Learning to Tune Optical WANs:
A Field Deployment of Noise Models in Optical Networks

Bhaskar Kataria[1], Howard Hua[1], Andrea D Amico[2], Bill Owens[3], Rachee Singh[1]
[1]Cornell University [2]NEC Labs [3]NYSERNet

Abstract

Accurately modeling optical signal transmission is critical for optimizing network performance, particularly in large-scale fiber optic networks operated by Internet Service Providers. In this work, we develop a Gaussian Noise model for a New York state ISP's optical backbone. Our model accounts for all major network components, including amplifiers, fiber spans, reconfigurable optical add-drop multiplexers, and transceivers. By accurately predicting end-to-end signal-to-noise ratio, our model provides a foundation for network performance analysis and optimization. Then, we leverage hyperparameter search techniques—commonly used in machine learning—to identify amplifier gain settings that improve signal quality. By treating the model as an opaque box, we systematically search for amplifier configurations that maximize the predicted end-to-end SNR while maintaining practical network constraints. We validate our approach through a field deployment by applying optimized amplifier gain settings in a live ISP network. Our results show a significant improvement in optical signal quality, achieving a 2 dB increase in SNR on a single wavelength [1].

1 Introduction

Modern long-haul connectivity is made possible by optical networks that modulate data on to wavelengths of light [8]. These networks consist of a series of optical components, including transceivers, amplifiers, and switches, that work together to transmit data over long distances using optical wavelengths. The end-to-end quality of these optical signals determines the efficiency and capacity of the optical network. For example, the achievable data rate of a given wavelength — whether 200 or 400 Gbps — is fundamentally determined by the signal quality of the wavelength. In fact, recent work has proposed dynamically improving the data rate of optical wavelengths based on measurements of optical signal quality in fiber, highlighting the importance of achieving optimal signal quality on fiber networks [34].

Complexity of achieving high signal quality operationally. As a result, optical network operators spend significant effort to develop networks that maximize the quality of signal transmission. This involves monitoring signals, debugging equipment configuration and rolling out human crews to fix malfunctioning equipment in a timely manner. However, much of these efforts are manual and ad-hoc. For example, operators often have to rely on their experience to determine the best configuration of amplifiers in a network. This is a difficult task since making changes in one part of the optical path has implications for the entire system, and the effects of these changes are often difficult to predict. Similarly, adding new data-carrying wavelengths to augment network capacity is considered a risky operation that can lead to potential outages [20, 38]. Operators manage this uncertainty by using long maintenance time windows to make incremental changes, which can be costly and inefficient.

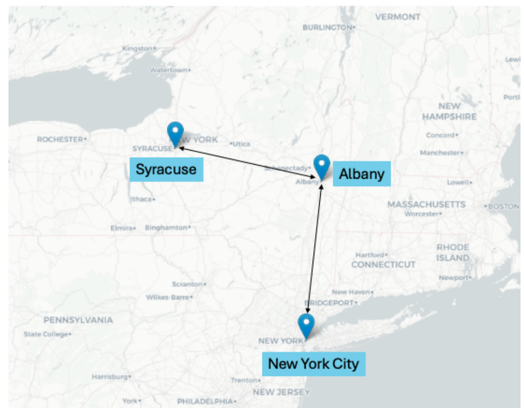

Figure 1: Map of optical paths in the NYSERnet network that we model in this work.

System-level modeling of networks. In recent work, operators have modeled datacenter networks as an end-to-end system to manage operational complexity [21]. Network models abstract intricate details of individual components to approximate the system behavior. Such models can enable operators to anticipate system-wide implications of local configuration changes, providing visibility into network behavior that is otherwise challenging to reason about manually. By systematically evaluating *what-if* scenarios on the model, operators can proactively optimize network configurations, efficiently allocate capacity, and troubleshoot performance issues.

Accurately modeling optical networks. Inspired by modeling efforts in other network types, we develop an analytical model of optical paths in a production-scale optical network operated by a New York state ISP, NYSERnet (§2). Our model uses Gaussian Noise (GN) modeling which has shown promise in accurately predicting the end-to-end signal quality of optical networks in laboratory settings [12]. A GN model simplifies

[1]Our artifacts are open: https://github.com/artifact-release/optical-model

complex physical phenomena — like amplifier-induced noise on optical fiber — by representing the collective impact of these phenomenon as additive Gaussian-distributed noise. However, creating accurate GN models is challenging since it requires detailed knowledge of the network layout, the optical equipment used, and their configurations. Many of these important details are proprietary, making it challenging to build models that accurately reflect real-world network behavior.

Model of production New York ISP. In this paper, we tackle this challenge by developing a high-fidelity GN model of NYSERnet. We gather comprehensive network topology information, operational parameters of optical equipment, and real-world signal quality measurements. To overcome gaps due to missing proprietary information, we approximate equipment using information in openly available vendor datasheets (§3). Finally, we collect end-to-end signal quality measurements from the live network to compare the model's predicted signal quality against measured values and show that the model can accurately predict the end-to-end signal quality (§5).

Field deployment of the model. Using the model as an opaque box, we then apply machine learning-inspired hyperparameter search techniques to optimize the configuration of amplifiers in the network. By systematically exploring the parameter space, we identify configurations that maximize the predicted end-to-end signal-to-noise ratio (SNR) while adhering to practical constraints (§4). We deployed the optimized amplifier configuration in the production network and show that it leads to a measured signal quality increase of approximately 2 dB on targeted wavelengths.

Artifact release. We are releasing the GN model we have developed and our comprehensive dataset — including network topology, equipment configuration and measured signal quality — to the research community. Code and data are available at `https://github.com/artifact-release/optical-model`. By sharing these resources, we aim to support further research in the management and optimization of optical networked systems, ultimately improving reliability and efficiency across large-scale infrastructures.

2 Composition of an optical backbone network

In this section, we describe hardware components of an optical backbone network. The behavior of these hardware components is managed by software configuration programmed by network operators. We will discuss important software settings that impact the end-to-end quality of signal transmision in optical wide-area networks.

Optical wavelengths. Network links between routers in a wide-area network are underpinned by one or more optical paths. Specifically, router ports transmit data bits through optical transceivers which modulate these bits on a wavelengths of light. These wavelengths travel on the underlying optical path until they reach another router port where they are demodulated and converted back to electrical signals.

Optical paths. An optical path is a series of interconnected *segments*, each with its own set of hardware components including (1) optical amplifiers, (2) fiber spans, and (3) reconfigurable add/drop multiplexers (ROADMs). A segment terminates at a ROADMs on each end. The performance of an optical path is influenced by factors like the type of fiber used, the configuration of the amplifiers, and the loss introduced by optical components like ROADMs.

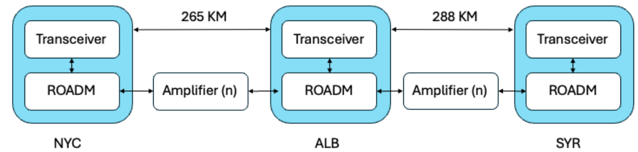

Figure 2: Logical diagram of optical paths in NYSERnet.

Wavelength-division multiplexing. Optical backbone networks rely on wavelength-division multiplexing (WDM), a technique that enables multiple optical signals to be transmitted simultaneously over a single fiber by assigning each signal a unique wavelength (or channel) of light. This approach increases fiber capacity without requiring additional physical infrastructure. In long-haul networks, the C-band (1530 nm to 1565 nm) is commonly used because it offers low propagation loss in optical fibers [32]. The C-band is subdivided into tightly spaced channels, typically 50 GHz apart, allowing tens of distinct data-carrying channels to be multiplexed onto a single fiber. Each wavelength carries its own independent stream of data, and all wavelengths propagate together along the same optical path.

