Renewable Energy-Aware Manycast Overlays

Thilo Schöndienst[*], Dylan A.P. Davis[†], Jeremy M. Plante[*], and Vinod M. Vokkarane[*]

Vinod_Vokkarane@uml.edu

[*]Department of Electrical & Computer Engineering, University of Massachusetts Lowell
[†]Computer and Information Science Department, University of Massachusetts Dartmouth

Abstract—**Manycasting allows a single source to reach multiple destinations while providing flexibility in destination selection. Our goal in this work is to improve the cost of the Manycast Drop at Member Node (MA-DMN) overlay algorithm in terms of energy consumption and associated greenhouse gas (GHG) emissions. To reduce the environmental impact, ideally a large percentage of the network nodes along the transmission and the chosen destinations need to be green. We present a novel energy-conservative emission-aware variant of the MA-DMN algorithm. We then propose further modifications to increase the utilization of those destinations that are powered by renewable energy sources: Manycast Drop at Greenest Nodes (MA-DGN). The potential for emission reduction by those algorithms is two-fold: the data is transported in the most efficient way, and processed at the greenest available data centers. We compare the approaches by simulating realistic quantities of dynamic traffic. We assume heterogeneously distributed and time-dependent availability of renewable energy sources to power nodes throughout the network. We find that the energy source-aware algorithms lower both energy-consumption and GHG emissions at stable network performance levels, in some cases even lowers blocking rate.**

Index Terms—**Manycasting, Greenhouse Gas, Emissions, WDM, Overlay, Split-Incapable.**

I. INTRODUCTION

Large-scale scientific domains are becoming increasingly collaborative. Modern big science cannot succeed without cooperation and communication among several otherwise independent laboratories, recent trends call for localized experimentation and distributed result analysis. The availability of renewable energy at certain data processing locations further encourages the prioritization of using some sites over others. As data yield frequents the range of petabytes or exabytes of information, sophisticated communication paradigms like multicast and manycast become necessary. The advantages of supporting point-to-multipoint communication paradigms in the optical layer have been discussed in [1], [2].

Multicast communication establishes connections from a single source, i.e. experiment laboratory, to multiple destinations, i.e. super-computing sites, data analysis laboratories, or super-storage facilities. If one such destination cannot be used, or is unreachable given the current network or resource state (too much competing traffic, prior computational commitments), the overall distribution task is unsuccessful. Manycast [3] by contrast allows for connections to be established from a single source to a non-deterministic subset of candidate destinations. Unlike multicast communication, manycasting decrees that *not all* proposed destination sites must be reached for successful provisioning [4]. Under the manycast paradigm,

each destination provides its own unique service, or one fraction of a parallel service distributed across each of the chosen destination sites. The problem of optimally identifying the minimum cost Steiner tree for a manycast request has been shown to be NP-hard in [5].

Since not all destinations must be connected to the source for an individual manycast session, the destination selection becomes critical to the cost and success rate of a set of service requests. Candidate destinations may be prioritized based on cost or end-to-end delay. Nodes with the fewest prior commitments or greatest processing power may be incentivized. If destination resources are homogeneous, the subset of selected nodes may merely specify that some number of candidates must be reached, possibly irrespective of their costs/locations. In this study, we reduce network-wide greenhouse gas emissions by intensified use of nodes with localized high availability of renewable energy sources, which will minimize the carbon output of a manycast session.

Despite the growing necessity for point-to-multipoint connectivity in large-scale networks, not all backbone architectures support it. Networks such as the DOE's Energy Sciences network (ESnet) contain only OXCs which do not support splitting at the optical (physical) layer. In this type of *split-incapable* network, manycast communication can *only* be supported as a logical overlay to the unicast-only network capability. Prior works [6], [7] visit the manycast provisioning problem in split-incapable networks through consideration of both optimal (ILP) and approximate (heuristic) solutions on WDM networks under static and dynamic traffic scenarios. The manycast overlay approaches therein are proposed as energy-efficient alternatives to comparable multicast overlay approaches. In this work, we present a quantitative analysis of the superior overlay approach with regard to both energy consumption and carbon emission yield under dynamic traffic scenarios on real-world core network topologies. Additionally, we propose emission-efficient variations on the selected overlay approach and provide both qualitative and quantitave comparisons.

In [8] the authors present a broad overview of early efforts in saving energy in network operations. Energy-efficient routing schemes, network planning, and selectively turning equipment on and off are mentioned as strategies to reduce energy consumption.

A comprehensive power consumption model is proposed in [9]. Load dependent nodal power consumption is calculated by adding up the power consumption values of the node's layers.

Recent work [10] discusses a realistically constrained en-

ergy optimization. The main findings include that the main power consumers in an IP-over-WDM network are routers and line cards. Anticipating the emergence of equipment with more flexible capabilities the authors suggest in [11] that individual components of core networks, such as line cards, can be switched off to save energy in low-demand scenarios. The authors of [12] provide distributed solutions to turn off parts of the network during off-peak hours. The smaller topology can still support the lower demands while overall power consumption is significantly reduced. Another study on the possible reductions by putting equipment to sleep or turning it off entirely is presented by the authors of [13], [14]. These efforts and more research within the core network community culminate in larger government funded research, for example the TREND project [15].

In [16] the authors propose an approach to reduce power consumption by concentrating processing within the datacenter to as few racks as possible to turn off the switches sitting on top of those racks. Effectively the powered down switches remove these from the network topology and the reduced network size leads to reduced network power consumption.

One strategy to make the network *green*, or lower the produced greenhouse gas (GHG) emissions caused by the network, is to co-locate data centers and Grid resources close to sites with high availability of renewable energy resources, i.e., wind, solar, geothermal. The trade-off between rise in transport emissions and reduction through this green processing is analyzed in [17]. The authors propose anycast selection of processing data-centers according to which data center can be reached in the most energy-efficient way.[1] **An optimal solution for energy optimized selection of destinations in optical grids is presented in [18]. The chosen 2-step approach outperforms other available algorithms for energy-efficient anycast routing.**

Recently research interest has risen for the consideration of renewable energy sources in routing and network planning. These efforts marry the advantages of green data *processing* and green data *transport* [19]. Green power source aware routing with GMPLS is discussed in [20]. A multi-layer approach is presented in [21]. For each network location, the authors model partial renewable-energy supply and also partial fossil-energy supply. In [22] the authors study the combination of energy proportional routers with partial supply of green energy. The authors of [23] present one of the first approaches that makes use of dual power sources (clean renewable vs. fossil fuels) for routing and wavelength assignment. Unicast routing is optimized in terms of emissions yield on this partially green network. Work on dual power source-aware algorithms for network survivability is presented in [24].

Initial results for the least-impact tree manycast approach presented in this paper have been published in [25]**, the first paper to use manycast to leverage data network *and* power grid cross-optimization**. This paper extends the depth of

those results by studying further metrics, like the number of green destinations selected, and contains both theoretical and more realistic emission scenarios with dynamically and continuously changing emission factors. We further introduce a **rewritten algorithm that increases the scope of the approach: even before the tree selection step, we optimize the tree construction across power and network layers and add the objective of reaching a large number of green endpoints (=data centers).**

The rest of the paper is organized as follows: In Section II we summarize the proposed manycast overlay heuristics and compare them qualitatively. Section III introduces the node model used in evaluating the overlay approaches and provides formulas for evaluating the power consumption, energy consumption, and emissions of network nodes based on the proposed model. In Section IV, realistic manycast request sets are provisioned using the proposed overlay approaches and evaluated under dynamic traffic scenarios in networks with dual power sources. Section V expands the scope of the power sources in the network and presents two use-cases wherein green nodes may have variable impact on GHG emissions. We analyze and interpret the performances of the proposed algorithms on these variable emission scenarios in Section V. Finally we conclude the paper in Section VI.

II. OVERLAY APPROACHES

In this section we recapitulate the manycast overlay models for split-incapable networks as proposed in previous studies [6], [7]. Further, we qualitatively summarize the relevant comparisons between the proposed models. We also propose two novel, emission-aware modifications on the preferred algorithm, which aim to minimize the overall GHG emissions yield for a set of manycast requests.

Given a network $G = (V, E)$, a manycast request can be defined as $R = (s_r, D_r, K')$ where s_r is the source node of the request ($s_r \in V$) and D_r is the set of candidate destination nodes ($D_r \subseteq V - \{s_r\}$). For a manycast request with K destination members, we represent the destination set as $D_r = \{d_1, d_2, \ldots, d_K\}$. $K' \leq K$ is the number of nodes necessary to reach in order to service the manycast request. Note that if $K' = K$, the request is actually a multicast one. Collectively, we will refer to s_r and D_r as the *members* of a manycast request. To service a given manycast request we must find a set of lightpaths (combination of a route and a wavelength) that starts at the source node of the request and reaches at least K' nodes from the candidate destination set in the overlay network. We assume that the wavelength continuity constraint is enforced for all overlay models. This constraint enforces that a lightpath uses the same wavelength on all the physical links of the network it traverses.

A. Multiple-Hop Overlay Model

The typical approach to logical manycasting as an overlay service is to construct a logical tree for each request such that each branch in the tree is a unique lightpath originating at the source node. Such an approach is presented in [6] as a baseline for comparison. The authors of [26] predict that the inevitable

---

[1]Anycasting is a specialized case of manycasting that must reach exactly one candidate destination.

growth and complexity of next generation applications will render a such naive, single-hop manycast overlay mechanism inefficient and infeasible due to high bandwidth consumption. An alternative multiple-hop logical overlay called *Manycasting with Drop at Member Node (MA-DMN)*, was also proposed in [6]. MA-DMN establishes a set of lightpath routes to connect the source node to the selected K' destinations from D_r, such that not all lightpaths must be directly sourced at s_r. The MA-DMN approach allows lightpaths to be terminated (dropped to the electronic layer) or originate exclusively at member nodes of the manycast request, thereby establishing manycast trees in the overlay layer consisting of (possibly) multiple hops from the source. Each drop-member node requires an optical-electrical-optical (O-E-O) conversion in order to "split" the lightpath from a single input port to multiple output ports, which results in longer end-to-end delay per request than the single-hop model. Previous study has shown that allowing lightpaths to drop at any arbitrary node in the network results in only minimal resource consumption reduction but requires greater complexity than MA-DMN.[2]

Definition (MA-DMN): The overlay solution must assign, for each request, a set of lightpath-routes to reach at least K' destination nodes in such a way that the number of wavelengths required is minimized while satisfying the wave-length continuity constraint. The solution must also take into consideration that for any given manycast request, no lightpath terminates/originates at a non-member node of the request.

