

Deep-Neural-Network-Based Wavelength Selection and Switching in ROADMs

Weiyang Mo,  Craig L. Gutterman, Yao Li,  Shengxiang Zhu, Gil Zussman, and Daniel C. Kilper

Abstract—Recent advances in software and hardware greatly improve the multi-layer control and management of reconfigurable optical add-drop multiplexer (ROADM) systems facilitating wavelength switching. However, ensuring stable performance and reliable quality of transmission (QoT) remain difficult problems for dynamic operation. Optical power dynamics that arise from a variety of physical effects in the amplifiers and transmission fiber complicate the control and performance predictions in these systems. We present a deep-neural-network-based machine learning method to predict the power dynamics of a 90-channel ROADM system from data collection and training. We further show that the trained deep neural network can recommend wavelength assignments for wavelength switching with minimal power excursions.

Index Terms—Machine learning; Power excursions; ROADM systems; Wavelength switching.

I. INTRODUCTION

Growing dynamic traffic demands for Internet applications, including HD video rendering, cloud computing, and the Internet of Things (IoT), motivate more efficient networks capable of handling a wide range of applications [1]. Dynamic reconfigurable optical add-drop multiplexer (ROADM) systems, in which connections are established through real-time wavelength switching, have long been studied as a means to achieve greater scalability and increase the network resource utilization [2]. However, today's commercial ROADM systems remain "quasi-static," with wavelengths being provisioned to meet the peak traffic requirements and left in place [3]. While ROADMs are extensively deployed in today's wavelength-division multiplexing (WDM) systems, they are primarily used for flexible wavelength provisioning without real-time switching functionality. Software-defined networking (SDN) potentially provides software control capabilities that might be exploited to achieve real-time wavelength switching,

but its scalability and flexibility are limited by various types of physical layer impairments [4–6].

A key unresolved challenge in achieving dynamic ROADM systems through SDN is predicting and controlling the optical power dynamics resulting from wavelength switching operations. Power excursions can result from the interactions between the wavelength-dependent gain and automatic gain control of optical amplifiers, Raman scattering in the fiber, and other wavelength-dependent phenomena. Deviations of the channel powers outside pre-allocated system margins can potentially result in service disruption due to reduced quality of transmission (QoT) [7]. For this reason, today's commercial systems take minutes and even hours to provision a wavelength through time-consuming power adjustments along an optical path [8]. To realize dynamic ROADM systems, we implement a deep neural network that predicts power excursions resulting from wavelength switching operations. After training a 90-channel multi-hop ROADM system including eight erbium-doped fiber amplifiers (EDFAs) and five ROADM nodes with 67,200 training samples, the deep neural network is able to recommend wavelength assignments for wavelength switching in randomly loaded systems with over 99% precision.

The remainder of this paper is organized as follows. In Section II, we discuss the basics of power excursions and related work to address power excursions. In Section III, we introduce the principle of the proposed deep neural network approach. The experimental setup is discussed in Section IV. In Section V, we discuss the deep neural network architecture, data collection, training, and power excursion prediction. The performance of the deep neural network is evaluated against different metrics to show its effectiveness to mitigate power excursions in wavelength switching operations. In Section VI, we address the scalability of the proposed approach for large-scale ROADM systems. We conclude our findings in Section VII.

II. PROBLEM STATEMENT

Recent work has extensively investigated advanced modulation formats to improve the spectral efficiency and network capacity of WDM transmission systems [9]. However, these spectrally efficient modulation formats

Manuscript received February 21, 2018; revised March 31, 2018; accepted April 9, 2018; published May 8, 2018 (Doc. ID 322820).

W. Mo (e-mail: wmo@optics.arizona.edu), Y. Li, S. Zhu, and D. C. Kilper are with the College of Optical Sciences, University of Arizona, Tucson, Arizona 85721, USA.

C. L. Gutterman and G. Zussman are with the Department of Electrical Engineering, Columbia University, New York, New York 10027, USA.

<https://doi.org/10.1364/JOCN.10.0000D1>

require tighter QoT margins due to lower tolerance to both optical noise accumulation [impacting the optical signal-to-noise ratio (OSNR)] and fiber-nonlinearity-based impairments. As a result, the reduced transmission distances are further compromised by large margins that are needed to account for optical channel power variations or uncertainties. Thus, optical power dynamics that arise from wavelength switching operations become especially problematic in these systems. Furthermore, optical power dynamics often include the phenomena that switching a wavelength on one channel causes power changes on other channels [7].

One main manifestation of optical power dynamics is the transient effect in an optical amplifier. The transient effect is fast power overshoots and undershoots that arise from sudden changes in input power due to wavelength switching operations or upstream fiber cuts. For automatic gain controlled (AGC) EDFAs, a fast feedforward control loop can be implemented to augment the slower feedback control loop to effectively and rapidly suppress the transient effect. The feedforward control loop has a response time of 1 μ s that can immediately adjust the pump power based on a pre-defined relationship between the pump current and input power for a target gain [10].

A different form of optical power dynamics in optical amplifiers—power excursions that result from the interactions between the wavelength-dependent gain and AGC of optical amplifiers—can occur in wavelength switching operations. In the case of these power excursions, wavelength switching operations lead to persistent power differences on surviving channels, which are then corrected over long time scales using individual channel power controls in the ROADM nodes. Power excursions can grow in magnitude over cascaded amplifiers and cause substantial service disruptions. In recent work, 15-dB power excursions were reported in a WDM transmission system with recirculating loops totaling 2240 km [11]. For this reason, introducing or provisioning a new channel into a ROADM system is a time-consuming process that requires repetitive small-step power adjustments by sequentially actuating many optical components along an optical path to ensure that the powers of all surviving channels are within pre-allocated margins. In commercial-scale transmission systems, the fastest reported wavelength provisioning time is several minutes for a single 400 Gbps wavelength channel over a long-distance link [8].