New York State ISP. We model operational optical paths in the optical backbone of a New York state-wide educational and research network called NYSERnet. NYSERnet is a non-profit organization that provides high-speed internet access to educational institutions in New York State. NYSERnet operates a state-wide optical backbone that is composed of multiple optical paths that provide high-speed connectivity between different regions of the state. Figure 1 shows the physical location of the optical paths in NYSERnet that we model in this work. Specifically, we model optical paths between New York City (NYC) and Syracuse, routed through Albany. These paths have two optical segments — NYC to Albany and Albany to Syracuse — each with its own hardware components.

2.1 Optical hardware components

Figure 2 shows the logical diagram of the optical network connecting different components and Figure 3 shows the physical diagram of the optical network in significant detail. We detail the role and configuration of all hardware components in the NYSERnet segments and describe how they impact overall signal quality and system performance.

2.1.1 ROADMs

Reconfigurable Optical Add-Drop Multiplexers (ROADMs) are wavelength-selective optical switches that route optical signals in fiber networks. ROADMs interface with WAN routers using optical transceivers plugged into router ports. Routers transmit data by converting electrical signals to optical

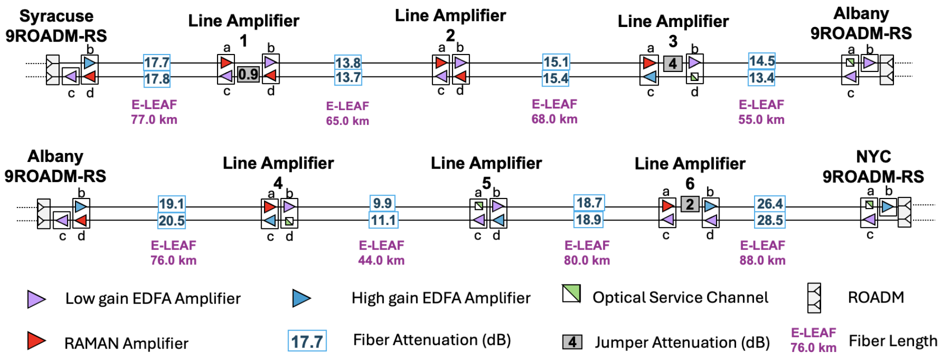

Figure 3: Optical network topology. The network consists of two segments: Syracuse to Albany and Albany to NYC. Each segment is connected by ROADMs, amplifiers, and fiber spans. Fiber attenuation is shown in the blue boxes along each fiber span, and were measured at the time of fiber installation using an optical time-domain reflectometer (OTDR). The jumpers are shown in gray boxes positioned after some Raman amplifiers, and are fixed optical attenuators used to adjust the amplifier gain since the model of Raman ampifier NYSERnet used does not contain a variable optical attenuator (VOA). Fiber distance is shown in the purple text below each fiber span and was provided to us by the network operator.

form via transceivers, and the resulting optical signals are then *added* to the optical network by ROADMs. Conversely, a wavelength is *dropped* at a ROADM node by passing the optical signal to a transceiver, converting it to an electrical signal that is processed by the router. Additionally, ROADMs can *pass* optical signals without modification, allowing them to traverse the network without being added or dropped. Thus, at each network node, ROADMs selectively insert (add), extract (drop), or transparently pass optical signals.

ROADMs introduce between 4 dB and 9.5 dB of signal attenuation known as insertion loss [25], causing a reduction in the optical signal's power level. To compensate for these insertion losses — which weaken optical signals and reduce transmission distances — ROADMs are paired with optical amplifiers such as erbium-doped fiber amplifiers (EDFAs) or Raman amplifiers placed immediately downstream. These amplifiers restore the optical signals to adequate power levels, enabling reliable long-haul fiber transmission.

In our Gaussian Noise model of the NYSERnet network (§3), we represent ROADMs as multiplexers and demultiplexers of optical wavelengths to capture their wavelength routing behavior. We explicitly capture the insertion loss of ROADMs across different wavelengths of light. Accurately modeling these parameters is critical for predicting and optimizing end-to-end network performance.

2.1.2 Optical amplifiers

Optical backbone networks transmit light signals over hundreds to thousands of kilometers of fiber. While propagating along these paths, the signal degrades due to attenuation from scattering and absorption. To maintain signal quality over long distances despite this attenuation, networks deploy optical amplifiers, which boost signal power directly in the optical domain without optical-to-electrical conversion. Amplification works by using a *gain medium* that, when pumped with external laser energy, stimulates the emission of additional photons to amplify the input wavelengths. Amplifiers are essential to overcoming loss from fiber spans, ROADMs *etc.*

Amplification modes. Optical amplifiers operate primarily in two modes: *constant gain* and *constant output power* mode. In constant gain mode, the amplifier provides a fixed amount of amplification regardless of input power of optical signals. Therefore, the amplifier measures the aggregate input and output optical powers across all signals, subtracts them to calculate the instantaneous gain ($\Delta P_{dB} = P_{out} - P_{in}$), and then adjusts internal settings so that this gain stays fixed to the configured value. In constant output power mode, an amplifier measures the aggregate output power of all channels and adjusts internal settings so that the aggregate output power (P_{out}) stays at the configured value. Therefore, in constant power mode the amplifier keeps the absolute output power constant, adjusting gain as the input power changes.

Constant gain vs. constant power. Therefore, network operators can configure the amplifier to operate in either mode but the choice of mode affects the overall performance of the optical path and the quality of transmitted signals. Specifically, in constant gain mode every active wavelength is amplified by roughly the same dB. As a result, across all wavelengths on fiber, any wavelength leaves with the same relative power it had before getting amplified. This ensures that the optical power of each wavelength in the network is within a narrow range of every other wavelength, provided this property held before they are input to the amplifier. However, this also means the total power coming out of an amplifier rises and falls whenever channels are added or dropped on fiber, forcing downstream gear to absorb those changes. In contrast, amplifiers in constant power hold the total output power fixed. If fewer channels are present, the gain per channel becomes higher since the power is held constant. Thus, constant power mode lets the gain per channel drift.

Amplifier types. Given the pros and cons of the two amplification modes and the capability of optical equipment deployed in the network, NYSERnet uses amplifiers in constant gain mode. Optical long-haul networks deploy amplifiers of a few different types. The most common types in practice are Erbium-Doped Fiber Amplifiers (EDFAs) and

Raman amplifiers. NYSERnet uses both EDFAs and Raman amplifiers along optical paths (Figure 3). We describe how they work and their role in the network:

- **Erbium-doped fiber amplifiers (EDFAs):** An EDFA uses a short length of fiber that has been doped with erbium ions. When this fiber is pumped with light, the excited erbium ions transfer energy to passing C-band signals, boosting their power. Since light in the C-band suffers lowest attenuation in fiber, EDFAs have become the workhorse amplifiers for long-haul networks. NYSERnet deploys two EDFA variants, tuned for different fiber span lengths. *High-gain* EDFAs sit on the long, high-loss spans where a large boost per channel is essential. *Low-gain* EDFAs are used on short spans where only modest amplification is needed.

- **Raman amplifiers:** A Raman amplifier boosts optical signals by using the transmission fiber itself as the gain medium. Instead of relying on a separate doped fiber (like an EDFA), a pump laser injects light at a wavelength longer than the signal. As this pump light travels through the fiber, it interacts with the fiber molecules and transfers energy from the pump light to the signal, amplifying the signal as it propagates along the fiber. Raman amplifiers provide high gain values while adding lesser noise to optical signals compared to EDFAs. However, Raman amplification requires careful tuning of pump power and wavelength, making it more operationally challenging and sensitive to fiber conditions [2, 5].