B. Emission-Aware Overlay Variant

The most beneficial aspect of manycast communication is its ability to avoid inefficient or undesirable nodes during lightpath establishment. MA-DMN allows for reduced resource consumption as compared to the single-hop overlay model, but does not consider the GHG emissions resulting from the trees it constructs to satisfy a request. We therefore propose a variation to MA-DMN which is aware of which candidate destinations are powered by renewable energy sources. This *Manycasting with Drop at Greenest Node (MA-DGN)* overlay approach will evaluate each member node in the candidate set and incentivize the utilization of those nodes which are greenest. In Section III we describe how the source, drop-member, and destination nodes contribute to most of the network's GHG emissions. By maximizing the selection of green nodes to fill these premium roles for as many requests as possible, MA-DGN seeks to reduce the overall network-wide emissions caused by manycast requests. Note though that there is no guarantee that a green node will be available for serving as a drop-point. In such a case, MA-DGN simplifies to MA-DMN. A special consideration here is that most likely the green network nodes also indicate the availability of green data processing sites. If this is the case the benefits are even greater, as data centers are and are forecasted to keep contributing twice as much emission to the global footprint of ICT as wired networks are [27].

Definition (MA-DGN): The overlay solution must assign, for each request, a set of lightpath-routes to reach at least K' destination nodes in such a way that the greenest nodes are selected from among the candidate set. If multiple nodes possess equal levels of GHG emission potential, then the number of wavelengths required is minimized while satisfying the wavelength continuity constraint. The solution must also take into consideration that for any given manycast request, no lightpath terminates/originates at a non-member node of the request.

C. Emission-Aware Heuristics

The heuristic used to satisfy multiple-hop manycast overlay-trees using MA-DMN is reviewed in detail in [6]. For a manycast request with K candidate destinations and K' required destinations to satisfy the request, the referenced algorithm establishes K-alternate trees, wherein the first hop of the k^{th} tree is mandatorily the shortest path to the k^{th} candidate destination in the set. Of these K-alternate trees, the one which consumes the fewest additional resources (lightpaths) given the network's current state is selected as the provisioning solution to the manycast request. In short, the tree which consumes the fewest additional wavelengths is selected. A tie between multiple trees is broken by selecting the tree with the lowest total logical hop count, thus reducing delay. The worst-case runtime complexity for the MA-DMN heuristic is $O(K|V|^3)$ [6].

As stated previously, the MA-DMN heuristic is blind to energy consumption and the emissions production resulting from its manycast overlay-tree solutions. We therefore propose a novel emission-aware variant of the MA-DMN heuristic named *Manycasting with Drop at Member Node using Least-Impact Trees (MA-DMN-LIT)*. Like MA-DMN, this approach still produces K-alternate trees for each manycast request, but rather than greedily selecting the tree which adds the fewest wavelengths to the network, MA-DMN-LIT will inspect the emissions potentially produced by establishing each tree and provision the one with the least immediate impact in this regard. If multiple trees should produce the same emissions, the approach simplifies down to the MA-DMN algorithm and breaks the tie based on resource consumption.

An algorithm similar to that for MA-DMN can be used to produce trees for our novel MA-DGN approach. We will still create K-alternate trees for each manycast request, but will sort the list of paths between member nodes, which serves as the input to the algorithm, so that the destinations reached are guaranteed to be the greenest available. This will result in longer physical paths, but will reduce emissions by performing costly O-E-O conversions at nodes powered by renewable energy sources. The Least-Impact Tree extension can also be applied to MA-DGN to produce a variant which is aware not only of the destination emissions, but the overall tree emissions as well. We refer to this heuristic as *Manycasting with Drop at Greenest Node using Least-Impact Trees (MA-DGN-LIT)*.[3]

---

[2]We have also compared MA-DMN to the single-hop overlay approach in terms of energy- and emission-efficiency, but omit those results because as expected, emissions and energy consumption are much lower for the multiple-hop approach.

[3]Our results show that using the LIT approach is beneficial for overall network performance irrespective of the availability of green nodes. Therefore we do not explicitly present the findings for DGN without LIT.

Algorithm 1: MA-DGN using Least-Impact Trees

> **input** : Manycast Request: $R = (s_r, D_r, K')$,
> $D_r = \{d_1, d_2, \ldots, d_K\}$
>   : List of appropriate least-cost routes: $routeList$
> **output** : Manycast overlay-tree producing the lowest GHG emissions

```
 1  AltTrees[K] = NULL
 2  routedDestinations = 0
 3  sort(D_r) from lowest to highest ε
 4  sort(routeList, D_r)
 5  for each d_i ∈ D_r do
 6      Tree_k = NULL
 7      Tree_k.add(SP(s, d_i))
 8      update routedDestinations
 9      for each route_i ∈ routeList do
10          if routedDestinations ≥ K' then
11              break
12          if (route_i.source ∈ Tree_k) & (route_i.dst ∉
                Tree_k) & (route_i.dst ∈ D_r) then
13              Tree_k.add(route_i)
14              update routedDestinations
15      calculateEmissions(Tree_k)
16      AltTrees.add(Tree_k)
17  sortByEmissions(AltTrees)
18  return min(AltTrees)
```

Algorithm 1 describes the MA-DGN-LIT tree creation and selection process. This algorithm is a modified version of the MA-DMN-LIT algorithm found in [25]. Just as MA-DMN does, the algorithm generates K candidate trees, with the constraint that the first hop in the k^{th} tree is the least-cost path from the source node to the k^{th} candidate destination of the manycast request. Before the algorithm begins building trees however, we sort the input, $routeList$ which is a list of all weighted paths in the network between the source and each candidate destination member of the request, as well as all the weighted paths between pairs of member nodes. On Line 3, we sort the destination set, D_r so that the first destinations considered will be those with the smallest emission-factor, a value representing the quantity of CO_2 released per kWh of electricity consumed. Higher values of ϵ indicate a larger environmental impact; green nodes will have very small ϵ values. This factor is described in greater detail in Section III. We can now use the newly sorted D_r to sort $routeList$ on Line 4. The result is a list of least-cost paths sorted first by destination (to match the new ordering of D_r), and then by path weight. Now, despite the fact that the first hop in each tree will be to a different destination, each additional hop will be provisioned as efficiently as possible in terms of destination-driven emission production. The complexity for sorting the $routeList$ on Line 4 has a worst-case runtime of $O(2K^2)$.

The other differences from MA-DMN are found at lines 15 and 17 of Algorithm 1; these are the LIT modifications. Once the candidate trees have been constructed, their GHG emissions are evaluated, using Eqs. $(1)-(8)$ described later in Section III-B. The worst-case runtime complexity for this operation is the amount of time it takes to sum up the emissions at each node in the tree, $O(|V|)$. Once all the candidate trees have

been constructed and evaluated, they are then sorted according to their total GHG emissions, an operation which runs with a complexity of $O(K)$. Finally, the lowest-emissions tree is selected and provisioned on the network to satisfy the manycast request. If no tree was found that can satisfy the request, the request is considered blocked.

The overall complexity of the DGN-LIT algorithm can be represented by the sorting of $routeList$ based on candidate ordering plus the product of the number of routes traversed in $routeList$, $O(|V|^2)$, and the maximum number of nodes to traverse in a single tree, $O(|V|)$ for each of the K trees: $O(2K^2 + K|V|^3)$. Since D_r is a subset of V, and $|D_r| = K$, this complexity may be reduced to simply: $O(K|V|^3)$. Note that by omitting lines 3 and 4, Algorithm 1 describes MA-DMN-LIT. Removing Lines 15 and 17 simplifies it further to MA-DMN.

D. Physical Lightpath Routing

Algorithm 1 expects and manipulates a list of paths in the physical topology between pairs of nodes. We have so far only discussed overlay routing, and simply assumed least-cost path generation and wavelength availability. We have investigated two approaches to physical routing in the overlay-trees. The first approach is a shortest path routing scheme using Dijkstra's algorithm. Therefore, $routeList$ is a list of shortest paths between node-pairs, we refer to this as Shortest Path (SP) routing in the results in Section IV. The second approach aims to consider not just the physical length of the path, but the greenness as well. We weight links between nodes depending on whether they connect two black nodes (carbon-fueled), two green (renewable energy-fueled) nodes, or a black-green combination. The links between a pair of green nodes is least costly, and the links connecting a pair of black nodes are most costly. We then apply Dijkstra's algorithm once again to find the greenest paths. This results in a $routeList$ containing paths with as few fossil-fueled nodes as possible, we refer to this as Greenest Path (GP) routing in the performance evaluation section. Once these routes are passed to Algorithm 1 and added to the overlay-trees, wavelength assignment is performed in a first-fit manner.

E. Illustrative Example and Comparison of Heuristics

Figure 1 shows a simple five-node setup and the overlay-trees returned by the different algorithms for the request $R : \{S, \{D1, D2, D3, D4\}, 2\}$. Nodes S, D1, and D2 are assumed to be powered by fossil energy sources. Nodes N, D3, and D4 are powered by renewable energy sources. Node N is a regular node in the network, not part of the destination set; this node can only be bypassed and is unable to serve as a drop-member node. All algorithms start out with a precalculated set of least-cost paths for all node-pairs, based on minimum hop distance or another weight, such as path greenness. The algorithms will loop through K iterations, and during each iteration will construct a tree reaching the required number of destinations. In each subfigure, all of the alternate trees that could be used to satisfy the manycast request are shown
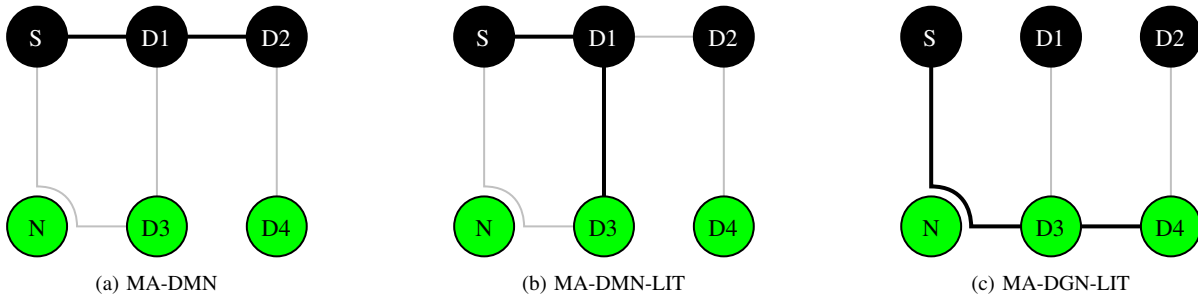
Fig. 1 – Logical overlay-trees selected by the different algorithms to satisfy a request $R : \{S, \{D1, D2, D3, D4\}, 2\}$.

together. The thick lines show the selected tree, and the thin gray lines are the potential paths of the unselected trees.