There have been a number of approaches to address these amplifier-based power excursions. A fast tunable source was implemented to distribute a single optical signal over two wavelengths—one with a high gain and the other with a low gain—to equalize the mean gain and cancel out the power excursions [12]. An analytical solution was studied in Ref. [13] to mitigate power excursions based on a pre-measured EDFA gain spectrum. However, the gain spectrum does not consider the tilt change during wavelength switching operations, and as a result only a 5%–15% power excursion reduction is achieved. An optical probing method was also investigated to measure the EDFA gain spectrum without causing power excursions on surviving

channels and thus recommend an optimal wavelength assignment with minimal power excursions [14]. Nevertheless, previous work relies on either specific system designs or specialized hardware, and as a result increases the total hardware cost. Conversely, machine learning offers a more flexible solution without special hardware requirements. Particularly, machine learning has been used to promote the development of intelligent optical communication systems [15,16]. Through the extensive data collection of the power excursions versus changing channel loadings, a machine learning model can be trained to accurately recommend new wavelength assignments that will not cause power excursions. Previous machine learning applications examined wavelength assignment and defragmentation to minimize the channel power divergence or standard deviation of surviving channels, which primarily arises from the static gain ripple and tilt of EDFAs [17,18]. Regression models, such as ridge regression and kernelized Bayesian regression, were investigated to predict the channel power divergence in a 24-channel single-hop ROADM system, but such regression models do not consider the interactions between WDM channels and are unlikely to accurately predict the power excursions in wavelength switching operations. In order to accurately predict power excursions for WDM transmission systems, including multiple ROADM hops and full C-band WDM channels, a more sophisticated machine learning model based on a deep neural network is investigated in this paper. In this work, we extend a recent analysis of neural-network-based wavelength switching in Ref. [19]. We provide additional analysis on the computational complexity, overfitting reduction, and early termination using the deep neural network. The performance is also compared against random forest, showing the advantage of the deep neural network in learning and predicting complex power excursions. The scalability of the deep neural network in large-scale transmission systems is discussed, and we propose two approaches as our future work.

III. PROPOSED MACHINE LEARNING METHODOLOGY

Machine learning has been developed to allow computers to learn to do a specific task without being explicitly instructed. Machine learning problems can be divided into two general categories—supervised learning problems and unsupervised learning problems. Supervised learning analyzes the training data and produces a relationship between an input object and the desired output object, which can be used for predicting the output of new input objects. Unsupervised learning problems try to draw inferences from datasets only consisting of input data. In this work, the focus is on developing a supervised machine learning model to predict power excursions based on an initial set of channels and the addition of a new set of channels. The data set includes the impact of complex interactions between channels that result in the power excursion response. A popular machine learning model for solving such complex problems is deep neural networks.

Deep neural networks are computational models that are inspired by the biological neural networks in the




Fig. 1. (a) Schematic diagram of a neuron. (b) Illustration of a deep neural network containing two hidden layers.

human brain [20]. The basic unit of a deep neural network is a neuron (also known as a node or unit) as shown in Fig. 1(a), which receives the input from other neurons and computes the output. In the real world, most data are nonlinear, and we want these neurons to learn complex nonlinear representations. Therefore, nonlinear activation functions are introduced to the output of neurons to improve neural network approximations. Common types of nonlinear activation functions include tanh (hyperbolic tangent), sigmoid, and ReLU (a unit ramp function). Recent work has reported the advantages of ReLU because ReLU does not cause the “gradient vanishing” problem (which can completely stop the neural network from further training) [21]. However, the optimal activation functions will still depend on the particular applications and need to be determined by trial and error during the training process. Deep neural networks combine many layers of neural networks to find complex relationships and abstractions from the input data to understand and approximate the output.

A deep neural network consists of three types of layers as shown in Fig. 1(b): (i) input layers, which contain input neurons that provide information from the outside world; (ii) hidden layers, which contain hidden neurons that perform nonlinear transformations from the input layer to the output layer (a deep neural network may contain multiple hidden layers) and (iii) output layers, which contain output neurons that predict the output to the outside world. The initial weights of the neural network are randomly set based on a probability distribution determined by the user. The first stage of training the neural network is forward propagation. The input vector is propagated through the neural network to determine the corresponding output. A cost function C is used to measure the accuracy of the predicted output \hat{y}_i and the corresponding true output y_i .

Common cost functions include mean square error (MSE) for regression [Eq. (1)] and cross entropy log loss for classification [Eq. (2)],

$$C = \frac{1}{2m} \sum_{i=1}^m (y_i - \hat{y}_i)^2, \quad (1)$$

$$C = \frac{1}{m} \sum_{i=1}^m (-y_i \log(\hat{y}_i) - (1 - y_i) \log(1 - \hat{y}_i)). \quad (2)$$

The next step is to determine how to update the network weights in order to minimize the cost function. A mini-batch (only a user-selected small subset of the training set in each training iteration) gradient descent is commonly used to determine the direction of steepest descent and how each weight in the neural network should be updated [22].

In order to optimize the performance of the deep neural network, validation data is used to compare the performance of a deep neural network with different parameters. The configuration that minimizes the specified loss function is chosen, and the test data is used to get an unbiased view of the performance of the deep neural network.

IV. EXPERIMENTAL SETUP

A metro-scale multi-hop ROADM system shown in Fig. 2 is built to study wavelength switching using the proposed machine learning approach. At the transmitter, a 90-channel comb source with spacing of 50-GHz is used to create 90 WDM channels from 191.60 to 196.05 THz (i.e., 1529.2–1564.7 nm in wavelength). The power of the transmitter is then equally divided into four equal outputs using a 1×4 splitter, and each output is sent to a different ROADM (ROADM 1–ROADM 4) to create different channel loadings. The ROADM system consists of five ROADMs, which are separated by four standard single-mode fiber (SSMF) spans. Each SSMF span contains two dual-stage AGC EDFAs to compensate for the loss of the ROADMs and the transmission fiber and one variable optical attenuator (VOA) to increase the span loss to match the average 18-dB amplifier gain. Two-stage EDFAs realize the tilt-control by adjusting the attenuation of the variable optical attenuator (VOA) in the first stage, taking advantage of the fact that the tilt is dependent on the internal gains of the individual stages [14]. The tilt of each EDFA is adjusted in




Fig. 2. Schematic of the experiment setup including five ROADM nodes, four fiber spans, and eight EDFAs with different gain characteristics. The training, validation, and test data are collected by reconfiguring the channel loadings and measuring the power excursions.