  NYSERnet operates Raman amplifiers at their maximum safe gain to take advantage of their lower noise compared to EDFAs. However, maximizing Raman amplifier gain can result in excessively high optical power levels entering the subsequent EDFA, potentially causing signal distortion (see the cascaded amplifiers shown in Figure 3). To manage these power levels precisely, we insert fixed optical attenuators — commonly known as jumpers — between the Raman amplifier and the following EDFA. These jumpers allow us to finely adjust the cumulative amplification. For instance, if a fiber span requires a total gain of 22 dB, and the Raman amplifier provides 15 dB while the downstream EDFA has a minimum gain setting of 11 dB, inserting a 4 dB jumper ensures the overall amplification precisely matches the span requirements without introducing distortion. In practice, we use jumpers as a way to captures both unintentional losses introduced during deployment—such as those due to imperfect splicing—and intentional losses introduced to ensure compatibility with amplifier gain requirements.

Operators configure various parameters of optical amplifiers to ensure that the optical signals are amplified to appropriate levels in the network. We first describe these parameters and later optimize their settings to improve the performance of optical paths in NYSERnet (§3):

- **Gain:** Amplifier gain quantifies the increase in signal power, defined as the ratio of output to input power and typically expressed in decibels (dB):

  $$\text{Gain (dB)} = 10 \cdot \log_{10}\left(\frac{P_{\text{out}}}{P_{\text{in}}}\right) \tag{1}$$

  Gain is controlled by adjusting the power of the laser pump inside the amplifier. For the NYSERnet deployment, we tune gain values between 6 dB and 26.5 dB, depending on fiber loss and system requirements. Improper gain settings can degrade signal quality: excessive gain introduces noise and nonlinear impairments, reducing signal quality, while insufficient gain results in weak signals and increased bit-error rates. Therefore, operators carefully balance the two effects to find the optimal gain setting for each amplifier in the network.

- **Tilt:** Tilt is the configured variation in amplifier gain across different optical wavelengths, typically expressed in decibels (dB). It measures the difference in amplification between the lowest and highest wavelength channels and is crucial for ensuring consistent gain across the entire optical spectrum. In NYSERnet, we typically configure tilt values ranging from -1.5 dB to 0 dB. A negative tilt indicates that lower-wavelength channels receive higher amplification compared to higher-wavelength channels. For instance, a tilt of -1.5 dB means that the gain at the lowest wavelength channel is 1.5 dB higher than at the highest wavelength channel. This negative tilt compensates for the gain characteristics of EDFAs, which amplify certain wavelength regions more than others.

- **Integrated variable optical attenuator (VOA)**: A VOA controls the amplifier's output power by attenuating the amplified signal to manage the signal power levels. In NYSERnet, we use fixed VOAs with attenuation values ranging from 0 dB to 9 dB. VOAs ensure that the output power is within the acceptable range for downstream components. The attenuation applied by a VOA is adjusted based on the specific requirements of the network and the characteristics of the optical signal.

2.1.3 Optical fiber

Optical fibers are the physical medium through which optical signals propagate in fiber-optic networks. Their performance is characterized by three key parameters: attenuation, dispersion, and effective area.

Attenuation refers to the gradual loss of signal power as the optical signal travels along the fiber, measured in decibels per kilometer (dB/km). Fiber attenuation directly limits the transmission distance by weakening signals over long spans. We calculate the fiber attenuation using the following relationship:

where α represents the fiber's loss coefficient in dB/km, and L denotes the fiber length in kilometers.

$$\text{Loss (dB)} = \alpha \cdot L \quad (2)$$

Dispersion is the phenomenon by which optical pulses spread out as they travel through the fiber. This spreading limits the maximum achievable transmission distance and data rate.

Effective area describes the cross-sectional area within the fiber through which the optical signal propagates. The effective area influences nonlinear optical effects that can degrade signal quality, particularly at high power levels. Larger effective areas can reduce these nonlinear impairments, improving overall system performance.

In NYSERnet, we use Corning E-LEAF fiber [17], a fiber type known for its optimized balance of attenuation, dispersion, and effective area. Throughout this paper, we report fiber attenuation values for each span based on measurements obtained using an industry-standard instrument for signal loss characterization: the optical time-domain reflectometer (OTDR).

2.1.4 Optical transceivers

Optical transceivers convert electrical signals from routers to optical signals and vice versa. They transmit and receive data over optical fibers. Optical transceivers are characterized by their data rate, wavelength, and modulation format. In NYSERnet, we use 200 Gbps transceivers with 16-state dual-polarization quadrature amplitude modulation (16DP-QAM) modulation format. This modulation scheme encodes four bits per symbol per polarization by varying both amplitude and phase of optical signals across two orthogonal polarization states. Although 16DP-QAM allows higher data throughput within limited spectral bandwidth [20], it requires higher signal-to-noise ratios due to increased susceptibility to noise and impairments [26]. As we show in §5, signal quality in NYSERnet can support 16DP-QAM modulation format.

2.1.5 Optical service channel

Optical Service Channels (OSCs) are channels dedicated for monitoring and managing the optical network. They are used for out-of-band signaling and control, allowing network operators to perform maintenance and diagnostics without interfering with the main data channels. OSCs are implemented as low-rate channels that can carry management information, alarms, and performance monitoring data. In NYSERnet, OSCs have a fixed loss of 1.2 dB and are used to monitor the performance of the optical amplifiers and other network components. OSCs are not used for data transmission but play a crucial role in ensuring the reliability and performance of the optical network.

2.2 Optical network topology

Figure 3 illustrates optical paths in the NYSERnet, comprising of two primary segments: Syracuse to Albany and Albany to New York City (NYC). The network utilizes Corning E-LEAF fiber throughout. This fiber type was chosen for its low attenuation, making it suitable for long-haul transmissions. Each network segment is composed of fiber spans, amplifiers, and ROADMs, collectively enabling high-capacity optical transmission over long distances.

Intermediate amplifier sites are spaced approximately 44–88 km apart along each segment. Each site hosts four amplifiers — two in each direction of data transmission — with combinations of Raman amplifiers and EDFAs. A single amplifier at a site is sometimes inadequate to fully compensate for the accumulated fiber span and ROADM insertion losses. To address this, NYSERnet deploys cascaded amplifiers at these sites: a Raman amplifier first provides gain with minimal noise, followed by an EDFA to amplify the residual signal power to the required output level.

Amplifier placement also accounts for asymmetries in the optical path. For example, in the direction receiving signals from NYC, the optical path requires two amplifiers in series to overcome both span loss and higher ROADM insertion losses. Conversely, transmission towards NYC accumulates lower losses due to the hardware differences that result in varying loss profiles, thus requiring only one amplifier. Additionally, some amplifier sites use Optical Service Channels in configurations where span losses are sufficiently low for a single amplifier. These OSCs provide out-of-band management channels, enabling continuous monitoring of amplifier performance, fiber conditions, and ROADM status, enhancing operational visibility and simplifying maintenance.

3 Optical network model

Optical backbone networks are a complex combination of various optical equipment, each of which has its own set of configuration parameters and noise contributions (§2). Optical equipment contributes to signal degradation by adding noise and impairments to the signal as it traverses the fiber. Therefore, managing these networks is a challenging task for network operators. For instance, making changes in one part of the optical path has implications for the entire system, and the effects of these changes are often difficult to predict for network operators. Similarly, adding new data-carrying wavelengths to augment network capacity is considered a risky operation that can lead to potential outages [20, 38]. Operators manage this uncertainty by using long maintenance time windows to make changes, which can be costly and inefficient.