Figure 1(a) shows the output of the MA-DMN algorithm. Following Algorithm 1 (but excluding lines 3, 4, 15 and 17), the first destination to reach during the first loop iteration must be D1. The next-nearest destination is D2, which is added to the tree. At this point the required number of destinations is reached, so the iteration is complete and one candidate tree has been built to satisfy the request. For the second tree, D2 must be reached first. D1 is already on the shortest path to D2, and thus two destinations have been reached, completing the iteration. The algorithm will continue this procedure for the last two iterations, which reach D3 and D4 first, respectively. Once all the trees are built, the algorithm finally chooses the shortest tree solution, which is the one found during both of the first two iterations/

In Fig. 1(b), MA-DMN-LIT is shown, which follows the same process of constructing trees used by MA-DMN. The same set of trees will be considered, except the algorithm will now select the most environmentally friendly tree. The third iteration reached D3 first, and thus includes a green destination node, unlike the tree chosen for MA-DMN.

The MA-DGN-LIT algorithm strives to maximize the number of green destinations in the tree. This outcome is illustrated in Fig. 1(c). Before the candidate tree search, we sort the destination candidates by their greenness. Thus we encourage finding trees as shown in Fig. 1(c): both destinations are green, node N is simply bypassed optically. This also shows two of the potential costs of picking low emission trees: the paths are longer (increasing delay), and may consume more energy. The additional energy however, is consumed at a green node and thus produces fewer emissions.

III. NODAL POWER CONSUMPTION MODEL

Here we describe the node model used in evaluating the overlay heuristics described in the previous section. We also present values and formulas for calculating power consumption, energy consumption, and emissions generated by each node in the network. In [9] the authors propose a comprehensive model for multilayer network power consumption which we use as foundation and combine with [28] to calculate GHG emissions. Fig. 2(a) presents an illustration of a generalized node, which can be specialized as an OXC, ROADM, etc. This model can be used to identify the power consumed by any heterogeneous node in the network based on its classification

TABLE I
NODE-SPECIFIC VARIABLES

Symbol	Description	Unit
A_{IN}	Input traffic from access network	$Gbps$
A_{OUT}	Output traffic to access network	$Gbps$
L_{D}	Lightpaths dropped from WDM to IP (O-E)	integer
L_{A}	Lightpaths added to WDM from IP (E-O)	integer
L_{IN}	Incoming lightpaths	integer
L_{OUT}	Outgoing lightpaths	integer
ϵ_{n}	Node-specific emission-factor	$kgCO_2e/kWh$

TABLE II
NETWORK-SPECIFIC VARIABLES

Symbol	Description	Value	Unit
r	Transponder transmission rate	10	$Gbps$
$\pi_{\text{TX}}(r)$	Transponder's O-E-O power consumption	0.05	kW
π_{OXC}	OXC power consumption on optical plane	0.1	kW
π_{IP}	IP layer power consumption	0.01	$kW/Gbps$

in terms of its role in satisfying a manycast overlay request, as well as its incoming and outgoing nodal degrees. The model illustrates not only the optical layer, but also the IP layer and the ingress/egress points to the access network. The component in Fig. 2(a) labeled "Transponder" represents the transmitter/receiver component of the switch on the optical plane. Table I describes the vector labels in Fig. 2(a), and provides their units of measurement.

The table also includes an emission-factor value, ϵ_n that represents the impact of power consumption on GHG emissions. The units for the emission-factor are kilograms of carbon dioxide equivalent per kiloWatt-hour (kgCO$_2$e/kWh). The emission-factor varies with location and time, $\epsilon = \epsilon(n,t)$ with n = node location and t = time. The location variation is dictated by environmental factors, such as feasibility and implementation of certain kinds of power generation. Examples are the required presence of rivers for hydropower, geothermal activity in sufficient shallow rock, etc. The variation in time depends on physical and environmental factors, e.g., solar panels can not operate at night and only have reduced output on cloudy days. **The emission factor not only represents GHG emissions in connection with electricity generation but comprises the entire life cycle emissions for the respective energy sources. Included in the emission factor are the environmental impact of upstream, fuel cycle, operation,**

and downstream factors [29].

Values for the life cycle emissions vary greatly in literature but harmonization efforts[4] yield median values from 0.007 (Hydropower) to 0.054 (Photovoltaic generation) for the renewables. For nonrenewable nuclear energy the median value is 0.016. The fossil sources are found to have medians from 0.477 (Natural gas) to 1.001 (Coal). The emission factor is expressed in kgCO$_2$e for simplicity: The *e* stands for equivalent; thus the impact of other pollutants and aerosols is included in the single factor ϵ and does not need to be considered separately [28]. The emissions embodied in network equipment is not considered as this is a traffic optimization paper rather than system design, that is we consider no monetary cost or GHG emissions for equipment purchases; we consider only operational expenditures (OPEX) and not capital expenditures (CAPEX). Table II contains additional node-specific values. These constants are uniform across every node in the network while nodes are switched on.

A. Manycast Tree Nodal Classification

The given node model is uniform across all nodes in the network, however the power and energy consumed by each node is unique, depending on its role in provisioning a manycast overlay-tree. Table III depicts the mutually exclusive types that each network node may be classified as for an individual manycast request by MA-DMN and MA-DMN-LIT. Nodes which are not destination members, but may be found along paths within the tree are intermediate nodes. The switches at these nodes do not require any electronic conversion, thus signals may bypass them all-optically, incurring only very minor power and energy costs. Source nodes and destination nodes are responsible for performing E-O and O-E conversions respectively, which consume far more power than an intermediate bypass node. Destination nodes that not only terminate a lightpath but also serve as the source for a connecting lightpath are defined as *drop-members* of the tree, and incur additional energy costs for splitting the signal electronically to be forwarded along separate lightpaths to one or more additional destination nodes. Unused nodes are those which are not included in the tree solution for provisioning a manycast request and therefore do not consume any power or energy for the given request. If a node is unused across all currently provisioned manycast overlay-trees, the node is non-essential to the network and may be switched off to conserve energy and prevent idling costs. All network nodes are idle before any requests arrive and again after all requests have departed (i.e. lightpaths freed) in dynamic traffic scenarios.

B. GHG Emissions Model

Equations. $(1) - (4)$ provide relationships between the values in Tables I and II and their respective effects on node-specific power consumption: : $p_{\text{IP}}, p_{\text{O-E-O}}$ and p_{WDM} are the instantaneous power values at the IP, OEO-Conversion and WDM layers respectively. For simplicity, Eqs. $(1)-(3)$ assume

[4]http://en.openei.org/apps/LCA/
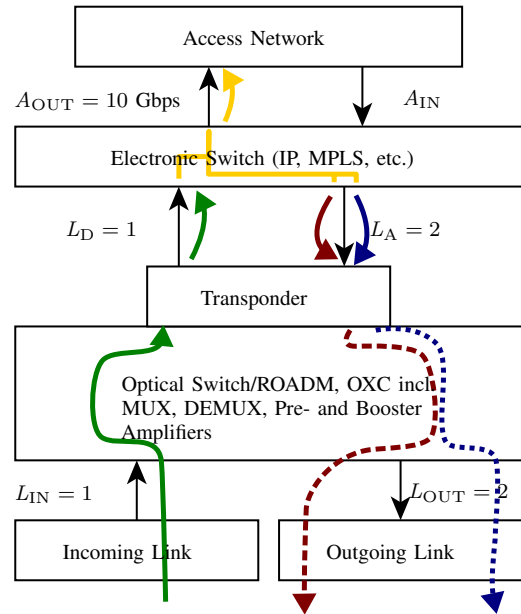
(a) Multilayer node model used in evaluating node-specific power consumption. Illustrative paths through the node are added to describe the behavior of a drop-member.
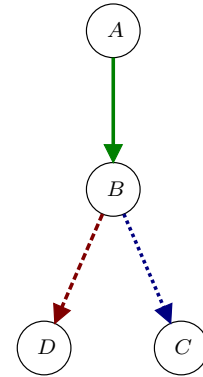
(b) Typical MA-DMN tree; B is classified as a drop-member in this overlay-tree.

Fig. 2 – Proposed node model and example overlay-tree.

that all operands are in terms of an instant of time, i.e., p_{IP} is more accurately described as $p_{\text{IP}}(n, t)$, where n is the node in question, and t is the instant in time of this calculation. In Eq. (3), $\alpha = 0.085$ kW is the power consumed by amplifiers, which are necessary even for the all-optical bypassing of a lightpath through a node. The value $\beta = 0.15$ kW is the constant power consumed for directing the lightpaths to the appropriate ports, for example using microelectromechanical systems (MEMS). Equation (4) describes the total power consumption for a node at an instant of time as the sum of the power consumed at the optical, IP, and access layers. ϕ_{n} and Φ calculate the node-specific and network-wide energy consumption respectively over the total time interval of request arrivals and departures. In Eqs. (7) and (8), γ_n and Γ the node-specific and network-wide GHG emissions respectively. We use these formulas as the basis by which to obtain the results

TABLE III
OVERLAY NODE CLASSIFICATION

Class	Indegree	Outdegree	Qualifying Parameters
Source	0	$[1, K']$	-
Destination	1	0	Request Member
Intermediate	$[1, K']$	$[1, K']$	Request Non-Member; In = Out
Drop-Member	1	$[1, K'-1]$	Request Member
Unused	0	0	Not in Tree

Fig. 3 – 14-node augmented ESnet topology.

presented in Sections IV and V and to compare the relative energy- and emission-efficiency of MA-DMN, MA-DMN-LIT and DGN-LIT.