order to create wavelength-dependent gain and study the power excursion mitigation; however, the peak-to-peak gain variation is kept within ± 0.5 dB, which is typical for line amplifiers. Two-stage EDFAs realize the tilt control by adjusting the attenuation of the variable optical attenuator (VOA) in the first stage, taking advantage of the fact that the tilt is dependent on the internal gains of the individual stages [14]. The VOA of the EDFA will be automatically adjusted by its internal controller based on the user-specified tilt requirement. The tilt is defined as the peak-to-peak gain variation of a least-squares-fitted line of the channel gains over the full signal band (ranging from 1529.2 to 1564.7 nm). Figure 3 shows the gain spectrum of the two EDFAs in the first span. The tilt of the first EDFA is set to -0.4 dB to compensate for the stimulated Raman scattering (SRS) in the transmission fiber. The tilt of the second EDFA is set to 1.0 dB. The EDFAs in the other three spans have the same gain and tilt settings, thus giving rise to a similar gain spectrum. After cascading four transmission spans, the cumulative peak-to-peak gain variation across the C-band is measured to be 3.8 dB. Note that 4 – 6 dB gain variation is typically allowed between ROADMs nodes depending on the system design. Such peak-to-peak gain variation along with the AGC operation results in substantial power excursions in wavelength switching operations. Each ROADM is comprised of multiple wavelength selective switches (WSSs), per-channel VOAs, and per-channel optical channel monitors (OCMs). The power at the drop point is tapped to the OCM for per-channel power measurement. To measure the power excursion, we first measure the per-channel power of initial channels at the channel drop point before the wavelength switching operation. Then, we measure the per-channel power of initial channels at the channel drop point after the wavelength switching operation (i.e., a new wavelength channel is added into the system). The power excursions are measured by taking the differences between them. Note that in general ROADMs will perform gain equalization for all the output ports to remove the power excursion accumulated in the previous link. This operation is time consuming and requires spectral analysis; it would therefore come after the channel add or drop event and would be used to remove any residual power excursions. Using the deep-neural-network-based wavelength

switching thus minimizes these gain equalization operations, resulting in more stable system operation and faster turn-up times for new channels.

The effectiveness of machine learning is evaluated for power excursions that occur on top of the static power divergence due to the EDFA gain ripple and tilt. Thus, the system is initially configured to remove the static power divergence. Two types of channels are identified for wavelength switching operations—initial channels and new channels. First, the VOAs in each ROADM are initialized to ensure uniform 0 dBm launch power per channel into the transmission fiber (i.e., 19.5 dBm total power) with the 90 initial channels (i.e., with 90 -channel WDM input). These attenuation values are stored as a reference for newly added channels. Note that the VOA initialization can largely mitigate the channel power divergence due to static wavelength-dependent gain in the EDFAs and static SRS in the transmission fiber. However, the initial VOA values cannot guarantee uniform 0 dBm power per channel when the initial channel loading is changed in later experiments due to EDFA tilt change, EDFA power excursions, and dynamic SRS. Thus, VOA adjustment is executed to remove any power variations before wavelength switching, as would normally be done in system operation.

V. RESULTS AND DISCUSSION

In this section, we describe the data collection process, deep neural network architecture, training process, power excursion estimation, and wavelength assignment recommendation using the trained deep neural network. A deep neural network is first built and trained. Its performance is evaluated with regard to the number of training samples, the speed of the training process, and the accuracy in power excursion prediction and wavelength assignments. The performance of ridge regression and random forest methods are evaluated against the deep neural network for comparison purposes.

A. Data Collection

The first step in learning the complex optical power excursion response is extensive data collection of the power excursion response under a variety of channel loadings. However, such a data collection process is time-consuming due to the speed limitation of hardware actuation and software control. In this experiment, collection of each data sample takes approximately 3 s on average, including the latency of control signaling, WSS actuation along an optical path, VOA adjustment, and power excursion measurement. Note that the VOA adjustment is executed only once on each initial channel for each initial channel loading to remove any power variations, as would normally be done in system operation. The WSS actuation to turn on the new channels (with no additional VOA adjustments) and the power excursion measurement are executed for each wavelength switching operation. Collecting data might even take longer in commercial large-scale systems, and




Fig. 3. EDFAs in the first span with different gain spectra. (a) Wavelength-dependent gain spectrum of the first EDFA. (b) Wavelength-dependent gain spectrum of the second EDFA.

thus potentially imposes an obstacle to using machine learning in practical ROADM systems. Methods to overcome or to address these implementation issues will be discussed later.

In this experiment, 1680 training cases are used to train a deep neural network, each of which contains 40 power excursion measurements (i.e., 67,200 training samples in total) as the following process: 40 available wavelength positions for adding a new channel are randomly selected, and the maximal power excursion among all initial channels is measured by switching on and off these 40 wavelength positions one by one. In addition, 210 validation cases are collected for evaluating how well the deep neural network is trained and which parameters provide optimal prediction performance. Finally, another 210 testing cases are collected for evaluating the prediction accuracy and the performance of wavelength switching using the deep neural network. In total, 84,000 data samples are used in this experiment, taking approximately 70 h for collection. Note that in this experiment all channels are sent through the longest route over four spans (i.e., ROADM 1-2-3-4-5), but the machine learning methodology is applicable to the multi-route case by recording the added ROADM and the dropped ROADM as additional input and collecting more training samples. Strategies to minimize the number of training samples while ensuring the prediction accuracy will be detailed in the training section.

B. Deep Neural Network Architecture

A deep neural network is built to predict the power excursion that occurs when adding a new channel into the multi-hop ROADM system. The input of the deep neural network is a 180-element binary vector as shown in Eq. (3). The first 90 binary input features (which correspond to 90 wavelength locations) are used to represent the wavelength locations of initial channels. A “1” represents that the wavelength is initially lit or occupied, and a “0” represents that the wavelength is not initially lit. The next 90 binary input features represent the wavelength locations of new channels being added into the systems. The “1s” represent the new wavelength locations of new channels added into the system. Note that in this experiment, we focus on the channel add operation, and we will show that minimized power excursions are guaranteed for channel add operations. Since the channel drop operation is the reverse process of the channel add operation, minimized power excursions are also guaranteed for the corresponding channel drop operations. The output that we aim to predict is the maximum power excursion among all initial channels as shown in Eq. (4),

$$\vec{x} = [\vec{x}_{\text{initial}}, \vec{x}_{\text{new}}] = [x_1, x_2, \dots, x_{90}, x_{91}, \dots, x_{180}] \in \{0,1\}^{180}, \quad (3)$$

$$\vec{y} = \max_{j \in \text{Initial channel}} (\Delta P_j). \quad (4)$$

The optimal neural network architecture is determined by varying a number of parameters. The optimized parameters include: the number of hidden layers, the number of




Fig. 4. Architecture of the deep neural network. The input layer contains 180 features representing the “on” or “off” state of initial channels and new channels. The output layer contains a single output, representing the maximal power excursion among all initial channels.

neurons per layer, the activation function of hidden layers, the number of iterations (or epochs), the learning rate, the L2 regularization term, and the dropout rate. The performance is determined by minimizing the root mean square error (RMSE) against the validation set. The RMSE can be interpreted as the standard deviation of the difference between observed and predicted values [in decibels (dB)]. A lower RMSE indicates a more accurate prediction. The architecture with the lowest validation RMSE depicted in Fig. 4 includes four hidden layers using a combination of tanh and ReLU activation functions. Other parameters are summarized in Table I, and the details of training the deep neural network will be discussed in the next section.