Need for analytical models for optical networks. One way to tackle the operational complexity of optical backbone networks is to develop analytical models that faithfully mimic the network's behavior and can be used to predict the effects of configuration changes. These models can help network operators understand how different components interact and how they affect the overall performance of the network. By simulating various *what-if* scenarios, operators can make informed decisions about network configurations changes, capacity planning, and troubleshooting. In this section, we discuss how we model optical paths in the NYSERnet network to accurately capture the end-to-end behavior of the network. We use the

Label	Model	Type	Gain	Tilt	VOA
SYR a	9ROADM-RS	ROADM			
SYR b	AMP-S20H-C15	EDFA	15	0	2.5
SYR c	EDFA-S20L	EDFA	11.8	-0.5	2.5
SYR d	AMP-S20H-C15	RAMAN	11	-1	
LA 1 a	AMP-S20L-C15	RAMAN	12	-1.5	
LA 1 b	AMP-S20L-C15	EDFA	8.5	0	5.5
LA 1 c	AMP-S20L-C15	EDFA	8	0	2.5
LA 1 d	AMP-S20-C15	RAMAN	12	-1.5	
LA 2 a	AMP-S20L-C15	RAMAN	12	-1	
LA 2 b	AMP-S20L-C15	EDFA	7	0	4.5
LA 2 c	AMP-S20L-C15	EDFA	7	0	5.5
LA 2 d	AMP-S20L-C15	RAMAN	12	-1	
LA 3 a	AMP-S20H-C15	RAMAN	15	-1	
LA 3 b	EDFA-S20L	EDFA	6	0	6
LA 3 c	AMP-S20H-C15	EDFA	16.5	0	4.5
LA 3 d	EDFA-S20L	OSC			
ALB S a	EDFA-S20L	OSC			
ALB S b	AMP-S20L-C15	EDFA	19.5	-1.5	0
ALB S c	EDFA-S20L	EDFA	15	-1	2.5
ALB S d	9ROADM-RS	ROADM			
ALB N a	9ROADM-RS	ROADM			
ALB N b	AMP-S20H-C15	EDFA	15	0	3.5
ALB N c	EDFA-S20L	EDFA	11.5	-0.5	0.8
ALB N d	AMP-S20H-C15	RAMAN	15	-1	
LA 4 a	AMP-S20H-C15	RAMAN	15	-1	
LA 4 b	EDFA-S20L	EDFA	12	-1	9
LA 4 c	AMP-S20H-C15	EDFA	14	0	3.5
LA 4 d	EDFA-S20L	OSC			
LA 5 a	EDFA-S20L	OSC			
LA 5 b	EDFA-S20L	EDFA	14.5	-0.5	3.5
LA 5 c	EDFA-S20L	EDFA	11.5	-0.5	5
LA 5 d	EDFA-S20L	EDFA	12	0	0
LA 6 a	AMP-S20L-C15	RAMAN	15	-1.5	
LA 6 b	EDFA-S20L	EDFA	11	-1	1
LA 6 c	AMP-S20L-C15	EDFA	11	0	3.5
LA 6 d	EDFA-S20L	EDFA	16	0	0
NYC a	EDFA-S20L	OSC			
NYC b	EDFA-S20H	EDFA	26.5	-1	0
NYC c	EDFA-S20L	EDFA	15	-1.5	0
NYC d	9ROADM-RS	ROADM			

Table 1: Network Parameters. All blank entries indicate that the parameter is not applicable. For example, the model of Raman amplifier we used does not contain a VOA at its output. Additionally, OSCs and ROADMs do not have gain, tilt, nor VOA parameters. The gain and tilt values are in dB. The VOA values are in dB and indicate the amount of attenuation applied to the signal.

model to improve network performance in §4 and evaluate our model in §5 with the field deployment.

3.1 Gaussian noise modeling for optical networks

Recent work has proposed Gaussian noise (GN) modeling as an analytical technique for modeling the effects of noise in optical networks [6, 7]. The GN model simplifies the complex interactions between different noise sources by treating them as independent, additive Gaussian-distributed processes. This allows for analytical predictions of the Optical Signal-to-Noise Ratio (OSNR) of wavelengths in an optical network. The GN model has been shown to produce reliable estimates of OSNR across a wide range of optical network configurations [29]. Moreover, the computational simplicity makes GN modeling particularly suitable for operational deployment.

Accurate Gaussian noise modeling. However, to effectively use GN modeling for operational purposes, it is critical to create an accurate representation of the optical network. Each network component — including EDFAs, Raman amplifiers, ROADMs, and optical fibers — has distinct noise contributions that must be carefully parameterized in the model. By precisely capturing and aggregating these component-specific impairments across each span of the optical path, the GN model can yield reliable, end-to-end predictions of the OSNR. Without detailed representation of the noise characteristics of all components, the utility of the GN model will be limited, potentially leading to unreliable predictions of network performance and incorrect decision-making based on them. We tackle the challenge of developing an accurate GN model for the NYSERnet optical network.

Predicting end-to-end SNR using GN models. A GN model predicts the end-to-end SNR of a wavelength in an optical network by aggregating the contributions of various noise sources across the entire path. Such a model treats each noise source as an independent Gaussian process [13], allowing for analytical predictions of the overall SNR. There are two main types of noise sources in optical paths:

- **Amplified spontaneous emission (ASE) noise:** ASE noise arises from the spontaneous emission of photons within the gain medium of optical amplifiers. Both types of amplifiers used in our network — EDFAs and Raman amplifiers — are affected by this phenomenon. As the optical signal traverses the network, ASE noise is not only introduced but also amplified at every stage, resulting in a cumulative reduction in the SNR [2].

- **Nonlinear interference (NLI) noise:** NLI noise arises from nonlinear effects in the fiber, particularly when signals are transmitted at high launch powers. These *nonlinearities* occur because the refractive index of the fiber changes slightly with the intensity of the light passing through it [2], leading to unwanted interference and modulation of the signal. NLI becomes especially problematic in DWDM systems, where many closely spaced channels share the same fiber [10]. The extent of NLI depends on several factors, including fiber dispersion, channel spacing, and the number of spans in the link. Because it is highly sensitive to both launch power and channel spacing, NLI is a key limitation in the design and performance of long-haul optical systems [11].

A GN model of the network is parameterized by the network topology, fiber parameters, amplifier properties, and optical channel configurations. Using these, it calculates ASE noise introduced by amplifiers using their gain and noise figures. It can estimate nonlinear interference noise caused by fiber nonlinearities, computed based on optical channel power and fiber characteristics. By combining the ASE and nonlinear noise along each fiber span and amplifier in the path, GNPy computes the total noise power and derives the resulting channel SNR.

3.2 Modeling with GNPy

We use the open-source GN modeling tool, GNPy [14], to simulate our optical network. GNPy allows us to flexibly model essential physical-layer components, like amplifiers, fiber spans, splices, ROADMs, and transceivers [13]. Each component in our model is parameterized either from configuration settings used in the NYSERnet deployment or from vendor datasheets of the deployed equipment. In the following, we discuss the modeled components and their associated parameters in our GN model of the NYSERnet network.