$$p_{\text{IP}} = \pi_{\text{IP}}(A_{\text{IN}} + A_{\text{OUT}} + r(L_{\text{D}} + L_{\text{A}})). \tag{1}$$

$$p_{\text{OEO}} = \pi_{TX}(r)(L_{\text{D}} + L_{\text{A}}). \tag{2}$$

$$p_{\text{WDM}} = \pi_{OXC}(L_{\text{D}} + L_{\text{A}}) + \alpha(L_{\text{IN}} + L_{\text{OUT}}) + \beta. \tag{3}$$

$$p(n, t) = p_{\text{IP}} + p_{\text{OEO}} + p_{\text{WDM}}. \tag{4}$$

$$\phi_{\text{n}} = \int_{t_0}^{T} p(n, t)\, dt. \tag{5} \qquad \Phi = \sum_{\text{n}}^{N} \phi_{\text{n}}. \tag{6}$$

$$\gamma_{\text{n}} = \int_{t_0}^{T} p(n, t) \cdot \epsilon_{\text{n}}\, dt. \tag{7} \qquad \Gamma = \sum_{\text{n}}^{N} \gamma_{\text{n}}. \tag{8}$$

Consider the overlay-tree shown in Fig. 2(b), which may be a solution provided by MA-DMN for the manycast request, $R = (A, \{B, C, D\}, 3)$. According to the classifications presented in Table III, node B may be classified as a drop-member. Fig. 2(a) depicts the internal state of node B according to the described model. The number of lightpaths into the node, L_{IN}, is equal to 1, and that is the lightpath directed from node A. Since this node is a drop-member, the incoming signal must be considered at the IP level as well. The signal therefore passes through the photonic transponder device for the incoming lightpath's selected wavelength, in essence dropping the signal to the electronic layer, completing an O-E conversion. Thus, $L_{\text{D}} = 1$. In the node's IP switch, the signal is split electronically into multiple disparate signals. Since node B is in fact a reached destination of the request, one of these split signals must be passed to the access network and directed to its desired host. Therefore, assuming each incoming lightpath consumes an entire wavelength, the value of A_{OUT} will be increased by the outgoing signal's bandwidth, 10 Gbps. The other two signals, shown in Fig. 2(a) with a dotted and dashed line, represent the signals which will be sent to the other two destination nodes C and D, respectively. The signals must each be passed back to the optical plane and sent along a lightpath, thereby invoking two separate E-O conversions, and resulting in an L_{A} value of 2. The two signals are then forwarded to their respective destinations along their own lightpaths, completing the utilization of node B, and resulting in $L_{\text{OUT}} = 2$. In this particular example, no traffic arrives from the access network, so $A_{\text{IN}} = 0$ Gbps. All the values necessary to compute node B's consumed power and energy using Eqs. $(1) - (6)$ are now known.
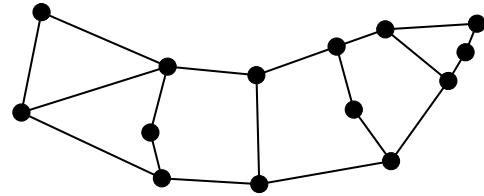
## IV. DUAL-SOURCE PERFORMANCE EVALUATION

In this section we subject both the emission-blind and the emission-aware overlay approaches to various network scenarios and evaluate their relative performances. Initially we study the impact of choosing and building efficient trees on the network performance in scenarios where various portions of the network are powered by renewable energy. Specifically, we consider heterogeneous networks consisting of two distinct types of nodes classified by their energy sources. Nodes powered by fossil fuels are referred to as *black nodes* and are assigned an ϵ_n value of 1 kgCO$_2$e/kWh. This value means that for each kW of power consumed by the node, greenhouse gases equivalent to one kg of CO$_2$ are emitted per hour. Network nodes powered exclusively by renewable energy sources are dubbed *green nodes*, and are awarded an emission-factor of 0.01 kgCO$_2$e/kWh, which indicates they are cleaner than black nodes and may consume more power with a very small adverse environmental effect.[5] The distribution of power sources is uniformly randomized which allows us to make observations on the robustness of the approaches.

Each set of manycast requests consists of 10^5 requests. Manycast requests arrive according to a Poisson process with average arrival rate λ and exponentially distributed holding times with an average service rate μ. The network-wide load in Erlangs is calculated as the ratio of the average arrival rate to the average service rate (λ/μ). For simplicity we define $\mu = 1$ and get a network wide load, that is equal to the arrival rate λ. The source node of each manycast request is uniformly distributed. For each request, the size of the candidate destination set (K) is uniformly distributed from 3 to D_{max} (a parameter that represents the maximum number of candidate destination nodes). For a particular manycast request R_r, with $|D_{R_r}| = K$, the minimum number of destination nodes (K') to reach in order to satisfy the request is set to $\lceil \frac{K}{2} \rceil$. The results presented in this section represent the average of 30 unique request sets. In what follows, we present the results for the augmented Energy-Sciences network topology (ESnet) shown in Fig. 3 for $D_{max} = 10$ and number of wavelengths per link, $W = 16$. We also have examined smaller $D_{max} = \{8, 6\}$ and have found that the results are the same when normalized to the resulting lower blocking probability. **All values for emissions, energy consumption, logical hops and physical hops are normalized to the heuristics exact blocking probability. This allows a fair**
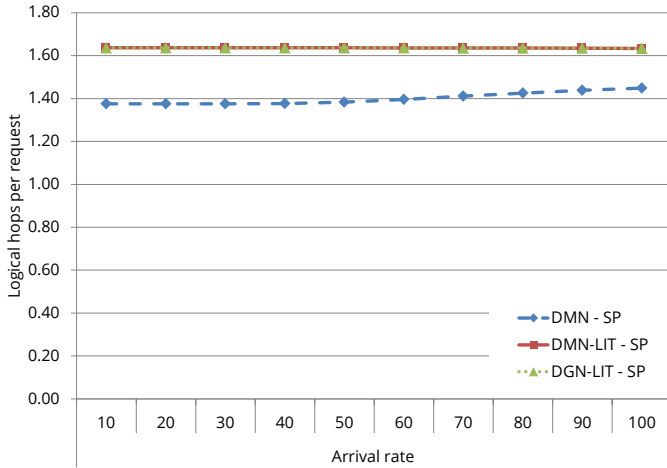
---

[5]We have simulated and examined scenarios with green node distributions throughout the network of 0%, 25%, 50%, and 75%. Due to space restrictions, we present only results for 0% and 50% distribution in this section, but their trends are consistent with the omitted results.
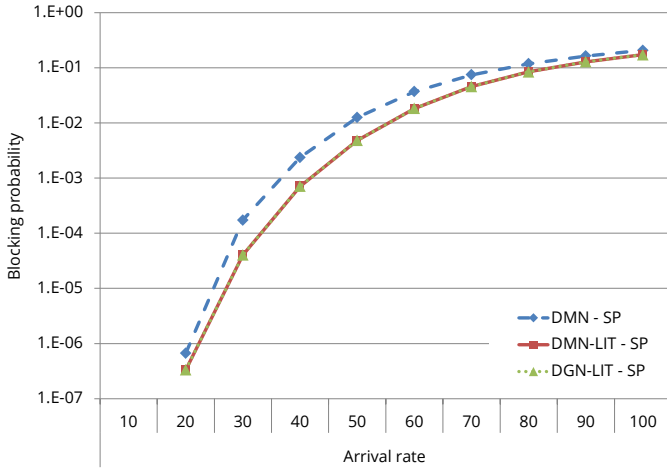
(a) Blocking probability.
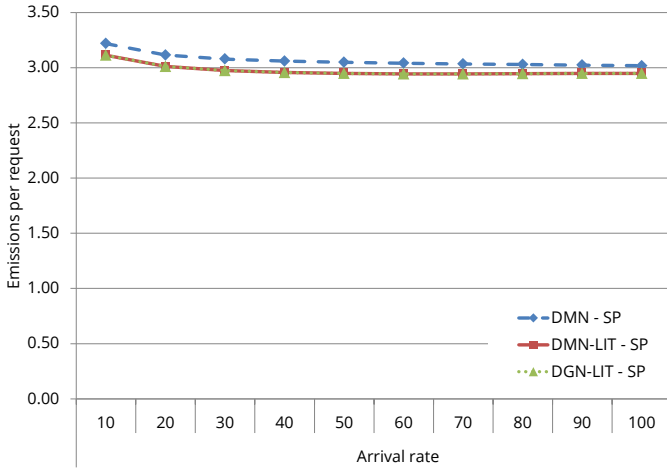
(b) GHG Emissions per request in kgCO$_2$e/kWh.

(c) Logical hops per request.

Fig. 4 – Comparison of overlay algorithms and physical routing techniques with 0% green node distribution throughout the network.

comparison, otherwise the average values for example for energy consumption would decrease simply due to the smaller number of provisioned requests. The normalized values are $\bar{v} = \frac{v}{(1-b_I) \times R}$, where $\bar{v}$ is the normalized value, v the observed total, b_I the blocking probability for an

instance I, and R the number of requests[6]. Requests are assumed to require bandwidth granularity of exactly one wavelength on established lightpaths.

A. Worst-Case Evaluation: No Green Node Availability

We first evaluate a scenario in which the network topology is completely void of green-energy sources for a baseline comparison of the proposed overlay approaches. Fig. 4 shows that the green MA-DMN-LIT approach is superior to the emission-blind MA-DMN even when no green nodes are introduced into the network. The blocking performance of MA-DMN and MA-DMN-LIT are compared in Fig. 4(a). MA-DMN aims to limit resource consumption in building and selecting trees to provision in the network by minimizing wavelength consumption. This results in a high likelihood of selecting many small trees consisting of short unicast paths in quick succession. Consequently, requests requiring overlay-trees built on longer paths will be unprovisionable. By contrast, MA-DMN-LIT aims to find the greenest trees irrespective of their sizes. Therefore, MA-DMN-LIT tends to add longer paths to the overlay-trees. By spreading the consumption of links across the network and using trees with more drop-members in them, fewer links are consumed at any given time, allowing for future requests to make use of the remaining resources, and lowering the overall blocking probability beyond what is capable in its emission-blind counterpart. Recall that in the absence of green nodes, MA-DGN reduces to MA-DMN. Consequently, MA-DGN-LIT performs identically to MA-DMN-LIT. Note also in this figure that we investigate both the shortest path and greenest path physical routing approaches but since there is a 0% distribution of green nodes, there is no distinction between the routes selected in either scheme.