C. Training

It is important to minimize the data collection time while still ensuring prediction accuracy. With a small number of available training samples, the deep neural network tends to overfit the specific training samples, resulting in a low training error. However, since these training cases do not well represent the full set of behaviors in the system, this causes a high variance that generates a high prediction error in the test samples as new data samples have not been seen by the deep neural network. Getting more training samples can effectively reduce the variance and better generalize the model, but the prediction performance might saturate at some point because of the existence of a small bias that limits further improvement of learning performance. The bias in this experiment mainly arises from actual system and measurement errors such as time-varying penalties (e.g., temperature change) and measurement uncertainty (e.g., OCM inaccuracy). For example, the

TABLE I
PARAMETERS OF THE OPTIMIZED DEEP NEURAL NETWORK

Parameter	Value
Neurons in hidden layers	(180,120,30,15)
Activation function	(tanh, tanh, ReLU, ReLU)
L2 regularization	0.001
Dropout rate	0.1
Initial learning rate	0.005
Number of epochs	217

accuracy of OCMs in this experiment is ± 0.1 dB, and it may happen that two data samples with the same features give rise to different target outputs.

In this experiment, online training is implemented in a control plane to determine the number of training samples that are needed as follows: 210 testing cases and 210 validation cases are first collected. The online training of the deep neural network contains repetitive processes. For each process, 168 more training cases (i.e., 6720 more training samples) are added to the training set to train the deep neural network and calculate the RMSE of the 210 testing cases. The online training continues until the test RMSE does not decrease with two consecutive processes. Figure 5 illustrates the training RMSE and test RMSE over a varying number of training samples (also called the learning curve) during online training of the deep neural network. The training error curve shows the difference between the prediction based on the training data compared against the actual training data, whereas the test error curve shows the error in predicting the power excursion for different sets of random data (the test set). With just 6720 training samples, there is a large difference between the training error and the test error due to large variance on different samples, indicating a poor generalization of the model. With the increased number of training samples, the training error and the test error start to converge, but a gap will persist due to the inherent variance of the samples. Online training stops at 67,200 training samples because there is no RMSE decrease over two consecutive processes.

The deep neural network shown in Fig. 4 is trained to minimize the RMSE using a mini-batch stochastic gradient descent (SGD) with a mini-batch size of 64. In order to prevent overfitting, regularization techniques, including L2 regularization [23] and dropout [24] are implemented. Several combinations were tested with a varying L2 regularization value and dropout rate in each hidden layer, and we found an L2 regularization of 0.001 and a dropout rate of 0.1 can effectively prevent overfitting and achieve a low RMSE. The initial learning rate is set to 0.005 and is

adapted in the training stage to allow for fine weight updates. The learning rate adaptation is multiplied by 0.99 every epoch.

It is important to reduce the training time, as a long training time for a system can add significant cost. One important metric that decides the training time is the number of epochs in the training stage, since the total training time is linearly proportional to the number of epochs. In this experiment, the trade-off between the prediction accuracy and the number of epochs is evaluated by comparing the accuracy of the neural network predication using the validation data set. Figure 6 shows the RMSE (in dB) of the training set and the validation set in the training stage as a function of the number of epochs. Initially, the RMSE of both the training set and the validation set significantly decreases with the increased number of epochs. After approximately 200 epochs, although small fluctuations exist, the RMSE of the validation set shows minimal improvement. In fact, the RMSE of the validation set is 0.103 dB after 200 epochs and 0.100 dB after 900 epochs. In this experiment, the training is terminated if three consecutive epochs fail to decrease the RMSE of the validation set. By introducing this rule to the training stage, the training stage is terminated at the 217th epoch with a validation RMSE of 0.104 dB. Compared with the validation RMSE of 0.100 dB at the 900th epoch, there is negligible performance difference, but the training time is reduced by more than a factor of four.

D. Performance Evaluation

After the training stage has completed, a check on the deep neural network performance is carried out against the test set using different metrics. For comparison purposes, ridge regression and random forest methods are also evaluated against the test set. For the ridge regression model, the regularization parameter was tuned to 0.01 through cross validation. For the random forest model, 200 trees (with 180 features being considered for each tree)




Fig. 5. Root-mean-square error (RMSE) of the training set and the test set as a function of the number of training samples during online training.




Fig. 6. RMSE as a function of the number of epochs in the training stage. The training stage is terminated at the 217th epoch with the validation set RMSE of 0.104 dB.

are found to provide the best performance while ensuring minimal training time. Note that we also evaluated the support vector machine (SVM), but its computational time does not scale well to a large number of training samples [25].

First, the RMSE and the maximal prediction error of the entire test set are evaluated, and the results are summarized in Table II. A lower RMSE indicates a more accurate prediction. Similarly, a lower maximal prediction error reveals a better fit under corner cases. The deep neural network outperforms ridge regression and random forest by more than a factor of two in the RMSE and the maximal error. Random forest results in a worse performance because it is not able to learn the complex nonlinear relationship among 180 features with the given 67,200 training samples. Ridge regression performs worse than the deep neural network because ridge regression does not take into account the interdependencies between input features. The prediction errors using different machine learning models can also be viewed in Fig. 7. The actual power excursion ranges from 0 to 3.5 dB, and black dashed lines indicate the perfect prediction. It is seen that the deep neural network obtains significantly lower errors between the actual power excursions and the predicted power excursions, and its accuracy is stable over the entire power excursion range. On the other hand, both ridge regression and random forest result in high prediction errors, particularly when the actual power excursion is above 2 dB.