3.2.1 Modeling amplifiers

All amplifiers in our deployment operate in gain mode (§2). In our GN model, we instantiate each amplifier shown in Figure 3 using gain parameters directly from the configured amplifier values in the NYSERnet network (Table 1). Additionally, for each amplifier type, we specify the supported gain range, maximum output power, and corresponding noise figure range, derived from vendor datasheets. While we enable the output VOA feature within our model, we set its attenuation to 0 dB during optimization to limit the search space, as detailed in §4.
Modeling Raman amplifiers: Raman amplifiers pose a modeling challenge compared to EDFAs because their key internal parameters, like pump wavelengths and power levels, are vendor-proprietary and thus unavailable to us. To address this limitation, we adopt a first-order approximation method supported by GNPy [14] where we model Raman amplifiers as modified EDFAs [1, 18]. Using publicly available datasheets, we found that the Raman amplifiers deployed in the NYSERnet network achieve a noise figure below -1dB at maximum gain, significantly lower than the typical EDFA noise figure of 5.8 to 8dB. Given this information, we approximate the Raman amplifier performance by using a generic EDFA model configured with an NF of -1dB. This approach avoids the need to reverse-engineer pump parameters, since amplifier gain and tilt values were directly available (Table 1), allowing us to approximate their performance without detailed characterization of the pump laser.

3.2.2 Modeling fiber spans

We represent each span of Corning LEAF fiber in NYSERnet using the standard single-mode fiber (SSMF) type within GNPy, since it is the most widely deployed fiber type in long-haul optical networks [4]. For each fiber span, we explicitly set the fiber component in GNPy to SSMF and use the true physical span lengths, indicated in purple text in Figure 3. Since fiber loss can vary significantly due to aging, splices, and environmental conditions [16], we independently measure each span's total loss using a diagnostic instrument called the Optical Time-Domain Reflectometer (OTDR) and incorporate these measured values directly into our model. Additionally, we rely on Corning's E-LEAF datasheets [17] for static fiber parameters like internal reflectance, chromatic dispersion, effective cross-sectional area, and polarization mode dispersion (PMD). These parameters influence nonlinear and dispersion characteristics of the fiber and specifying them in the GN model enables accurate modeling of signal degradation for each span.

3.2.3 Modeling transceivers

Accurately modeling transceivers in an optical network requires a clear understanding of both the transmit (Tx) and receive (Rx) components, as they exhibit distinct physical behaviors. The Tx transceiver involves understanding the characteristics of the output laser and is typically modeled using vendor-provided OSNR values and other publicly available parameters. In contrast, the Rx transceiver requires understanding the physical characteristics of the photodiode, which detects incoming light and converts it into an electrical signal. To address these differences, we model Tx and Rx transceivers separately.
Modeling Tx transceivers: We model the Tx transceiver using GNPy's transceiver modeling framework. For each channel, we set the minimum and maximum frequencies to the same value, effectively representing a fixed central frequency for the deployed network. Key transceiver parameters are configured based on datasheet specifications. These include the channel slot width (which defines the spectral allocation in the WDM system), baud rate (symbol transmission rate), roll-off factor (used for spectral shaping), and transmitter OSNR. To simulate idealized conditions, we set ΔPDB—which controls inter-channel power deviation—to zero, ensuring uniform launch power across all channels. These transceiver-level parameters are critical for accurate per-channel SNR estimation and serve as the input for calculating impairments such as nonlinear interference and ASE accumulation along the path.
Modeling Rx transceivers: Similar to the Tx transceiver, we model the receive-side transceiver using GNPy's built-in transceiver framework. However, accurately capturing Rx behavior requires understanding how noise and signal quality vary with input optical power. Prior work [22] has shown that the electrical SNR at the receiver is highly sensitive to received optical power levels. Unlike the Tx side—where vendors typically specify OSNR—detailed noise sensitivity metrics for the Rx side are rarely disclosed due to their proprietary nature. This lack of transparency complicates precise modeling and necessitates empirical adjustments. To address this, we introduce a power corrected value that aligns GNPy's OSNR estimates with the OSNR observed in practice. We describe the methodology for deriving the power corrected value in §3.3.

3.2.4 Modeling other components

Modeling ROADMs: We model ROADMs as passive components in GNPy, which means they do not introduce any additional gain to the signal. We retain GNPy's default ROADM configuration, as it sufficiently represents the behavior of the ROADMs deployed in our network. Specifically, we maintained the default wavelength add/drop OSNR of 38 dB, which indicates that a majority of the signal (almost 99 percent) passes through and is not filtered out. This was because the datasheet for the deployed ROADMs did not provide specific OSNR penalty values, and from literature [30] the OSNR add/drop penalty is almost 0 for between 0 to 3 filters. Since there is 1 filter at each ROADM in our network, and there are 3 ROADMs total, then we estimate that there are only 3 filters total along the path, and the total OSNR drop penalty should be approximately 0. Therefore, the default value in GNPy of 38 dB seems to be a reasonable approximation for typical ROADM performance.

Modeling optical service channel (OSC): We model the Optical Service Channel (OSC) as a static insertion loss applied at the end of the fiber span it traverses. Unlike losses introduced at the beginning of a span, losses at the end do not contribute additional nonlinearities. This makes modeling the OSC as an end-of-span insertion loss well-suited for capturing the static attenuation it experiences.

Modeling optical jumpers: We model optical jumpers as a GNPy "fused" component. Specifically, GNPy handles the fused component loss as a static parameter, which means that its attenuation is independent of signal frequency and amplitude. As a result, the attenuation factor is constant across all channels and does not vary with the input power levels. This accurately models the optical jumpers used in our network, which are designed to have a fixed loss characteristic.

3.2.5 Modeling end-to-end optical routes

GNPy lacks native support for simulating an end-to-end optical route, involving wavelength add/drop at intermediate nodes. Due to this, we separately modeled individual optical segments of the route. To obtain the end-to-end SNR (SNR_{E2E}), we implement a post-processing script to consolidate segment-wise SNR using inverse summation, where $\frac{1}{SNR_{E2E}} = \sum \frac{1}{SNR_{segment}}$.

The optical network carries five wavelengths across its spans, but NYSERnet logs performance metrics only for three of them: 193.35 THz, 193.5 THz, and 193.75 THz. This logging limitation is inherent to the deployed monitoring system. Among these, the 193.35 THz channel is a point-to-point connection that originates in Albany and terminates in NYC. In contrast, the 193.5 THz and 193.75 THz channels originate in Syracuse, pass through Albany without optical-electrical-optical (O-E-O) conversion, and terminate in NYC. These are referred to as pass-through channels because they are not dropped at the intermediate ROADM in Albany; they are simply forwarded along to the next span. We modeled all passthrough waves as point-to-point links across segments and computed the overall SNR as if the wave simply traversed the intermediate segments.

The physical path from NYC to Syracuse spans 553 km of Corning E-LEAF fiber and includes 16 amplifiers, 3 ROADMs, 2 OSCs, and six intermediate line amplifier sites. On the other hand, the reverse path—from Syracuse to NYC—includes 15 amplifiers, 3 ROADMs, and 2 OSCs. This asymmetry in noise and amplification highlights the need for separate modeling of each direction and channel.

3.3 Power correction for GSNR-to-SNR comparison

Receiver noise is a type of noise that originates locally at the signal receiver and arises from thermal noise, quantization errors, and other electronic impairments. It contributes to the overall degradation of signal quality in optical systems. GNPy's Generalized SNR (GSNR) calculation accounts only for ASE and NLI impairments, excluding receiver noise since it is independent of the transmission path characteristics.

This omission results in systematic overestimation of signal quality by GNPy when compared to real-world measurements from the modeled network, since receiver noise can degrade the observed SNR by about 1-2 dB depending on hardware specifications. To address this discrepancy, we introduce a calibrated receiver noise offset into our modeling pipeline. Incorporating this offset allows us to predict the real-world SNR given GNPy's GSNR estimates across diverse network scenarios. We discuss how to derive this offset in this section.

