduction factor of green paths ranges from 22% to 91% under ideal circumstances (100% dynamic power) and 5% to 23% under more realistic circumstances (25% dynamic power). Unfortunately most current routers have a dynamic power range of no more than 10% so the CO_2 reduction that can be achieved today is limited. However, the results show that there is a case to be made for more research into energy-aware networking technologies. One such technology, IEEE 802.3az, already exists for Ethernet links and provides a modest reduction in power consumption by switching off idle ports in routers [6].

It is important to note that these results only apply to the subset of the ESnet topology that we have chosen as well as the energy production mixes of the United States (see Table II), and cannot be extrapolated to the general case. For example, as shown in Table II there are only two states in the topology that have a significantly lower carbon footprint than the rest: California and New York. Both are coastal states, located at the borders of the topology. On the other hand, more centrally located states such as Colorado and Kansas have a very high carbon footprint because much of their energy depends on fossil fuels. At the same time they experience much more traffic on average precisely because they are in the core of the network topology. A cursory look at MyESnet [18] confirms this as well. From this we can surmise that it is generally more important for centrally located nodes to have a low carbon footprint than those on the borders, but the actual situation in ESnet is the exact opposite of that.

IX. CONCLUSIONS

We have shown in Section VII that it is possible to reduce the carbon footprint of a network like ESnet under three main conditions:

- That we can access up-to-date information about CO_2 emissions resulting from energy production;
- That we can acquire devices whose power profile depends on the load;
- That we can control the setup of network paths directly.

The first condition is the most difficult to fulfill, as this information is provided entirely at the discretion of energy companies. However some do publish up-to-date figures on energy pricing so it is not outside of the realm of possibility. In this paper we relied on averages but in reality even non-renewable energy sources can fluctuate depending on supply and demand.

NRENs will be dependent on hardware manufacturers to fulfill the second condition. Given the extensive ongoing work in this area we expect that purchase of more energy-savvy devices will be a concrete possibility in the near future.

The third condition is easier to fulfill given that nearly all NRENs have developed circuit provisioning systems for the creation of end-to-end paths. The tight coupling of these provisioning systems with power monitoring information allows to achieve this goal. In the case of ESnet this would concretely mean to couple OSCARS and perfSONAR data to create *green paths*. For other NRENs the availability of power measurements is still an issue.

Furthermore the current network evolution toward programmability by means of Software Defined Networking devices, using for example OpenFlow [27], seems to indicate that more granular control on the route taken by data will become default behavior. Thus studies like ours motivate NRENs to include energy awareness as driving criteria for their SDN developments.

In conclusion, we believe that energy aware devices in combination with path provisioning and power information systems to create *green paths* can be a promising method to reduce the carbon footprint of NRENs. Based on our simulations carried out on the ESnet network we make the conservative estimate that CO_2 reductions of 5% to 23% are attainable depending on the amount of network traffic.

A. Future work

We plan to repeat our simulations with other NRENs such as SURFnet and with more realistic traffic flow patterns. We also plan to investigate the optimality of networks with regards to overall carbon footprint coupled with how that is affected by changing the topology.

X. Acknowledgements

This research was sponsored by NWO (Project GreenClouds), SURFnet (GigaPort), and Cluster Green Software (EFRO), and partially funded by the USA Department of Energy.

References

[1] C. Guok, D. Robertson, M. Thompson, J. Lee, B. Tierney, and W. Johnston, "Intra and interdomain circuit provisioning using the OSCARS reservation system," in *Broadband Communications, Networks and Systems, 2006. BROADNETS 2006. 3rd International Conference on*. IEEE, Oct. 2006, pp. 1–8. [Online]. Available: http://dx.doi.org/10.1109/broadnets.2006.4374316

[2] GÉANT Bandwidth on Demand. http://services.geant.net/bod.

[3] W. Vereecken, W. Van Heddeghem, M. Deruyck, B. Puype, B. Lannoo, W. Joseph, D. Colle, L. Martens, and P. Demeester, "Power consumption in telecommunication networks: overview and reduction strategies," *Communications Magazine, IEEE*, vol. 49, no. 6, pp. 62–69, Jun. 2011. [Online]. Available: http://dx.doi.org/10.1109/mcom.2011.5783986

[4] M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide energy needs for ICT: The rise of power-aware networking," in *Advanced Networks and Telecommunication Systems, 2008. ANTS'08. 2nd International Symposium on*. IEEE, Dec. 2008, pp. 1–3. [Online]. Available: http://dx.doi.org/10.1109/ants.2008.4937762

[5] A. P. Bianzino, C. Chaudet, F. Larroca, D. Rossi, and J. Rougier, "Energy-aware routing: A reality check," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*. IEEE, Dec. 2010, pp. 1422–1427. [Online]. Available: http://dx.doi.org/10.1109/glocomw.2010.5700172

[6] D. Pavlov, J. Soeurt, P. Grosso, Z. Zhao, K. van der Veldt, H. Zhu, and C. De Laat, "Towards energy efficient data intensive computing using IEEE 802.3az," *High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion*, pp. 806–810, Nov. 2012. [Online]. Available: http://dx.doi.org/10.1109/sc.companion.2012.112

[7] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*. IEEE, Apr. 2008, pp. 457–465. [Online]. Available: http://dx.doi.org/10.1109/infocom.2008.93

[8] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. H. Andrew, "Greening geographical load balancing," in *Proceedings of the ACM SIGMETRICS Joint International Conference on Measurement and Modeling of Computer Systems*, ser. SIGMETRICS '11. New York, NY, USA: ACM, 2011, pp. 233–244. [Online]. Available: http://dx.doi.org/10.1145/1993744.1993767

[9] Y. Shang, D. Li, and M. Xu, "Energy-aware routing in data center network," in *Proceedings of the First ACM SIGCOMM Workshop on Green Networking*, ser. Green Networking '10. New York, NY, USA: ACM, 2010, pp. 1–8. [Online]. Available: http://dx.doi.org/10.1145/1851290.1851292

[10] P. Mahadevan, S. Banerjee, P. Sharma, A. Shah, and P. Ranganathan, "On energy efficiency for enterprise and data center networks," *Communications Magazine, IEEE*, vol. 49, no. 8, pp. 94–100, Aug. 2011. [Online]. Available: http://dx.doi.org/10.1109/mcom.2011.5978421

[11] S. Bird, A. Achuthan, O. A. Maatallah, W. Hu, K. Janoyan, A. Kwasinski, J. Matthews, D. Mayhew, J. Owen, and P. Marzocca, "Distributed (green) data centers: A new concept for energy, computing, and telecommunications," *Energy for Sustainable Development*, no. 0, 2014. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0973082613001129

[12] M. X. Makkes, A. Taal, A. Osseyran, and P. Grosso, "A decision framework for placement of applications in clouds that minimizes their carbon footprint," *Journal of Cloud Computing*, vol. 2, no. 1, pp. 1–13, 2013. [Online]. Available: http://dx.doi.org/10.1186/2192-113x-2-21

[13] C. Despins, F. Labeau, T. Le Ngoc, R. Labelle, M. Cheriet, C. Thibeault, F. Gagnon, A. Leon-Garcia, O. Cherkaoui, B. St. Arnaud, J. Mcneill, Y. Lemieux, and M. Lemay, "Leveraging green communications for carbon emission reductions: Techniques, testbeds, and emerging carbon footprint standards," *Communications Magazine, IEEE*, vol. 49, no. 8, pp. 101–109, Aug. 2011. [Online]. Available: http://dx.doi.org/10.1109/mcom.2011.5978422

[14] GÉANT and GreenStar Network: Powering Networks Through Renewable Energy. http://www.geant.net/Media_Centre/Media_Library/Media%20Library/GEANT_and_the_GreenStar_Network.pdf.

[15] B. Aksanli, T. S. Rosing, and I. Monga, "Benefits of green energy and proportionality in high speed wide area networks connecting data centers," pp. 175–180, Mar. 2012. [Online]. Available: http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6176458

[16] About ESnet. http://es.net/about.

[17] B. Tierney, J. Metzger, J. Boote, E. Boyd, A. Brown, R. Carlson, M. Zekauskas, J. Zurawski, M. Swany, and M. Grigoriev, "perfSONAR: Instantiating a global network measurement framework." [Online]. Available: http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.157.9027

[18] MyESnet. https://my.es.net.

[19] S. Ricciardi, D. Careglio, F. Palmieri, U. Fiore, G. Santos-Boada, and J. Solé-Pareta, "Energy-Oriented Models for WDM Networks," vol. 66, pp. 534–548, 2012. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-30376-0_39

[20] USA Economic and Energy data, Institute for Energy Research. http://www.instituteforenergyresearch.org/states.

[21] "CO2 emissions from fuel combustion," *International Energy Agency*, Jul. 2011.

[22] B. K. Sovacool, "Valuing the greenhouse gas emissions from nuclear power: A critical survey," *Energy Policy*, vol. 36, no. 8, pp. 2950–2963, Aug. 2008. [Online]. Available: http://dx.doi.org/10.1016/j.enpol.2008.04.017

[23] R. Sedgewick, *Algorithms in C, Part 5: Graph Algorithms*, Aug. 2001.

[24] SimPy. https://simpy.readthedocs.org.

[25] R. Ahuja, T. Magnanti, and J. Orlin, *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.

[26] C. de Laat, E. Radius, and S. Wallace, "The rationale of the current optical networking initiatives," *Future Generation Computer Systems*, vol. 19, no. 6, pp. 999–1008, Aug. 2003. [Online]. Available: http://dx.doi.org/10.1016/s0167-739x(03)00077-3

[27] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: http://dx.doi.org/10.1145/1355734.1355746