The second metric used to evaluate the performance is the mean square error of the channel (MSEC) at a particular wavelength. A low MSEC indicates a high prediction accuracy for the particular wavelength, while a high MSEC indicates that the particular wavelength may not be considered as a potential candidate for wavelength switching due to a substantial prediction error. Figure 8 shows the




Fig. 8. MSEC as a function of wavelength locations using different machine learning approaches. The deep neural network not only provides less prediction error but also more stable performance across the entire 90 channel spectrum.

MSEC of 10 different wavelengths. The deep neural network efficiently keeps the MSEC below 0.02 dB^2 across all 10 wavelength locations, while the maximal MSEC using ridge regression and random forest can be as large as 0.11 dB^2 and 0.13 dB^2 , respectively. Moreover, the MSEC is stable among all 90 channel wavelength locations with a standard deviation of 0.004 (note that only 10 wavelengths are shown in the figure), indicating that the deep neural network is trustworthy to make an accurate prediction over the entire spectrum. In contrast, ridge regression and random forest result in much higher standard deviations of 0.03 and 0.04, respectively.

Third, we evaluate the δ -recommendation accuracy, which is the proportion of test cases in which the deep neural network is able to recommend a wavelength with a power excursion within a δ margin from the minimal power excursion (which is achieved by switching on the optimal wavelength). A higher δ -recommendation accuracy indicates the model is able to accurately recommend wavelengths for wavelength switching operations within a tighter power excursion bound. Figure 9 shows the δ -recommendation accuracy as a function of the δ margin

TABLE II
TEST RMSE AND MAXIMAL PREDICTION ERROR

Machine Learning Model	RMSE (dB)	Maximal Error (dB)
Deep neural network	0.104	0.8
Ridge regression	0.273	2.3
Random forest	0.281	2.7

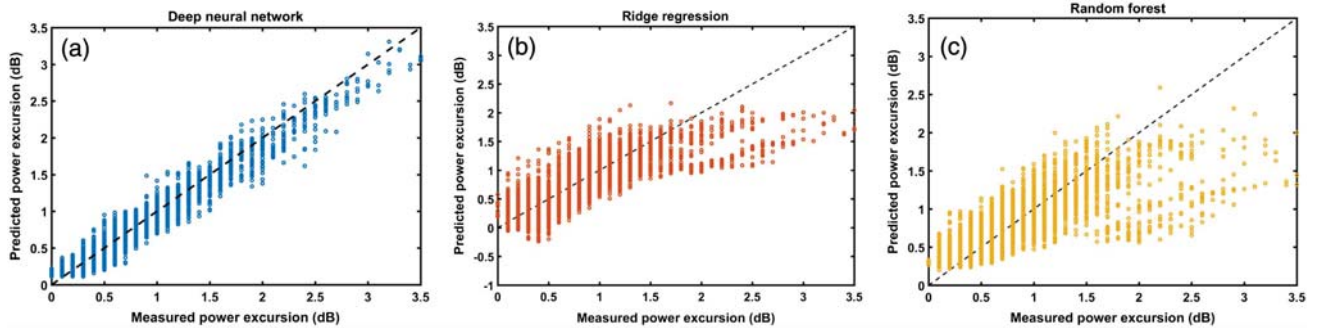


Fig. 7. Predicted power excursion versus measured power excursion over the test set. (a) Deep neural network. (b) Ridge regression. (c) Random forest. Both ridge regression and random forest underestimates the power excursion when the actual power excursion is above 2 dB.




Fig. 9. δ -recommendation accuracy as a function of δ margin from the actual minimal power excursion. The deep neural network is able to recommend the actual optimal wavelength 79.5% of 210 test cases.

using the deep neural network, ridge regression, and random forest. In this experiment, the minimum δ margin is 0.1 dB, taking into account the ± 0.1 dB precision in the power measurement. When the δ margin is set to 0.1 dB, such that the system has the strictest requirement of wavelength assignments (i.e., the exact optimal wavelength must be predicted by the model), the deep neural network can recommend the optimal wavelength among 40 wavelength candidates over 79.5% of the time (i.e., 167 test cases out of 210 test cases), while ridge regression and random forest only achieves 41.4% and 56.7% recommendation accuracy. When the δ margin increases to 0.4 dB (i.e., the actual power excursion of switching on the recommended wavelength must be within 0.4 dB from the minimal power excursion), the recommendation accuracy of the deep neural network is 100%, while the recommendation accuracy of ridge regression and random forest are only 89.5% and 96.2%. We also note that although random forest demonstrates a better accuracy than ridge regression under a small δ margin, its performance gets saturated after $\delta = 1.2$ dB. This result indicates that random forest leads to a higher variance over the new data set with high power excursions (i.e., predict some test cases pretty well but others poorly). Note that these tests are conducted over a finite sized, randomly generated data set within a very large space of possible values, and therefore this bound does not guarantee accuracy over the full range of possible events.

Next, the classification accuracy is assessed for different power excursion thresholds using receiver operating characteristic (ROC) curves as shown in Fig. 10. A better classification accuracy indicates a more powerful model that is able to separate good wavelength candidates from the bad ones for a given system power excursion threshold. In this experiment, the classification accuracy is checked against two different power excursion thresholds—0.5 dB and 1.5 dB—in according to the system QoT requirement reported in our previous work [12]. The classification is

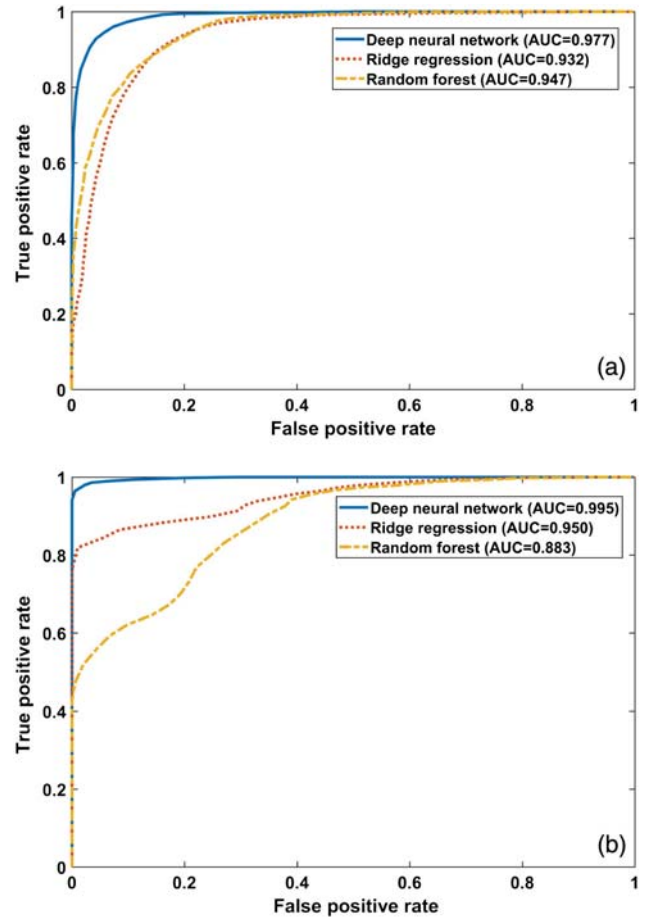


Fig. 10. Receiver operating characteristic (ROC) curves to assess the classification accuracy for different system power excursion thresholds. (a) 0.5-dB threshold; (b) 1.5-dB threshold.