GSNR provides a vendor-agnostic useful optical signal quality estimate, practical deployments rely on transceivers measuring and reporting the electrical Signal-to-Noise Ratio (eSNR), which reflects the total degradation—including both optical and receiver-induced noise. In contrast, GSNR specifically quantifies optical impairments and is computed as the harmonic sum of ASE and NLI contributions accumulated across all fiber spans along a path as $GSNR = \frac{1}{\sum\left(\frac{1}{ASE} + \frac{1}{NLI}\right)}$.

Specifically, the transceivers in our network report Pre-Forward Error Correction Bit Error Rate (Pre-FEC BER) values, which is directly related to eSNR. While some devices provide eSNR via APIs, the reported values are limited to a granularity of 0.1 dB. In contrast, computing eSNR from Pre-FEC BER yields higher resolution and consistency. We provide the equation for converting Pre-FEC BER to eSNR in §3.4.

Converting between OSNR (as reported by tools like GNPy's GSNR) and eSNR (as measured by transceivers) would require a correction factor. Prior work [22] has shown that this correction factor is input-power dependent. Specifically, for a fixed OSNR, increasing the input power also leads to a modest increase in eSNR. Therefore, to avoid introducing confounding effects when evaluating the impact of optimizations, it is important to maintain consistent input power levels. To derive the appropriate correction factor, we selected three channels (193.35, 193.5, and 193.75 THz) along the `Syr → Alb → NYC` path. For each, we measured the eSNR at the receiver and compared it against the segment-level GSNR values predicted by GNPy using the following relationship:

$$\frac{1}{eSNR_{E2E}} = \frac{1}{\text{Power Correction Factor}} + \sum \frac{1}{GSNR_{segment}} \tag{3}$$

We estimated the power correction factor to be approximately $1/135$ in linear units, or 21.3 dB which is similar to the what has been found in other works [9] but it is sensitive to vendor-specific transceiver characteristics and deployment scenarios. Notably, the input power levels before and after optimization remained within 1 dBm across all frequencies, hence this correction factor remained stable throughout our analysis.

We applied this correction to predict eSNR from GNPy's GSNR in the reverse direction (`NYC → Alb → Syr`) and evaluated both pre- and post-optimization signal quality. Since input power remained nearly unchanged, observed improvements in eSNR can be attributed to better optical path performance (i.e., improved OSNR) resulting from amplifier gain optimization.

3.4 Evaluating accuracy of the GN model

We compare the predicted end-to-end signal quality with the measured metrics in the live network. The transceivers used in NYSERnet measure hardware performance, including the Pre-FEC BER. Thus, we use the measured BER to calculate the received eSNR for comparison with the estimated eSNR from the GN model. In this network, all wavelengths operate at a 200 Gbps data rate and the transceivers use 16-state dual-polarization quadrature amplitude modulation (16DP-QAM). Given the modulation format, the eSNR can be derived as a function of the Pre-FEC BER using the following [22]:

$$\text{SNR} = 10\left[\text{erfc}^{-1}\left(\frac{8}{3}\cdot\text{Pre-FEC BER}\right)\right]^2 \tag{4}$$

Here, $\text{erfc}^{-1}()$ represents the complementary error function. We show that the predicted SNR closely resembles the model estimated SNR in the evaluation (§5).

4 Optimizing quality of transmission

A key purpose of developing an accurate model of a live optical network is the ability to conduct *what-if* analyses to optimize real-world performance. To achieve this, we apply hyperparameter search techniques—commonly used in machine learning [35]—to determine optimal amplifier gain settings. Specifically, we use `Optuna` [3, 27] to search for gain configurations that maximize end-to-end signal quality, treating the GN model as a opaque box. At each iteration, `Optuna` samples a candidate gain configuration across all amplifiers, and these parameters are passed to GNPy to simulate the resulting end-to-end GSNR.

Motivating optimization through amplifier gain sweeps: To motivate the reason why optimization of the network is feasible even at scale (with numerous amplifiers, each with their own gain target), we highlight some single-amplifier gain sweeps shown in Figures 4 and 5. The figure shows the GNPy-predicted GSNR as a function of the gain setting for each amplifier. Each curve was generated by sweeping a single amplifier from its minimum possible gain setting to its maximum possible gain setting as specified in vendor datasheets [25], while leaving all other amplifiers' gain at their in-network value. We observe that for many of the EDFA amplifiers, the GSNR is not monotonic with respect to the gain setting. This is due to the fact that the noise figure of the amplifier is not constant

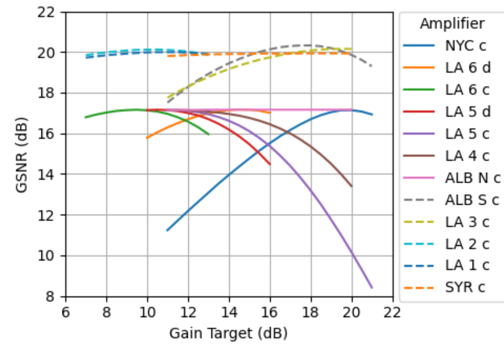

Figure 4: Observed GNPy GSNR for 193.5 THz pass-through channel using EDFA amplifiers as a function of gain setting. Solid lines indicate that the amplifier is along the NYC→Albany path, while dashed lines indicate that the amplifier is along the Albany→Syracuse path. The GSNR is not monotonic with respect to the gain setting, indicating that the noise contribution of the amplifier is not constant across all gain settings.

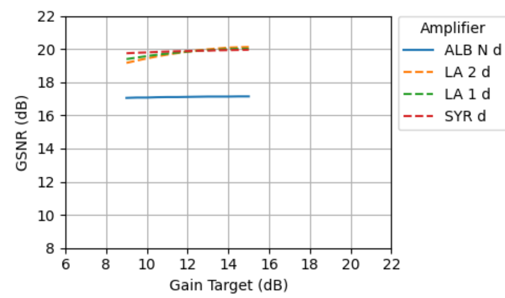

Figure 5: Observed GNPy GSNR for 193.5 THz pass-through channel using Raman amplifiers as a function of gain setting. Solid lines indicate that the amplifier is along the NYC→Albany path, while dashed lines indicate that the amplifier is along the Albany→Syracuse path. The intuition behind the relative flatness is that Raman amplifiers introduce very little noise across their entire gain operating range compared to EDFA amplifiers, thus their impact to the overall GSNR with respect to gain is very small.

across all gain settings, and thus the noise contribution of the amplifier is also not constant. The non-monotonicity is a key reason why tuning the gain of amplifiers in a network can lead to significant improvements in end-to-end SNR. Additionally, the concavity of each curve suggests that the optimization process can find the global GSNR maximum given the maximum and minimum gain settings for each amplifier.

Scope of in-network gain optimization: Before conducting any optimization steps, we first identified potential configuration change windows with the network operator's maintenance schedules. Due to the operator's requirement for scheduling maintenance at least two weeks in advance, and due to the operator disallowing multiple spans from being concurrently scheduled for maintenance, we could only optimize the gain targets for the NYC→Albany span. Thus for the pass-through channels, any improvement to SNR was caused solely by tuning the gain targets of the amplifiers in NYC→Albany span.

However, we show in Figures 4 and 5 that tuning amplifiers in the Albany→Syracuse span (the dashed lines) has significant positive effect (2 dB) on the GNPy-predicted end-to-end GSNR for the 193.5 THz pass-through channel, even when the amplifiers in the NYC→Albany span are left untuned. Future work should explore the potential of tuning amplifiers using the same procedure we detail for the Albany→Syracuse span.