Figure 4(b) depicts the total energy consumed, which is equal to the emissions yield (since $\epsilon = 1$), caused in the network when normalized with each heuristic's blocking probability. By reducing the number of splits and reducing the power needed in the electronic layer, MA-DMN-LIT is able to outperform MA-DMN in the production network-wide GHG emission, even in networks with no renewable energy sources. With all network nodes powered by carbon-based fuel sources, MA-DMN-LIT reduces its produced emissions by **3-5%** over MA-DMN across the shown network loads. Once again, MA-DMN-LIT and MA-DGN-LIT behave identically.

Figure 4(c) provides more clarity to the types of trees used to provision manycast requests under both algorithms. This figure depicts the average logical hop count of provisioned overlay-trees. MA-DMN-LIT has longer trees than MA-DMN, indicating more drop-member nodes. Simply put, MA-DMN-LIT focuses less on building small trees but instead focuses on reducing the number of costly O-E-O splitting at drop-members. As is evident from Eqs. $(1)-(4)$ in Section III, nodes which require O-E-O conversions are the most expensive in terms of power consumption. Nodes which must electronically split incoming signals to multiple output ports electronically consume even more power. Therefore, in terms of power consumption and ultimately energy consumption, those trees

[6]For all presented results we use $R = 100,000$

which use drop-members as splitting points are more expensive than those that use them simply as straightforward drop points. MA-DMN-LIT builds longer, straighter trees than MA-DMN.

B. Uniform 50:50 Distribution of Black and Green Nodes

Figure 5 contains performance results for scenarios wherein half of the nodes in the network are randomly chosen to be green, according to a uniform distribution. The percentage of green destinations emphasizes the advantage that manycast has over multicast; we do not force a certain set of destinations but add the flexibility of choice in heterogeneous topologies. This added flexibility greatly improves chances of successful request provisioning.

Figure 5(a) compares the various heuristic approaches in regard to request blocking. Unsurprisingly, MA-DMN gives an identical performance to its 0% distribution results when incorporated with the shortest path (SP) physical routing scheme. Recall that MA-DMN is emission-blind so no quantity of green nodal distribution will affect its tree building and selection behavior. MA-DMN-LIT outperforms MA-DMN at higher loads, but fails to deliver such success at low loads; MA-DMN results in two orders of magnitude less blocking than its green-aware counterpart. Of the three algorithms, MA-DGN-LIT performs the worst consistently across network loads. MA-DGN-LIT will consistently build and select longer trees than either of its rivals. This claim can be verified from Figs. 5(c) and 5(d) which show the average logical and physical lengths of lightpaths in the network when normalized with blocking. These longer paths lead to greater request starvation and higher blocking. In all cases, greenest path (GP) routing results in longer paths and higher blocking than SP routing **across all traffic loads.** The relative difference between MA-DMN and MA-DMN-LIT is **comparable to shortest path routing** when employing greenest path routing.

Figure 5(b) depicts the network-wide GHG emissions yield for each algorithm normalized against their respective blocking probabilities. As expected, the blind MA-DMN approach with no concern for GP routing produces the most emissions by a fair margin. Even when employing GP routing, MA-DMN still performs no better than the other algorithms. The best-performing algorithm is MA-DGN-LIT. **The margin of improvement over MA-DMN-LIT is more dramatic when employing GP routing.** All three overlay solutions emit more CO$_2$ when employing SP routing. **Notably**, MA-DMN-LIT outperforms MA-DMN in consideration of emissions output**, while maintaining a comparable blocking probability at high loads.** In fact, MA-DMN-LIT produces **10-23%** fewer emissions than its blind alternative. Note that the algorithms yielding the lowest emissions rates are only able to do so at the expense of much higher blocking probabilities**, particularly at low loads.**

When 50% of the nodes are powered by renewable energy sources, the overall emissions are obviously lower than they would be if all nodes were carbon-supplied, but should approach the network's inherent 50% reduction (a savings of approximately 50 metric tons of CO$_2$). MA-DMN with SP routing makes efficient use of the low-emissions availability and is able to reduce its total emissions by approximately 50.1% for all network loads shown. MA-DMN-LIT has a relative reduction of **55-60%**, and MA-DGN-LIT has a relative reduction of **65-67%**. When using GP routing, MA-DMN leads to as much as **51-52.5%** emissions reduction, MA-DMN-LIT leads to **57-64%** reduction, and MA-DGN-LIT leads to **67-74%** emissions reduction. Figure 5(e) illustrates this relative reduction in graphical form and shows what percentage of destination nodes selected throughout the entire duration of the entire request set are green. The percentage of green destinations gains significance depending on the processing job the data is required for. If the destination is green, chances are that the data center at the destination is also powered by renewable sources. This is extremely beneficial to the case where data that is being intensively processed rather than stored. In this case the energy needed can be much larger than the energy needed for the manycast transmission and the green destinations gain more importance and trade-offs in network efficiency might become acceptable.

Lastly, Fig. 5(f) shows the relative energy consumption of the overlay algorithms. MA-DMN-LIT using SP routing is most energy-efficient, whereas MA-DGN-LIT is least efficient. Notice that in cases where GP routing is enabled, paths grow longer on the physical topology, and the amount of energy (regardless of renewable or non-renewable source) is greater than using SP routing.

Figure 6 shows the relative comparison of the scenarios with 50% green nodal distribution compared to 0% for an analysis of how the heterogeneity of the network affects each algorithm. Note that for simplicity, all results shown in the figure assume SP routing only. Studying Figs. 6(b) and 6(c) we see the cost at which the GHG emission reduction comes: increased energy consumption, particularly at low network loads. At a load of just 30 Erlangs, the energy consumption of MA-DGN-LIT is increased by **4%** in order to obtain an emission reduction of **17%** beyond MA-DMN. Since MA-DMN is completely unaware of the degree to which the topology is green, its behavior, and thus its energy consumption is completely unaltered between distributions. As previously noted, the 50% reduction in emissions is due only to the heterogeneity of the network. From these results we learn that for the maximum emission savings, the energy consumption of the network must increase considerably. If for example, a network operator must pay a premium for green energy, the MA-DMN-LIT approach can lead to a **5-10%** emission reduction at only a **1.5-2%** energy consumption increase. If emission savings through an increased number of green data centers/destinations (Fig. 5(e)) is the main objective, those merit the additional **4%** in energy consumption required for MA-DGN-LIT.

### C. Summary of Evaluations: How Green is Too Green?

In this section we have compared overlay approaches which select green trees in combination with algorithms that greedily build trees to the greenest destinations in combination with routing algorithms that establish lightpaths along the greenest paths. In general, MA-DMN-LIT outperforms MA-DMN at
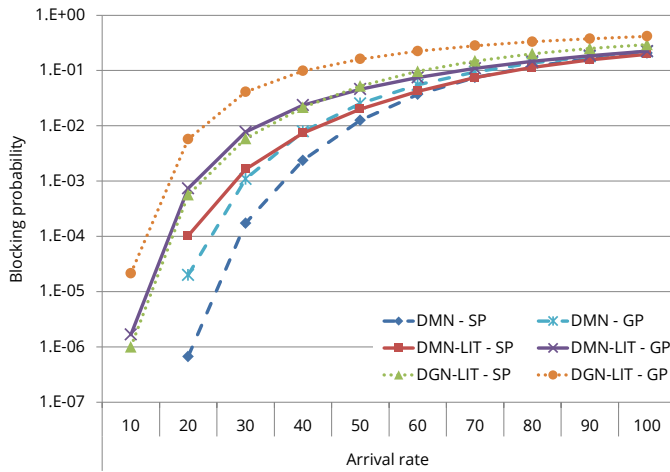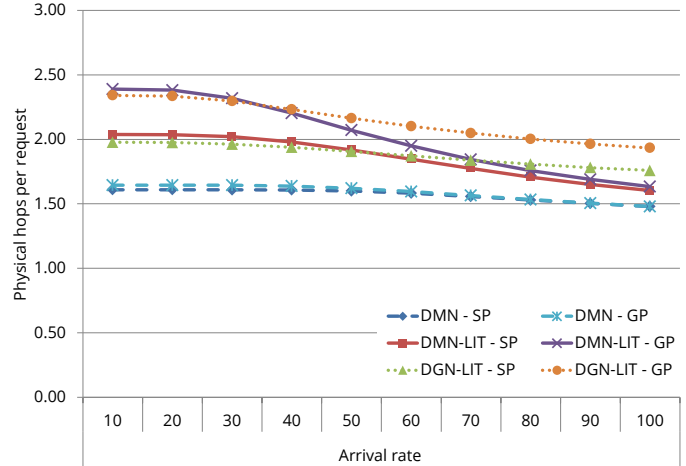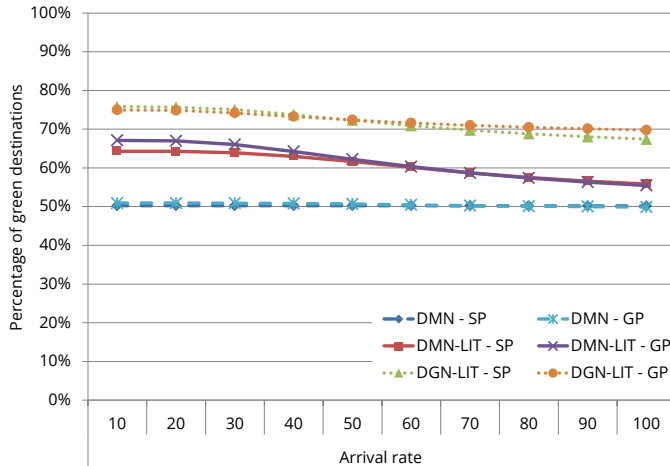
(a) Blocking probability.
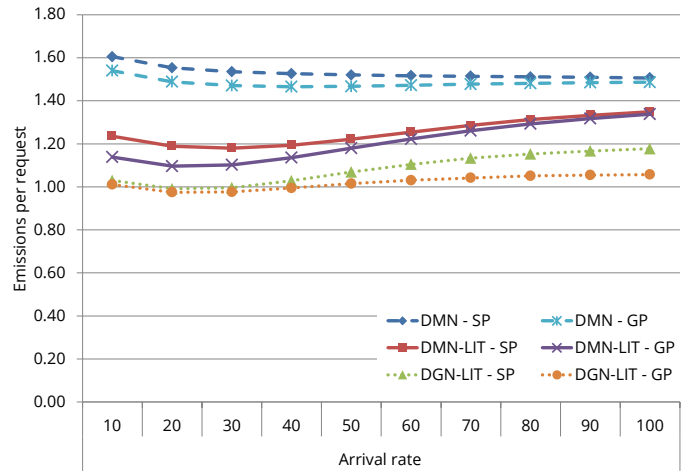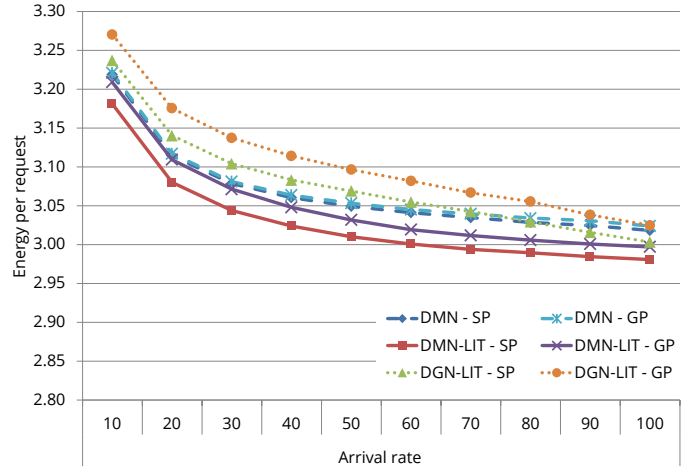
(b) GHG Emissions per request in kgCO$_2$e/kWh.
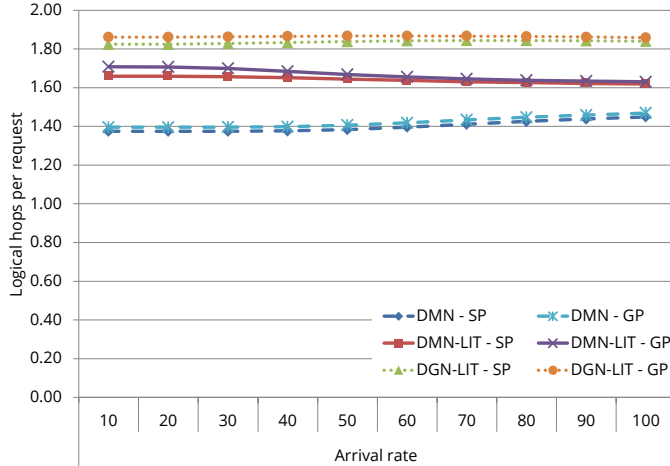
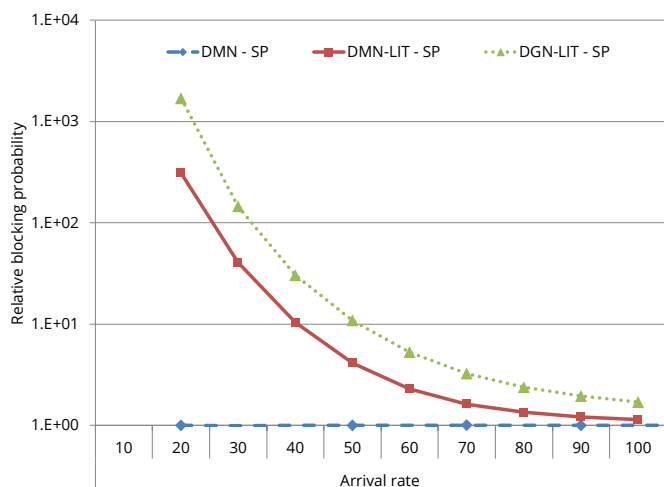(c) Logical hops per request.

(d) Physical hops per request.

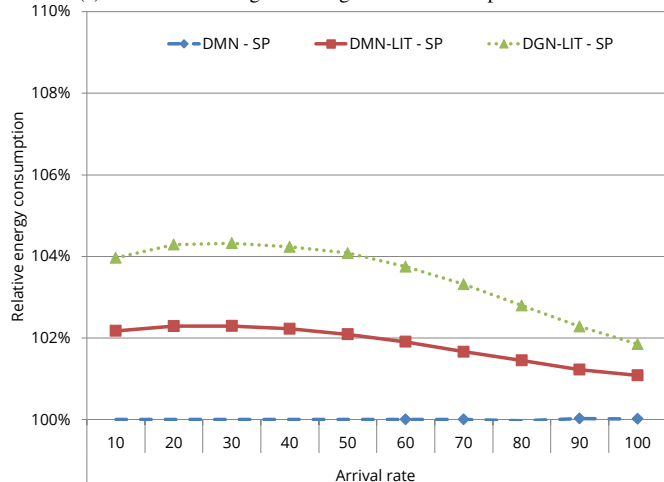(e) Percentage of green destinations.
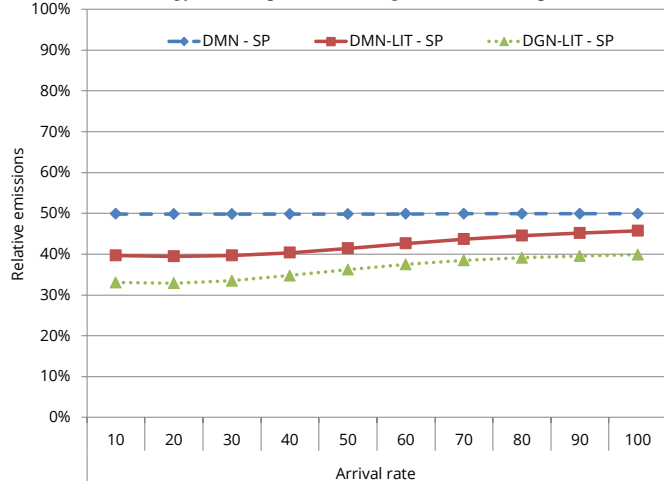
(f) Energy consumption per request in kWh.

Fig. 5 – Comparison of overlay algorithms and physical routing techniques with 50% green node distribution throughout the network.

(a) Relative blocking of 50% green nodes compared to 0%.

(b) Relative energy consumption of 50% green nodes compared to 0%.

(c) Relative GHG emissions of 50% green nodes compared to 0%.

Fig. 6 – Relative performance comparison: 50% green vs. 0% green.

every turn. Even in a completely carbon-fueled network topology, the Least-Impact Tree approach was able to find trees which can reduce blocking and minimize emissions output. When half the network became green, this trend was found to still hold true at high loads. MA-DMN-LIT will build longer trees resulting in longer end-to-end delay, but provides the benefits of emission reduction,**while experiencing comparable blocking probability at higher loads,** to offset this small delay. MA-DMN-LIT is able to find 10% more green nodes than MA-DMN. While MA-DGN-LIT reduces emissions even further than MA-DMN-LIT, it does so at the expense of building much longer trees, and even greater blocking. In terms of overall energy consumption, it also costs more. Only if it can be assumed that at green destinations the data processing can also be realized in renewable energy powered facilities the MA-DGN-LIT approach wins great significance. The only metrics in which using greenest path routing won out over shortest path routing were the number of green nodes reached per request and the emissions yield. The greenest path routing approach is far too costly in terms of blocking probability to reasonably declare it a winning solution. Combining the green routing with the green destination selection and the green tree selection yielded the least impressive results. From these findings, we can conclude that care must be taken in how green resources are utilized in the network to avoid hampering provisioning performance. We have found that the lowest trade-off comes from building short overlay-trees containing short lightpaths and then selecting those trees which turn out to be the greenest overall.

V. CONTINUOUS AND DYNAMIC EMISSION-FACTOR PERFORMANCE EVALUATION

As introduced in the previous section, we model the power-/energy-dependent emissions of networking equipment by multiplying with the greenness factor ϵ where ϵ is a function of space (node location) and time. In the future, the smart Grid will hold information about the amount of energy supplied from renewable sources in real-time at any location [30], [31]. A similar approach is taken by the authors of [32], where information about the current cost of energy is used to choose from distributed data centers. In anticipation of this infrastructure becoming available in the near future, we study the impact of variations to ϵ in and effort to develop robust energy source-aware algorithms.

To determine the appropriate range of simulated emission-factors, we take into account that the average ϵ for the US for example was .665 kg CO_2/ kWh in 2009 according to the Environmental Protection Agency.[7] An average of .396 kg CO_2/ kWh for the whole EU in 2009 has Norway on the low end with just .004 and Estonia with .990 at its other extreme.[8] We therefore model ϵ as a value between 0.01 and 1 in our simulations. We use continuous time-dependent functions as shown in Fig. 8(d) [33] or local distributions as in Fig. 7(d) [34].

Besides the variable availability at different locations, another particularity of renewable energy sources are temporal fluctuations. Examples are photovoltaic sources that are most productive in clear conditions, or wind sources that are unavailable on calm days. In our simulations, the emission-factors are updated at every connection arrival. Since we consider optical circuit reservations, we do not change the

[7]http://www.epa.gov/cleanenergy/energy-resources/egrid/index.html
[8]http://www.eea.europa.eu/data-and-maps/figures/co2-electricity-g-per-kwh

11

route after provisioning. We use a sinusoidal function as shown for adjacent time zones in Fig. 8(d) [33]. This kind of function is typical for availability of solar energy throughout the day, and on a larger scale throughout the year due to changing sun radiation intensity due to the changing angle during different seasons (except on the equator).

A. Wind Power: Geographically Distributed Emission-Factors

In contrast to the previous uniformly random distribution, we now distribute continuous emission-factors across the topology according to data about environmental behavior. The purpose of the presented approach is to study network performance at a point where the potential for renewable energy is exploited to actually generate the energy from renewable source. Our first approach is based on the heatmap shown in Fig. 7(d). Here the colors on the blue/violet end of the spectrum indicate high average annual wind speeds. The colors on the green/yellow end are low speeds. Instead of a uniform, or a balanced (sweeping from east to west) distribution we have localized high availability and localized very low availability. We distribute the emission-factors between 0.01 and 1 according to the colors indicated on the map. Our average ϵ results in an optimistic 0.36 $kg\,CO_2e/\,kWh$ across the country (assuming all the potential wind energy is harvested).

We then use the previously proposed algorithms to provision manycast requests with the same parameters as in Section IV. We only evaluate the shortest path routing versions of the heuristics, since we rely on precomputed shortest paths. In dynamic scenarios with variable ϵ values, the greenest paths between nodes are subject to change, so it would be impractical to use the proposed green-aware static routing approach. Additionally, the results shown in Section IV have shown that GP routing comes at the price of very high blocking even for low network load, which invalidates it as a desirable option.

In Fig. 7(a) we plot the blocking probability for the three algorithms in this scenario. We can observe that the general pattern as observed in Section IV is followed and the MA-DMN-LIT algorithm performs best for medium and high arrival rates. Different from the binary (green vs. black) scenario is that in this case the blocking probabilities are very close for all approaches, and for higher arrival rates they all converge. For Figs. 7(b) and 7(c), which show the average emissions per request and the associated energy consumption respectively, the observations are very similar to the uniformly distributed case too. In fact with the average emission-factor lying between the average ϵs for uniform distributions of 50% ($\bar{\epsilon} = 0.505$) and 75% ($\bar{\epsilon} = 0.2575$) we find that the non uniformity of the regional distribution in the case of the studied topology does not have a large impact on the result. For planning and estimation purposes we can therefore approximate the energy/emission outcomes by assuming a simplified random distribution, as in Section IV, as long as we are considering the target range with the average emission-factor.

B. Solar Power: Dynamically Variable Emission-Factors

Unlike the static emission-factors that are only location dependent and do not vary with time, we now study the impact of dynamically adjusting the emission-factor ϵ throughout the simulation. We subject our manycast algorithms to a similar dynamic scenario as is examined for virtual machines that *follow the sun* to solar powered data centers in [19]. We simulate the changing availability of solar radiation throughout the day with a sine function [33], and offset the four time-zones covered by the ESnet topology. The function is plotted for one 24-hour period in Fig. 8(d). The formula we use to estimate the emission-factors is

$$\epsilon = 1 - 1/2(\sin((2\pi)/24t) + |\sin((2\pi)/24t)|)$$

where t is the average time of sunrise in the respective time-zone. Since the simulations are limited by the number of requests, the entire simulation horizon covers different time-spans at different rates. Assuming that the minimum time unit is 1 hour, the simulated 100,000 requests at a load of 10 Erlangs covers more than a year, but less than two months at 100 Erlangs. Note that nowhere in Fig. 8 do we plot lines for green paths; we build the overlay-trees by combining precalculated routes, thus in a dynamic green scenario, our only option for static routing is to use the shortest paths. After a request is provisioned, the emission factors of its member nodes may change. While the request is still active, we need to repeatedly measure the emissions and keep track of increases and decreases over time. We do so by querying all the nodes in the network for their power consumption, whenever we adjust the emission factor of any of the nodes. We do not consider rerouting trees based on those changes, as we want to provision uninterrupted optical circuits.

Figure 8(a) shows the relative blocking performances of MA-DMN, MA-DMN-LIT and MA-DGN-LIT. We observe that the green destination selection results in the highest blocking probability. This is consistent with previous results, but the difference here is larger compared to the results for static ϵ distributions. The reason is that the longer trees, built to reach the greenest destinations, become increasingly long as the radiation sweeps across the topology and thus use up a large amount of network resources.

In Fig. 8(b), the two lines for the green-aware approaches widely overlap and clearly show savings in the order of **5% to 10%** when compared to the blind MA-DMN approach. Looking at Fig. 8(c) we observe that in this dynamic scenario, the energy consumption is **slightly** lowered by the green-aware algorithms. We conclude that in this dynamic scenario, the two emission-aware approaches perform equally well in terms of energy and GHG yield. MA-DGN-LIT still exhibits an undesirable increase in blocking probability for low network loads however.

VI. Conclusion

Manycast communication is becoming an essential paradigm for large-scale science applications, particularly in scenarios where routing decisions must be made based on cost and availability. In split-incapable networks, these decisions become even more important since the underlying optical core does not support the paradigm directly. Previous work has resulted in
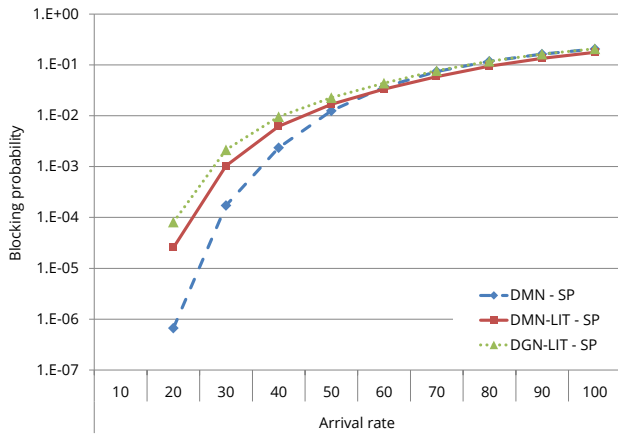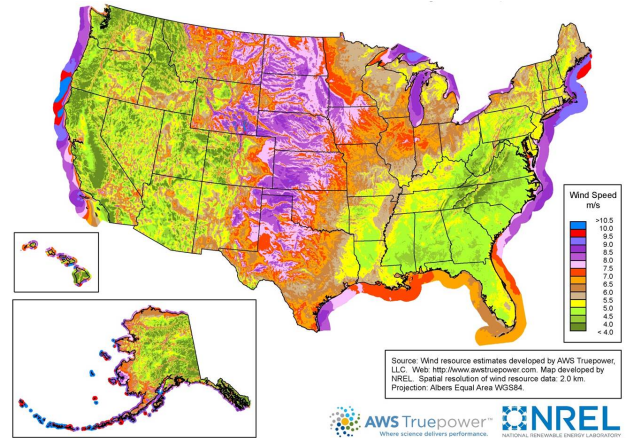
(a) Blocking probability.
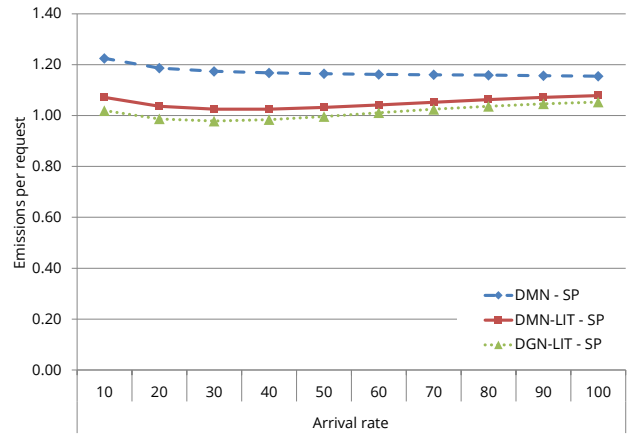
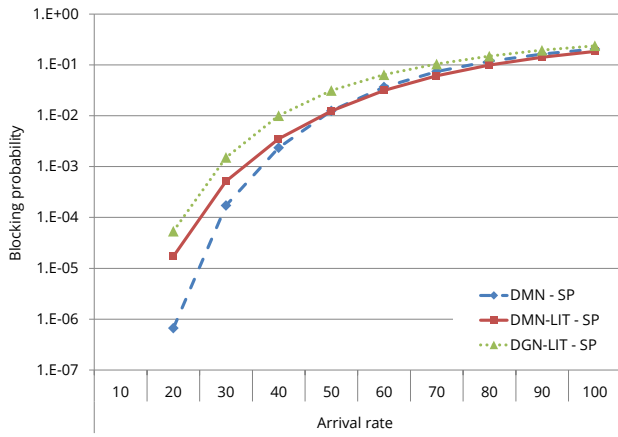(b) Emissions per request in kgCO$_2$e/kWh.
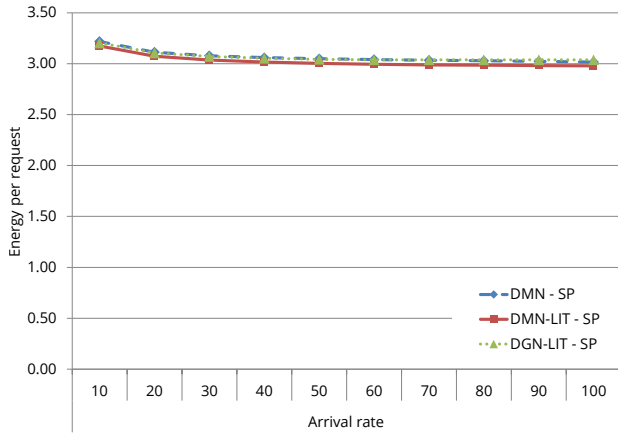
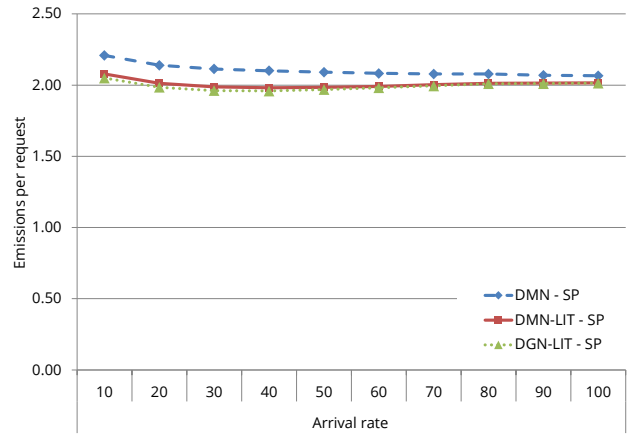(c) Energy per request in kWh.

(d) Wind energy potential map.

Fig. 7 – Overlay algorithm performance for distributed staticϵ according to 7(d). The average ϵ across the country is 0.36 kgCO$_2$e/kWh.

the well-balanced MA-DMN overlay scheme to build overlay-trees that are able to employ O-E-O conversions to split signals into logical branches. While found to be much more efficient than the simpler, single-hop alternatives, MA-DMN had never previously been thoroughly analyzed in situations where energy-efficiency and emissions-efficiency are important factors to consider. Throughout this work we have proposed two emissions-aware alternatives. The first, MA-DMN-LIT, builds trees identically to MA-DMN but selects those trees resulting in the lowest overall potential GHG emissions. The second alternative, MA-DGN-LIT, constructs overlay-trees such that the majority of the energy-intensive O-E-O conversions are handled at network nodes powered by renewable energy resources. By using a realistic, advanced node model to trace the power consumption throughout the network, we have evaluated these approaches and compared them both qualitatively and quantitatively. **Interestingly, under the traditional assumption of a homogeneous power supply, the power considerations alone lowered not only the network's emission production but also the request blocking probability. The proposed MA-DMN-LIT approach is therefore a more economical alternative to the previous MA-DMN approach.** In situations where the network is powered by dual-sources of energy (renewable and non-renewable), we have found that MA-DMN-LIT performs best in terms of trade-offs between

blocking, delay, energy consumption, and emissions yield, and even outperforms the green-blind MA-DMN approach in most of these areas. We have found that MA-DGN-LIT is too greedy in its search for green destinations to outperform MA-DMN-LIT in terms of network performance. If however, large additional savings are possible by an increased number of destinations with green data processing available, MA-DGN-LIT gains significance and the increased blocking rate can be offset by small capacity extensions. We have also subjected these heuristics to scenarios in which there are more than two power sources, each with a distinct geographically-distributed emission impact and found our conclusions for the binary case still hold true. Further, we have considered scenarios in which the greenness of individual nodes varies according to the time of day (analogous to solar power) and have again concluded that MA-DMN-LIT is the preferable approach. **A surprising observation is that the performance of our algorithms in the presence of continuous and even the time dependent emission factors can be approximated roughly by (simpler) static uniformly distributed and binary black/green setups. This enables a forecast of the potential reductions in emissions even with limited knowledge about the actual distribution of renewable energy as long as an** *average* **emission factor is available.** In the future information about the availability of renewable energy at certain locations during
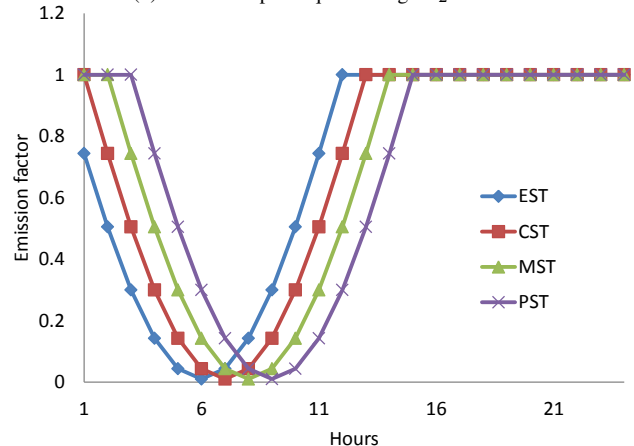
(a) Blocking probability.

(b) Emissions per request in $kgCO_2e/kWh$.
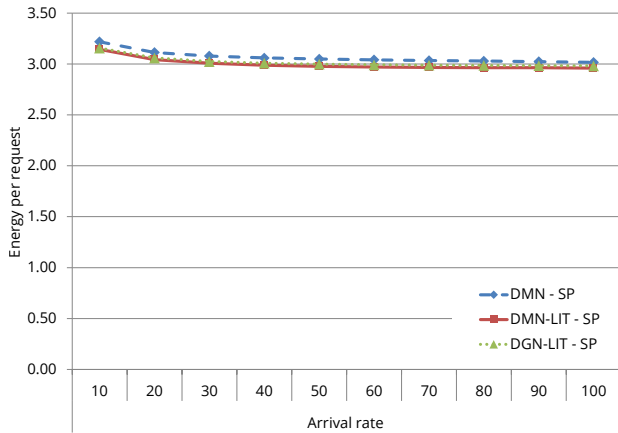
(c) Energy per request in kWh.

(d) Emission-factor function.

Fig. 8 – Overlay algorithm performance for dynamically variable ϵ according to the sinusoidal distribution in (d). The average ϵ during the day is 0.69 $kgCO_2e/kWh$.

certain times will be available through the smart grid. We propose future extensions to fully utilize the potential of emerging software-defined networking to dynamically adapt the network flows with this information and our green routing algorithms. Manycast communication is attractive because of the level of flexibility it offers in selecting communication channels. Future work utilizing this flexibility will be increased manycast survivability, protecting the proposed algorithms against cases of traditional link and node failures as well as failures due to temporary unavailability of renewable sources. The findings of this study provide strong evidence toward the notion that even split-incapable networks can support advanced communication paradigms without compromising their ability to be reliable, efficient, and clean.

REFERENCES

[1] R. Malli, X. Zhang, and C. Qiao, "Benefit of multicasting in all-optical networks," in *Proceedings, SPIE All-Optical Networks*, Nov. 1998.

[2] G. N. Rouskas, "Optical layer multicast: rationale, building blocks, and challenges," vol. 17, no. 1, pp. 60–65, Jan./Feb. 2003.

[3] S. Y. Cheung and A. Kumar, "Efficient quorumcast routing algorithms," in *Proc. IEEE INFOCOM*, Jun. 1994, pp. 840–847.

[4] B. Chen and J. Wang, "Efficient routing and wavelength assignment for multicast in WDM networks," *IEEE J. Select. Areas Commun.*, vol. 20, no. 1, pp. 97–109, Jan. 2002.

[5] R. Ravi et al., "Spanning trees short or small," in *Proceedings, ACM-SIAM symposium on Discrete algorithms*, 1994, pp. 546–555.

[6] A. Gadkar, J. Plante, and V. M. Vokkarane, "Manycasting: Energy-ffficient multicasting in WDM optical unicast networks," in *Proc. Globecom*, July 2011.

[7] J. Plante, A. Gadkar, and V. M. Vokkarane, "Dynamic manycasting in optical split-incapable WDM networks for supporting high-bandwidth applications," Feb. 2012.

[8] Y. Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Energy efficiency in telecom optical networks," *Communications Surveys Tutorials, IEEE*, vol. 12, no. 4, pp. 441–458, 2010.

[9] W. Van Heddeghem *et al.*, "Power consumption modeling in optical multilayer networks," *Photonic Network Communications*, 2012.

[10] F. Idzikowski, L. Chiaraviglio, and F. Portoso, "Optimal design of green multi-layer core networks," in *2012 Third International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy)*, May 2012, pp. 1–9.

[11] F. Idzikowski *et al.*, "Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios," in *2010 14th Conference on Optical Network Design and Modeling (ONDM)*, Feb. 2010, pp. 1–6.

[12] A. Bianzino, L. Chiaraviglio, and M. Mellia, "Distributed algorithms for green IP networks," in *2012 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, Mar. 2012, pp. 121–126.

[13] L. Chiaraviglio, M. Mellia, and F. Neri, "Minimizing ISP network energy cost: Formulation and solutions," *IEEE/ACM Transactions on Networking*, vol. 20, no. 2, pp. 463–476, Apr. 2012.

[14] L. Chiaraviglio and A. Cianfrani, "On the effectiveness of sleep modes in backbone networks with limited configurations," in *2012 20th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, Sep. 2012, pp. 1–6.

[15] M. Marsan *et al.*, "TREND: toward real energy-efficient network design,"

in *Sustainable Internet and ICT for Sustainability (SustainIT), 2012*, Oct. 2012, pp. 1–6.

[16] L. Wang *et al.*, "GreenDCN: a general framework for achieving energy efficiency in data center networks," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 1, pp. 4–15, Jan. 2014.

[17] M. Gattulli, M. Tornatore, R. Fiandray, and A. Pattavina, "Low-carbon routing algorithms for cloud computing services in IP-over-WDM networks," in *IEEE International Conference on Communications*, 2012.

[18] Y. Chen and A. Jaekel, "Energy optimization in optical grids through any-casting," in *Communications (ICC), 2013 IEEE International Conference on*.   IEEE, 2013, pp. 3835–3839.

[19] M. Lemay, K.-K. Nguyen, B. St. Arnaud, and M. Cheriet, "Toward a zero-carbon network: Converging cloud computing and network virtualization," *IEEE Internet Computing*, vol. 16, no. 6, Nov. 2012.

[20] J. Wang *et al.*, "Green-aware routing in GMPLS networks," in *International Conference on Computing, Networking and Communications (ICNC)*, 2012.

[21] X. Dong, T. El-Gorashi, and J. Elmirghani, "IP over WDM networks employing renewable energy sources," *Journal of Lightwave Technology*, vol. 29, no. 1, pp. 3 –14, Jan. 2011.

[22] B. Aksanli, T. Rosing, and I. Monga, "Benefits of green energy and proportionality in high speed wide area networks connecting data centers," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2012*, Mar. 2012, pp. 175–180.

[23] S. Ricciardi *et al.*, "An energy-aware dynamic RWA framework for next-generation wavelength-routed networks," *Computer Networks*, vol. 56, no. 10, pp. 2420–2442, Jul. 2012.

[24] M. P. Boddie, T. Schöndienst, and V. M. Vokkarane, "Dual power source aware algorithms for green optical network survivability," in *IEEE Online Conference on Green Communications*, 2012.

[25] T. Schöndienst, J. Plante, D. Davis, and V. M. Vokkarane, "Energy source-aware manycast overlay in WDM networks," in *Globecom 2013 - Optical Networks and Systems Symposium (GC13 ONS)*, Atlanta, USA, Dec. 2013.

[26] T. Szymanski, "Future Internet video multicasting with essentially perfect resource utilization and QoS guarantees," in *Proc. IEEE 19th International Workshop on Quality of Service*, 2011.

[27] L. Neves, J. Krajewski, P. Jung, and M. Bockemuehl, "SMARTer2020," Global e-Sustainability Initiative, Tech. Rep., Dec. 2012.

[28] World Resources Institute (WRI), World Business Council for Sustainable Development (WBCSD), Carbon Trust, and Global e-Sustainability Initiative (GeSI), "GHG protocol product life cycle accounting and reporting standard ICT sector guidance," January 2013, DRAFT.

[29] J. Sathaye *et al.*, "Renewable energy in the context of sustainable development," *IPCC Special Report on Renewable Energy Sources and Climate Change Mitigation, Geneva, Switzerland*, p. 707, 2011.

[30] U.S. Department of Energy. Recovery act smart grid programs. [Online]. Available: www.smartgrid.gov/recovery_act/project_information

[31] ——. Microgrid workshop agenda and presentations. [Online]. Available: e2rg.com/microgrid-2012/Workshop_Agenda-08092012.pdf

[32] L. Rao, X. Liu, L. Xie, and W. Liu, "Coordinated energy cost management of distributed internet data centers in smart grid," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 50–58, Mar. 2012.

[33] Afshin, Michael Andreas, "NREL: measurement and instrumentation data center (MIDC) home page." [Online]. Available: http://www.nrel.gov/midc/

[34] "NREL: dynamic maps, GIS data, and analysis tools - wind maps," Sep. 2013. [Online]. Available: http://www.nrel.gov/gis/wind.html