evaluated by two metrics: (i) The ability to separate positive cases from negative cases, which is quantified by the true positive rate (TPR) at a given false positive rate (FPR). A positive case means that the deep neural network recommends a channel as being within a given decision threshold, and a negative case means that the deep neural network rejects a channel as being outside a given decision threshold. Note that the decision threshold is used by the machine learning model to determine whether a potential wavelength is positive or negative, which is different from the system power excursion threshold. TPR is the ratio of correct positive predictions to all actual positives, and FPR is the ratio of incorrect positive predictions to all actual negative predictions. A perfect classification model is able to obtain 100% TPR while maintaining 0% FPR. The ROC curve is formed by connecting all TPR/FPR pairs, each of which corresponds to a different decision threshold. (ii) The area under the ROC curve (AUC). The AUC varies from 0.5 to 1, where 0.5 is the performance of a random classification model and 1 is the performance of a perfect classification model. Figure 10 shows the classification accuracy under 0.5 dB and 1.5 dB thresholds using the deep neural network, and the performance is compared to ridge regression and random forest. With a 0.5-dB power

excursion threshold, the deep neural network obtains the best classification accuracy with a TPR of 80.4% while ensuring the FPR less than 1% and the AUC of 0.977. When the system power excursion threshold is increased to 1.5 dB, the deep neural network obtains the TPR of 97.1% with less than 1% FPR and the AUC is 0.995. We also note the interesting behavior of random forest for which the classification accuracy goes down (with an AUC from 0.947 to 0.883) when the system power excursion threshold is increased from 0.5 to 1.5 dB. This indicates that random forest tends to estimate the power excursion to be less than 1.5 dB, when the actual power excursion is above 1.5 dB.

Finally, we evaluate the PTPR, which is defined as the precision at a specific TPR under a system power excursion threshold. The precision is the ratio of true positives to the number of total positive values predicted. Keeping a high PTPR is important because minimizing the chance of adding a channel with a power excursion beyond the system margin (which may disrupt the whole transmission system) is more important than missing a possible valid channel candidate. Thus, a high PTPR guarantees reliable system operations with a minimal possibility of system disruption due to wavelength switching operations. Figure 11 shows the PTPR curve with 0.5-dB and 1.5-dB thresholds

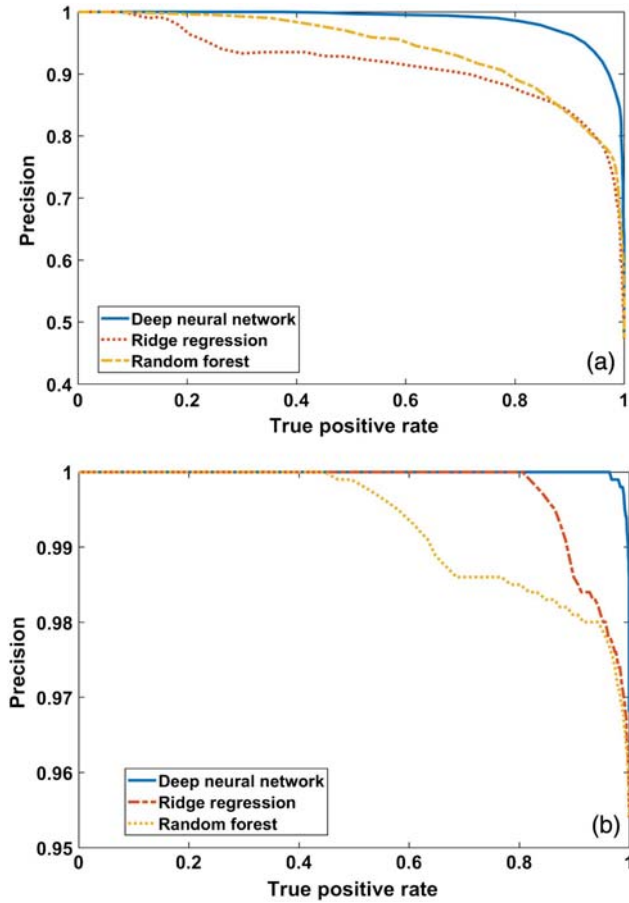


Fig. 11. PTPR curves using different machine learning models with two different system power excursion thresholds. (a) 0.5-dB threshold; (b) 1.5-dB threshold.

using different machine learning models. With a 0.5-dB power excursion threshold, the deep neural network obtains a precision of over 99% while ensuring a TPR of greater than 76% (i.e., ensure less than 1% false positives but also misses roughly 24% valid wavelength candidates). For comparison, ridge regression and random forest only obtain the TPR of 13.1% and 35.4%, respectively, in order to achieve the same precision. When the power excursion threshold is increased to 1.5 dB, the deep neural network is able to obtain a 100% precision while obtaining a 96.4% TPR (i.e., ensure zero false positives while missing only 3.6% valid wavelength candidates).

VI. SCALABILITY OF THE MACHINE LEARNING APPROACH

The results proposed in the previous sections reveal that a deep neural network can efficiently reduce, and in some situations, bound the power excursion in a 90-channel DWDM transmission system including four SSMF spans and eight EDFAs, thus allowing rapid wavelength switching operations. Future work needs to consider the scalability and implementation in large-scale mesh networks. Mesh optical networks involve a larger number of spans and optical amplifiers resulting in an increase in complexity. Methods to address this in future work can be twofold as follows.