Modeling constraints and simplifications: The amplifier gain optimization is subject to several modeling constraints and simplifications that shape the search space. Each amplifier's gain is bounded by its allowable range, again as specified in vendor datasheets [25], and discretized into 0.5 dB increments (e.g., 10, 10.5, 11) to match the real-world amplifier's configuration granularity. For the ROADM devices, we keep the total input signal power under the maximum 10 dBm as allowed by the datasheet. We constrain the optimization such that if GNPy predicts that the selected gain targets will result in a total input power exceeding 10 dBm, then the trial is void.

To reduce modeling complexity, we set all Variable Optical Attenuators (VOAs) in the network to zero. This simplificatio n comes with the drawback that at low gain, EDFA amplifiers suffer from increased noise figures compared to high-gain operation. This can be problematic, since EDFA amplifiers that operate at high gain and then have their signal attenuated by a VOA will exhibit a lower noise figure than the same EDFA amplifier operating at low gain without subsequent attenuation [28]. We continued to pursue this simplification, however, since it reduced the number of tunable parameters from $O(n^2)$ to $O(n)$. We acknowledge that optimizing the VOA parameter could lead to a more optimal solution.

To further reduce modeling complexity, we fix all Raman amplifiers to operate at their maximum gain. Due to their inherently lower noise figures compared to EDFAs, maximizing Raman gain improves the end-to-end predicted GSNR, as achieving similar amplification with EDFAs would add significantly more noise. Within this constrained space, we keep the objective focused on identifying the gain profile that yields the highest output GSNR and reduces running time to few minutes.

Limitations of built-in optimizer Unfortunately, the gain target optimizer built into GNPy could not be used in our scenario, since it occasionally outputs gain targets which were beyond the capabilities of each amplifier. As an example, one of the low-gain amplifiers used in the network was limited to a gain factor between 7 and 11 dB. Yet, the built-in optimizer in GNPy reported utilizing a gain of 4 dB for this amplifier, which was impossible. Furthermore the default optimizer also has no way to specify the maximum input power for ROADMs. This is a critical consideration, since exceeding the maximum input power can damage the ROADM optics.

Optimization performance and deployment: We ran the tuning process for nearly ten thousand iterations, evaluating different amplifier gain profiles and observing corresponding GSNR values, as shown in Figure 6. We observed consistent improvements in the highest observed GSNR during early iterations, but the gains plateaued soon. We discovered the max GNPy-predicted GSNR was **17.24 dB at a center frequency of 193.5 THz**. This GSNR value is higher than the maximum GSNR (17.16 dB) the baseline finds by sweeping the gain parameter of a single amplifier at a time, as shown in Figures 4 and 5. Thus, our optimization process maximizes the GSNR by varying the parameters across all amplifiers in our experiments. Although the GNPy-derived final total input power was within the 10 dBm limit for the ROADM, the network operator and maintenance team was unsure whether the power was still too high, especially considering the age of the ROADM components. Thus, we settled on a configuration which returned a GNPy-predicted 17.14 dB final GSNR for the 193.5 THz channel. This is the configuration we applied to the live network, where SNR improvements were directly measured. We did not observe any cross channel interference in the three logged channels under the changed configuration. In §5, we present empirical results that demonstrate the effectiveness of approach and configuration.

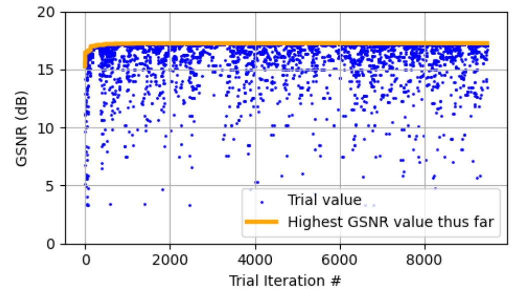

Figure 6: Observed GNPy GSNR for each trial iteration. The blue dots indicates the GSNR for the trial iteration, while the orange line represents the highest GSNR across all trials.

5 Evaluation

Table 2 and Table 3 compare the signal quality estimated by the GN model with measured signal quality along the two directions of the Syracuse → NYC optical route. Values in the Model Estimate column are the estimated GSNR values from the model. Values in the Power-corrected column are the SNR values obtained by accounting for the input power-dependent transceiver noise into the GSNR, resulting in a combined SNR. Values in the Directly Measured column are the SNR values measured in the network. We note that the photodetectors built into transceivers are only precise to ±0.5 dB.

To quantify accuracy, we include two error columns: `Power-corrected Error` and `Measured Error`.

The errors are computed as:

$$\text{Power-corrected Error (\%)} = \frac{\text{Model Estimate} - \text{Power-corrected}}{\text{Power-corrected}} \times 100 \quad (5)$$

$$\text{Measured Error (\%)} = \frac{\text{Model Estimate} - \text{Directly Measured}}{\text{Directly Measured}} \times 100 \quad (6)$$

Both the power-corrected and measured errors remain below 11% across all wavelengths and in both directions, with power-corrected errors consistently lower than the measured ones. This highlights the importance of incorporating the receiver noise into the model. The predicted power-corrected signal quality stays within 1.5 dB and within 10% relative error of the measured values in the network.

5.1 Deployment results

After validating our GN model, we used it to optimize amplifier gain settings. As described earlier in §4, we treat the GN model as a opaque box and apply hyperparameter tuning to identify gain configurations that maximize end-to-end OSNR. The hyperparameter-tuned model reports promising improvements in SNR, especially in channels 193.5 THz and 193.75 THz which propagate over longer distances and are amplified along the line more frequently (see Model Optimized values in Table 5). Motivated from the optimized results, we applied the computed amplifier gain targets to the three wavelengths in the live network. To do this, the network operator opened a maintenance window. During this maintenance window, we applied amplifier gain changes in NYC→Albany→Syracuse direction in increments of 0.5 dB to ensure that the changes do not cause an outage. For each increment, we monitored the SNR for one hour. In case of any performance degradation, we planned to revert back. After applying the proposed changes, SNR measurements showed an overall improvement of 4.2 dB across all three wavelengths, with individual wavelength improvements reaching up to 1.8 dB. Using OSNR, we estimate the improved network capacity using the Shannon-Hartley theorem [31], showing an increase of 10.4% for the link with 1.8 dB increment and 10.3% on average across all links. Follow-up measurements taken after the maintenance window confirmed that the improvements remained stable, demonstrating their robustness over time. Specifically, the 193.35 THz channel saw a stable improvement of 1.8 dB, as shown in Figure 7a. Similarly, the 193.5 THz and 193.75 THz channels, which bypass Albany, exhibited stable improvements of 1.34 dB and 1.33 dB, respectively, as illustrated in Figure 7b. Since the amplifier updates were applied only in the NYC→Albany→Syracuse direction—and the signal originates in NYC—the SNR at the NYC location remains unchanged, as shown in Figure 7c.

Table 4 shows the improvement in signal quality resulting from model optimization, comparing both directly measured and power-corrected SNR values. To quantify these improvements, we compute percentage gains using the following:

$$\begin{array}{c}\text{Power-corrected}\\ \text{Improvement (\%)}\end{array} = \frac{\text{Power-corrected} - \text{Power-corrected Original}}{\text{Power-corrected Original}} \times 100 \quad (7)$$

$$\begin{array}{c}\text{Measured}\\ \text{Improvement (\%)}\end{array} = \frac{\text{Measured} - \text{Measured Original}}{\text{Measured Original}} \times 100 \quad (8)$$

The results show that tuning the model improves both measured and estimated signal quality across all wavelengths. We observe 1.6 dB (10.68%) improvement in directly measured SNR values and 1.8 dB (12.2%) in power-corrected estimates.

In Table 5, we compare the modeling error against both power-corrected and directly measured values after tuning the network. While tuning improves absolute signal quality, we see a modest increase in relative error for both sets of values.

6 Related work

Role of SNR and modeling in optical WAN: SNR is a key metrics in planning and optimization of optical networks, playing a critical role in determining signal quality and thus the link capacity. Many recent works focus on utilizing SNR to optimize network operations which could benefit from accurate SNR estimates and optical network modeling. **RADWAN** [34] adapts link capacity based on observed SNR, but it treats SNR as an external input, limiting its ability to anticipate signal degradation. An optical network model could enable predictive rate adaptation under dynamic conditions. **Shoofly** [33] reduces hardware costs by using optical bypass, but their feasibility depends on maintaining acceptable SNR. An optical network model can evaluate the impact of bypasses on signal quality and guide placement decisions. **Arrow** [38] addresses recovery from fiber cuts by reconfiguring wavelengths onto healthy fibers. An optical network model can help optimize fiber selection and place wavelengths on fibers with the highest SNR. **OpTel** [23] streams fine-grained optical telemetry data, which could be complemented by a model to provide deeper insights into SNR variations during short-lived optical events. **OWAN** [19] optimizes optical WAN bandwidth and an optical network model could help improve spectrum utilization and reduce unnecessary signal regeneration by better understanding the physical feasibility of different configurations. **FlexWAN** [24] focuses on cost-effective planning of optical infrastructure. An optical network model could provide more accurate end-to-end SNR of the optical infrastructure which can improve decisions. In such systems, SNR modeling can improve decision-making and optimize network operations.

Diagnosing and predicting optical failures: Several recent studies highlight the importance of understanding the failures in optical WANs. **Ghobadi et al.** [15] show that optical signal quality can be used to predict future outages and that link availabilities vary over time. **Zhang et al.** [37] present a diagnosis framework that leverages large-scale telemetry for root-cause inference across layers. Our work complements both the efforts. An optical network model enables predictive analysis of performance degradation, even in the absence of live measurements, and can improve failure diagnosis and management by providing insights on expected physical-layer behavior.

Modeling optical network with GNPy: Several prior works have established the utility of GNPy for modeling optical networks. **Curri et al.** [6] outlines the foundational components—including models for fiber spans, amplifiers, ROADMs, and transceivers, and proposes its use as a digital twin for network planning of optical systems. GNPy has been experimentally validated by **Ferrari et al.** [13] using mixed-fiber, Raman-amplified, multivendor testbeds, showing high accuracy in predicting SNR. Lab-based studies **D'Amico et al.** [7] have further demonstrated GNPy's precision in a lab environment. Additionally, recent efforts like **Zhang et al.** [36] combine the GN model with deep learning to improve scalability and prediction quality in CL-band optical networks.

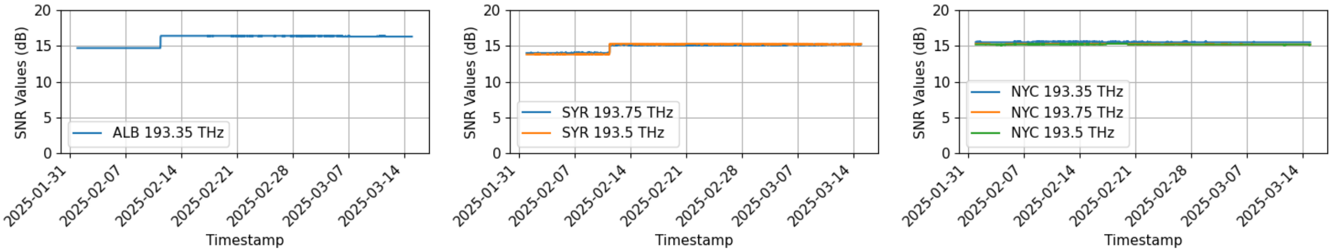

(a) SNR change over time at Albany transceiver for the end-to-end path NYC→Albany.

(b) SNR change over time at Syracuse transceiver for the end-to-end path NYC→Syracuse.

(c) SNR change over time at NYC transceiver for NYC→Albany→Syracuse.

Figure 7: SNR changes over time at different transceivers for various paths in the network.

Channel (THz)	Model Estimate (dB)	Power-corrected (dB)	Directly Measured (dB)	Power-corrected Error (%)	Measured Error (%)
193.35	16.91	15.56	15.5	8.68%	9.10%
193.5	16.07	14.93	15.1	7.65%	6.43%
193.75	16.25	15.07	15.2	7.83%	6.91%

Table 2: Comparison of model estimates and measured dB values for $Syr \rightarrow Alb \rightarrow NYC$.

Channel (THz)	Model Estimate (dB)	Power-corrected (dB)	Directly Measured (dB)	Power-corrected Error (%)	Measured Error (%)
193.35	16.23	15.06	14.7	7.77%	10.41%
193.5	15.39	14.39	13.9	6.96%	10.72%
193.75	15.46	14.46	14.1	6.91%	9.68%

Table 3: Comparison of model estimates and measured dB values for $NYC \rightarrow Alb \rightarrow Syr$.

Channel (THz)	Power-corrected Original (dB)	Power-corrected Optimized (dB)	Measured Original (dB)	Measured Optimized (dB)	Power-corrected Improvement (%)	Measured Improvement (%)
193.35	15.06	16.67	14.70	16.5	10.68%	12.24%
193.5	14.39	15.73	13.90	15.3	9.32%	10.07%
193.75	14.46	15.79	14.10	15.1	9.20%	7.09%

Table 4: Improvement in power-corrected and measured SNR values between original and optimized models.

Channel (THz)	Model Optimized (dB)	Power-corrected (dB)	Directly Measured (dB)	Power-corrected Error (%)	Measured Error (%)
193.35	18.49	16.67	16.5	10.92%	12.06%
193.5	17.14	15.73	15.3	8.96%	12.01%
193.75	17.22	15.79	15.1	9.06%	14.01%

Table 5: Comparison of post optimization model estimates and measured dB values for $NYC \rightarrow Alb \rightarrow Syr$.

While these studies validate GNPy in simulation and controlled environments, our work is the first to model and optimize a field-deployed optical WAN using GNPy.

7 Operational experience and lessons learned

In developing the model and deploying the optimized configuration, we encountered several operational challenges and gained insights into the process of applying the Gaussian Noise model to optical networks.

Data collection challenges: While developing the model, we found that obtaining accurate data about network components was a significant challenge. Many equipment details were proprietary, requiring reliance on vendor datasheets. Even with collaboration with network operators, some gaps due to the vendor restrictions may remain.

Deployment constraints: Operators face challenges implementing network changes, which requires careful planning. Even though the model can produce optimal configurations, some parameters in the network cannot be adjusted without risking stability due to equipment age or other factors. During deployment, we had to adhere to maintenance windows and ensure all changes within specifications.

Model accuracy vs. utility: Although models are not perfectly accurate, they remain valuable for optimization and "what-if" analysis. Even approximate models can provide actionable insights and help operators identify better configurations.

Generalizability to diverse networks: Real-world networks often include components that differ across deployments. However, with accurate data and appropriate corrections, the model can be adapted to diverse networks.

8 Conclusion

We develop a Gaussian Noise model of a New York state ISP's optical network and apply hyperparameter search techniques to optimize amplifier gain settings, improving end-to-end signal quality. A field deployment of the optimized settings shows an SNR improvement of nearly 2 dB per wavelength.

9 Acknowledgements

We sincerely thank our shepherd Prof. