Energy-Awareness in Dynamic Traffic Grooming

M. M. Hasan[1], **Farid Farahmand**[2], **and Jason P. Jue**[1]

[1]*Dept. of Computer Science, The University of Texas at Dallas, Richardson, Texas 75083*
[2]*Dept. of Engineering Science, Sonoma State University, Rohnert Park, California 94928*
Email: hasanmm@utdallas.edu, farid.farahmand@sonoma.edu, jjue@utdallas.edu

Abstract: We introduce an energy-efficient traffic grooming scheme for promoting greener optical networks. The scheme considers a modular node architecture, reuses already active components during request allocations, and conserves total energy consumption in the network.

OCIS codes: (060.4251) Networks, assignment and routing algorithms; (060.4256) Networks, network optimization

1. Introduction

To cope with future rapid growth, resources in backbone networks (e.g., core routers) are often over-provisioned [1]. In most cases, resources remain active to their full capacities for 24 hours a day regardless of diurnal traffic variations and lower demand. Therefore, resources remain under-utilized by a wide margin, resulting in unnecessary energy waste, increased operating cost (OPEX), and detrimental environmental effects. Thus far, less attention has been given to energy-awareness in optical networks compared to energy-awareness in wireless and copper counterparts (a literature review is available in [2]). The broad prospect of the IEEE 802.3az Task Force is to "enable new system level energy management techniques that will save energy beyond the network interface" [3]. As part of the larger goal, this study focuses on energy-aware protocols for wired optical networks.

Modern routers are becoming more modular (i.e., functional independency of components) such that they can operate in a distributed manner, e.g., Cisco 12000, Cisco 7600, and Cisco CRS-1 routers [4]. The Cisco Integrated Services Router solution has a module-based sleep mode where power delivery is automatically assigned to network modules, and idle modules neither charge nor dissipate. Such a modular architecture not only mitigates bandwidth scaling and dynamic power management problems, but also provides better fault tolerance and facilitates air-cooling at the points of presence [5]. In this study, we assume modular nodes. A node is divided into two main sections (Fig. 1(a)): photonic and electronic. Typically, the electronic section supports all electronic processing, switching, and routing operations. The photonic section can include optical cross-connects (passing through and switching lightpaths in optical form) or optical add/drop multiplexer (ADM) (terminating lightpaths and passing the wavelengths to the electronic ADM). We investigate the impact of having a modular architecture in the photonic ADM section of a node in terms of energy consumption (Fig. 1(b)). The modular architecture allows incoming traffic to be directed to already provisioned ports, modules, and chassis to minimize energy consumption of the node; unused equipments can be powered off. As more traffic enters the node, additional chassis, modules, or ports can be provisioned as needed. Since traffic requests in wavelength division multiplexing (WDM) networks are usually of sub-wavelength granularity, they are often groomed together in the nodes along their routes [6]. The main contribution of this paper is that we introduce energy-awareness in dynamic traffic grooming and show that modular nodes offer substantial energy-savings in network operation. Using an auxiliary graph based heuristic, we associate routing to the physical node architecture, rather than the logical node architecture traditionally used in optical network research. We provide experimental results comparing our approach to traditional approaches.

2. Problem Statement and Auxiliary Graph based Schemes

We assume that the physical topology of a network is already designed and, hence, capital expenditure (CAPEX) is fixed. We are given the node modularity, i.e., the number of chassis per node, the number of modules (or shelf cards) per chassis, and the number of ports (or transceivers) per module. Power consumption rates of components depends on state (e.g., active and idle states as in [3,5]), and the current state of lightpaths is known. Our target is to reduce OPEX via minimizing power consumption for dynamic requests. That is, we need to find routing and wavelength assignment for incoming requests, possibly grooming with existing lightpaths, that minimizes the total power consumption in the network.

Since the traffic grooming problem is NP-hard [6], we approach our case with heuristics. We first build an auxiliary graph (AG) representing the physical architecture of a given network, assign cost proportional to energy consumption rate of components, and then apply the shortest (i.e., the minimum cost) path algorithm on the AG. As shown in Fig. 2, the proposed AG model has W+5 layers for each node, where W is the number of wavelengths per fiber. Each layer has

---

This work has been supported in part by the National Science Foundation (NSF) under CNS grant number 0627128.

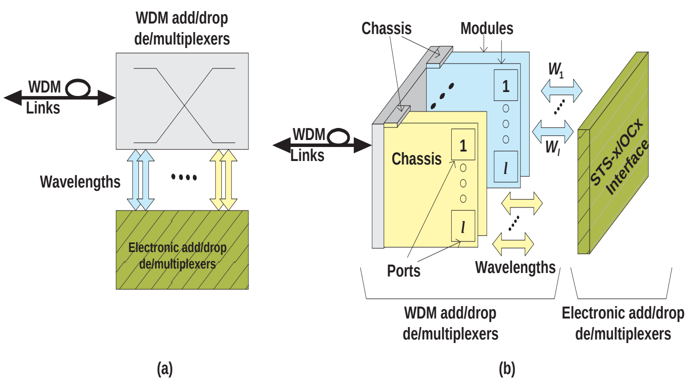

Fig. 1. An example of modular physical architecture of a node.

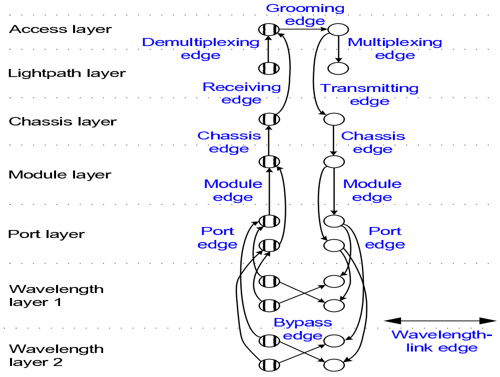

Fig. 2. An auxiliary graph model of a node for W = 2.

two groups of vertices: input and output. In the access layer, an electronic traffic flow begins and ends. In the lightpath layer, a lightpath originates and terminates. For each chassis, module, and port in a node, there is an input-output vertex-pair in the chassis layer, module layer, and port layer, respectively. For each wavelength, there is a dedicated layer, as shown in Fig. 2 for $W = 2$. A wavelength-link connects two nodes in the wavelength layer. The total number of vertices in the AG for one node is 2 × (2 + (number of chassis, modules, and ports) + ($W \times$ node-degree)). This number is a constant and, therefore, the AG model remains tractable.

Once the AG is constructed, the edges are weighted based on optimization objectives. When energy-awareness is not a concern (i.e., "traditional" grooming), we use high weight (or cost) for transmitter and receiver edges. This means that the shortest path based algorithm on the AG will find paths for a request where the number of lightpath set up is reduced. When energy-awareness is a concern, we assign cost proportional to the energy consumption of different components of a switch. For example, in the Cisco Catalyst 6500 series, the approximate power usage for a port is 3 watts, a line card/module is 315 watts, and a chassis is 375 watts (including power consumption for switching fabric, fan, and other parts) [7]. We use these power consumption rates in our study and proportionally set the weight of edges among chassis, module, and port in the AG. For a request arrival, the least cost path is chosen dynamically. If the path traverses the wavelength layer, the corresponding wavelength-link edges are removed from the AG and a new lightpath-link edge is established in the lightpath layer of the AG between the two nodes. For a request departure, the wavelength-link edges used in the lightpath setup are deallocated, if no other requests are using that lightpath.

3. Experimental Results

We examine the performance of traffic grooming, with and without energy-awareness, in terms of energy consumption. We consider two well-known NSF networks: 14-node, 21-link network and 24-node, 43-link network [6]. Each node in the 14-node network is assigned at most 2 chassis, 2 modules per chassis, and 2 ports per module. Each node in the 24-node network is assigned at most 3 chassis, 4 modules per chassis, and 4 ports per module. Wavelength capacity is OC-192 and traffic requests are OC-1, OC-3, OC-12, OC-48, or OC-192. The default value for W is 4. The dynamic traffic fluctuations in backbone networks follow a predictable daily cycle and have peaks and valleys as the sun moves [8]. We imitate such a traffic profile taken from a 10 GE Atlanta-Houston connection on a working day in the Internet2 backbone network [9]. We scale the traffic profile based on the network under consideration such that the network capacity remains more than twice the peak hour demand [1]. The network capacity is measured by (2 × number of links × number of wavelengths per link × wavelength capacity) and the load added to the network by a traffic request is measured by (request size × shortest hop distance between the source and the destination). Requests are uniformly distributed among nodes. Rather than using the "time-zone dynamism" as in [8], we use the time of the day dynamism in this study. Results are presented with a 95% confidence interval.

As shown in Fig. 3 for the 14-node network and Fig. 5 for the 24-node network, the total energy consumption is always lower in the proposed energy-aware approach than in the traditional approach. At off-peak hours both approaches have fewer requests to be groomed and the energy-saving is not significant. At peak hours in the energy-aware approach, the energy-saving is higher since more traffic requests are routed through already used (or activated) components by existing traffic. In the traditional case, requests that can not be groomed may or may not be routed through already activated components. This attribute is further illustrated in Fig. 4, which shows that the average physical length (or hop count) of a lightpath is indifferent to traffic pattern in the traditional case. On the other hand, in the energy-aware case, the physical length of a lightpath is increased at peak hours since requests follow longer paths to be provisioned through already activated components. Consequently, we observe as much as 30% energy-saving at peak hours. The

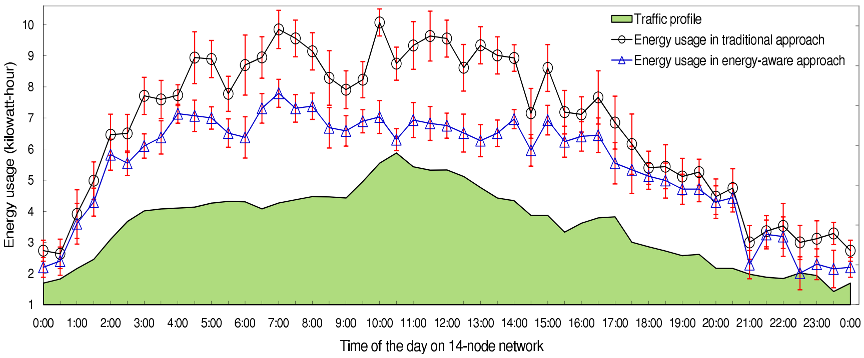

Fig. 3. Energy consumption for the 14-node network.

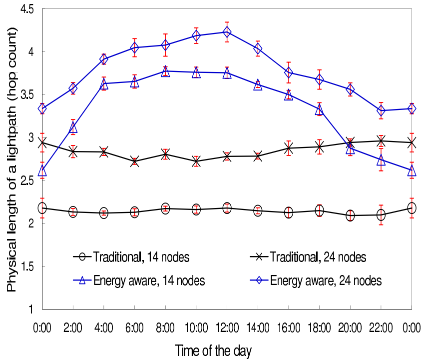

Fig. 4. Average lightpath length.

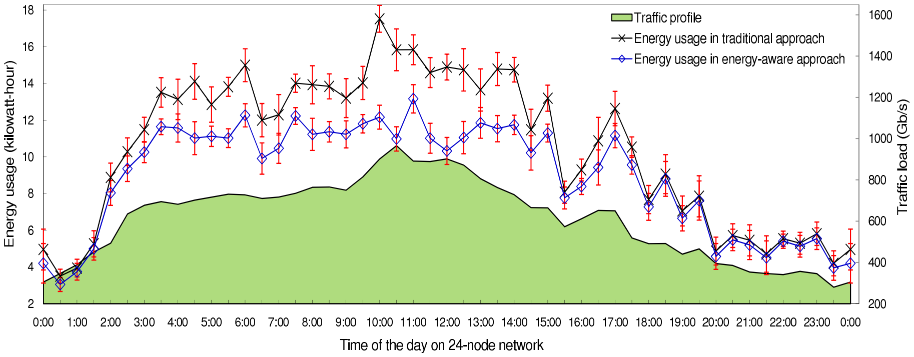

Fig. 5. Energy consumption for the 24-node network.

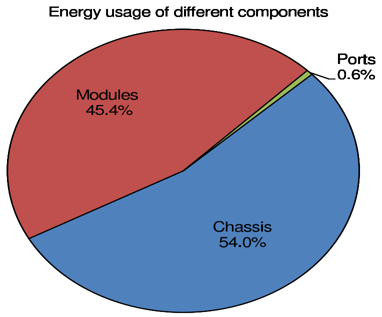

Fig. 6. Energy consumption by different components.

bigger the network is, the more total energy (and money) can be saved. The results in Fig. 3 and Fig. 5, however, do not include energy expense for components that are not used by any requests. In practice, all devices, even if not being used, are usually kept active in current telecommunication networks. Comparing to such cases, our energy-aware approach should provide even more energy-savings.

The distribution of energy expense among chassis, modules, and ports are shown in Fig. 6 for the 14-node network at a peak-hour (10:00 am). As expected, the dominant components are chassis and modules. Therefore, rather than reducing total number of ports, the careful choice of ports during routing (i.e., locations of ports on modules and chassis) may generate substantial energy conservation via minimizing active modules and chassis. One can also optimize energy expense in a node by using more modules per chassis and more ports per module, unless there are impositions from physical characteristics (e.g., design modularity, scalability, and air cooling [5]).

4. Conclusion

In this paper, we investigate energy-aware dynamic traffic grooming for WDM optical networks. We assume that switches are modular in terms of power management and operations. We apply simple heuristics over an auxiliary graph model that represents the relationship among physical architecture, lightpath level, and wavelength level. For the daily traffic variations, when compared to approaches lacking energy-awareness, the proposed approach can reduce energy consumption in the network that is significant both in monetary and ecological points of view.

References

1. E. B. Desurvire, “Capacity Demand and Technology Challenges for Lightwave Systems in the Next Two Decades,” IEEE/OSA Journal of Lightwave Technology **24,** 4697–4710 (2006).
2. M. M. Hasan, F. Farahmand, and J. P. Jue, “Traffic Grooming for Energy-Aware Optical Networks,” In *Tech. Report, UTDCS-31-09*, (Dept. of Computer Science, The University of Texas at Dallas, 2009).
3. IEEE P802.3az Energy Efficient Ethernet Task Force, Available online: http://www.ieee802.org/3/az/index.html, Accessed: Feb 2009.
4. G. Schudel and D. J. Smith, “Chapter 1: Internet Protocol Operations Fundamentals,” In *Network World, Cisco Press*, (2008).
5. J. Chabarek et al., “Power Awareness in Network Design and Routing,” In *Proc. of IEEE INFOCOM*, pp. 457–465 (Phoenix, AZ, 2008).
6. B. Mukherjee, *Optical WDM Networks* (Springer, New York, 2006).
7. Cisco® Power Calculator, Available online: http://tools.cisco.com/cpc, Accessed: Feb 2009.
8. C. Cavdar et al., “CATZ: Time-Zone-Aware Bandwidth Allocation in Layer 1 VPNs,” IEEE Communications Magazine **45,** 60–66 (2007).
9. Internet2 networking consortium, Available online: http://www.internet2.edu/observatory/archive/data-views.html, Accessed: Sep 2009.