First, practical WDM transmission systems are developed in a mesh topology containing multiple multi-degree ROADMs. As shown in Fig. 12, network edges that connect pairs of ROADM nodes may carry channels with different wavelengths (i.e., the set of wavelengths entering each EDFA is different), and each edge may also contain a different number of EDFAs. In these situations, machine learning that is trained in a distributed manner can be used to reduce the learning complexity as shown in Fig. 12. An individual machine learning model can be applied to predict the output power along individual edges. Using a centralized SDN controller, the resulting power excursions can be predicted by combining the individual prediction algorithms along the edges of the path. Moreover, the SDN




Fig. 12. Illustration of the machine learning approach that individual machine learning models are trained in a distributed manner. Different colors represent amplifiers from different vendors. For a new end-to-end connection request (dashed red line), the output power can be estimated by combining the prediction of individual neural network models.

controller needs a coordinate lightpath setup that traverses multiple edges, taking into account the wavelength continuity and the system QoT.

Second, it is important to address how the machine learning approach will scale with a different network, especially for a practical transmission system with the increased dimensions and complexity. It is very common to have commercial systems upgrade to support higher capacity, and it is impractical to train a new system every time from scratch. In order to avoid the time-consuming training data re-collection and machine learning model re-training, transfer learning can be applied to migrate the knowledge trained under one transmission system to a new transmission system. Previous work has used transfer learning to predict different 16 QAM systems using a trained deep neural network under a QPSK/16QAM system with just 20 new training samples [26]. Thus, a reference system can be trained in the test lab before deployment and transfer learning can be used on systems at the time of initial deployment. Online learning can further refine the machine learning model over time as the system performance evolves. The prediction of the power excursion response in different transmission systems using transfer learning and using online learning will be investigated in future work.

VII. CONCLUSION

A deep neural network is implemented to predict the dynamic power excursion of a 90-channel DWDM transmission system, containing four SSMF spans and eight EDFAs. The deep neural network is able to learn the complex optical power excursion response with 67,200 training samples and obtains a 0.1-dB RMSE for 8400 random test samples. Based on the predicted power excursions, the deep neural network can recommend valid wavelengths for wavelength switching with a precision over 99% over the tested samples. The deep neural network was also shown to be far more effective than regression and random forest models. This work is a first step in applying a deep neural network to rapid wavelength switching. The future work will investigate the deep neural network approach in large-scale networks along with transfer and online learning techniques for practical implementations.

ACKNOWLEDGMENT

This work is supported by National Science Foundation (NSF) grants CNS-1650669, CNS-1423105, CNS-1650685, and PFI-1601784.

REFERENCES

- [1] D. C. Kilper, K. Bergman, V. W. S. Chan, I. Monga, G. Porter, and K. Rauschenbach, "Optical networks come of age," *Opt. Photon. News*, vol. 25, no. 9, pp. 50–57, 2014.
- [2] A. Mahimkar, A. Chiu, R. Doverspike, M. D. Feuer, P. Magil, E. Mavrogioris, J. Pastor, S. L. Woodward, and J. Yates, "Bandwidth on demand for inter-data center communication," in *10th ACM Workshop on Hot Topics in Networks*, 2011, p. 23.
- [3] I. Tomkos, S. Azodolmolky, J. Sole-Pareta, D. Careglio, and E. Palkopoulou, "A tutorial on the flexible optical networking paradigm: state of the art, trends, and research challenges," *Proc. IEEE*, vol. 102, no. 9, pp. 1317–1337, 2014.
- [4] Y. Li, W. Mo, S. Zhu, Y. Shen, J. Yu, P. Samadi, K. Bergman, and D. C. Kilper, "tSDX: enabling impairment-aware all-optical inter-domain exchange," *J. Lightwave Technol.*, vol. 36, no. 1, pp. 142–154, 2017.
- [5] S. Yan, F. N. Khan, A. Mavromatis, D. Gkounis, Q. Fan, F. Ntavou, K. Nikolovgenis, F. Meng, E. H. Salas, C. Guo, C. Lu, A. P. T. Lau, R. Nejabati, and D. Simeonidou, "Field trial of machine-learning-assisted and SDN-based optical network planning with network-scale monitoring database," in *European Conf. on Optical Communication*, 2017.
- [6] Y. Li and D. C. Kilper, "Optical physical layer SDN," *J. Opt. Commun. Netw.*, vol. 10, no. 1, pp. A110–A121, 2018.
- [7] D. C. Kilper, H. Rastegarfar, M. Bhopalwala, and W. Mo, "Optical power dynamics in wavelength layer software defined networking," in *Advanced Photonics*, 2015, paper NeT2F-2.
- [8] L. E. Nelson, G. Zhang, N. Padi, C. Skolnick, K. Benson, T. Kaylor, S. Iwamatsu, R. Inderst, F. Marques, D. Fonseca, M. Du, T. Downs, T. Scherer, C. Cole, Y. Zhou, P. Brooks, and A. Schubert, "SDN-Controlled 400 GbE end-to-end service using a CFP8 client over a deployed, commercial flexible ROADM system," in *Optical Fiber Communication Conf.*, 2017, paper Th5A-1.
- [9] T. J. Xia, G. A. Wellbrock, M. Huang, S. Zhang, Y. Huang, D. Chang, S. Burtsev, W. Pelouch, E. Zak, G. Pedro, W. Szeto, and H. Fevrier, "Transmission of 400G PM-16QAM channels over long-haul distance with commercial all-distributed Raman amplification system and aged standard SMF in field," in *Optical Fiber Communications Conf. and Exhibition*, 2017, paper Tu2B-1.
- [10] C. Tian and S. Kinoshita, "Analysis and control of transient dynamics of EDFA pumped by 1480-and 980-nm lasers," *J. Lightwave Technol.*, vol. 21, no. 8, pp. 1728–1734, 2003.
- [11] F. Smyth, D. C. Kilper, S. Chandrasekhar, and L. P. Barry, "Applied constant gain amplification in circulating loop experiments," *J. Lightwave Technol.*, vol. 27, no. 21, pp. 4686–4696, 2009.
- [12] W. Mo, S. Zhu, Y. Li, and D. C. Kilper, "Dual-wavelength source based optical circuit switching and wavelength reconfiguration in multi-hop ROADM systems," *Opt. Express*, vol. 25, no. 22, pp. 27736–27749, 2017.
- [13] I. Kiyoo, J. Kurumida, and S. Namiki, "Experimental investigation of gain offset behavior of feedforward-controlled WDM AGC EDFA under various dynamic wavelength allocations," *IEEE Photon. J.*, vol. 8, no. 1, p. 7901713, 2016.
- [14] W. Mo, S. Zhu, Y. Li, and D. C. Kilper, "EDFA wavelength dependent gain spectrum measurement using weak optical probe sampling," *IEEE Photon. Technol. Lett.*, vol. 30, no. 2, pp. 177–180, 2017.
- [15] L. Barletta, A. Giusti, C. Rottondi, and M. Tornatore, "QoT estimation for unestablished lighpaths using machine learning," in *Optical Fiber Communications Conf. and Exhibition*, 2017, paper Th1J-1.
- [16] S. Yan, F. N. Khan, A. Mavromatis, D. Gkounis, Q. Fan, F. Ntavou, K. Nikolovgenis, F. Meng, E. H. Salas, C. Guo, C. Lu, A. P. T. Lau, R. Nejabati, and D. Simeonidou, "Field trial of machine-learning-assisted and SDN-based optical network planning with network-scale monitoring database," in *European Conf. and Exhibition on Optical Communication*, 2017, paper Th.PDPA.3.

- [17] Y. Huang, C. L. Gutterman, P. Samadi, P. B. Cho, W. Samoud, C. Ware, M. Lourdiane, G. Zussman, and K. Bergman, "Dynamic mitigation of EDFA power excursions with machine learning," *Opt. Express*, vol. 25, no. 3, pp. 2245–2258, 2017.
- [18] Y. Huang, P. B. Cho, P. Samadi, and K. Bergman, "Power excursion mitigation for flexgrid defragmentation with machine learning," *J. Opt. Commun. Netw.*, vol. 10, no. 1, pp. A69–A76, 2018.
- [19] C. L. Gutterman, W. Mo, S. Zhu, Y. Li, D. C. Kilper, and G. Zussman, "Neural network based wavelength assignment in optical switching," in *Workshop on Big Data Analytics and Machine Learning for Data Communication Networks*, ACM, 2017, pp. 37–42.
- [20] J. Schmidhuber, "Deep learning in neural networks: an overview," *Neural Netw.*, vol. 61, pp. 85–117, 2015.
- [21] G. Xavier, A. Bordes, and Y. Bengio, "Deep sparse rectifier neural networks," in *14th Int. Conf. on Artificial Intelligence and Statistics*, 2011.
- [22] R. Hecht-Nielsen, "Theory of the backpropagation neural network," in *Neural Networks*, 1988, pp. 445–448.
- [23] A. Y. Ng, "Feature selection, L 1 vs. L 2 regularization, and rotational invariance," in *21st Int. Conf. on Machine Learning*, ACM, 2004, p. 78.
- [24] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout from higher education: a theoretical synthesis of recent research," *Rev. Educ. Res.*, vol. 45, no. 1, pp. 89–125, 1975.
- [25] F. Melgani and L. Bruzzone, "Classification of hyperspectral remote sensing images with support vector machines," *Trans. Geosci. Remote Sens.*, vol. 42, no. 8, pp. 1778–1790, 2004.
- [26] W. Mo, Y. K. Huang, S. Zhang, E. Ip, D. C. Kilper, Y. Aono, and T. Tajima, "ANN-based transfer learning for QoT prediction in real-time mixed line-rate systems," in *Optical Fiber Communication Conf.*, 2018, paper W4F.3.

Weiyang Mo received his B.S. in Physics from Southeast University, Nanjing, Jiangsu, China in 2011, and his M.S. from the College of Optical Sciences, University of Arizona, Tucson, Arizona in 2014. He is currently a Ph.D. candidate under the advisement of Dr. Daniel Kilper at the College of Optical Sciences, University of Arizona. His research focuses on machine learning, software-defined networking, and advanced coherent DWDM technologies.

Craig L. Gutterman received a B.Sc. in Electrical Engineering from Rutgers University, New Jersey, USA, in 2012, and a M.S. in Electrical Engineering from Columbia University, New York, New York, USA, in 2014, where he is currently pursuing a Ph.D. His interests include mobile and wireless networks, optimization

algorithms, and machine learning. His research focuses on developing algorithms to improve the performance of networked systems. He received the NSF GRFP and the From Data to Solutions NSF IGERT fellowships.

Yao Li received B.S. and Ph.D. degrees in Electronic Engineering from Tsinghua University, Beijing, China, in 2010 and 2016, respectively. She is currently a Post-Doctoral Research Scientist at the College of Optical Sciences, University of Arizona, Tucson, Arizona, USA. She has coauthored more than 20 articles in peer-reviewed conferences and journals. Her current research interests include software-defined networking for optical networks, data center networks, spatial division multiplexing networks, and coherent transmission.

Shengxiang Zhu is a Ph.D. student in the Department of Electrical and Computer Engineering at University of Arizona, Tucson, AZ, USA. He received his bachelor's degree in Optoelectronics in Shenzhen University, Shenzhen, Guangdong, China. His research interests include coherent optical transmission, optical network control, software-defined networks, and photonics.

Gil Zussman received a Ph.D. in Electrical Engineering from the Technion in 2004 and was a postdoctoral associate at MIT in 2004–2007. He is currently an Associate Professor of Electrical Engineering at Columbia University. He is a co-recipient of seven paper awards including the ACM SIGMETRICS'06 Best Paper Award, the 2011 IEEE Communications Society Award for Advances in Communication, and the ACM CoNEXT'16 Best Paper Award. He received the Fulbright Fellowship, the DTRA Young Investigator Award, and the NSF CAREER Award, and was a member of a team that won first place in the 2009 Vodafone Foundation Wireless Innovation Project competition.

Daniel C. Kilper is a research professor in the College of Optical Sciences at the University of Arizona, Tucson. He holds a joint appointment in Electrical and Computer Engineering at the University of Arizona and an adjunct faculty position in Electrical Engineering at Columbia University. He is also founder and CTO of Palo Verde Networks, Inc. He received a Ph.D. in Physics from the University of Michigan in 1996. From 2000 to 2013, he was a member of technical staff at Bell Labs. He is a senior member of IEEE and is an editor for the IEEE Transactions on Green Communications and Networking and a steering committee member for the IEEE Green ICT Initiative. He currently serves as administrative director for the Center for Integrated Access Networks, an NSF Engineering Research Center. His work has been recognized with the Bell Labs President's Gold Medal Award and he served on the Bell Labs Presidents Advisory Council on Research. He holds seven patents and authored five book chapters and more than one hundred thirty peer-reviewed publications.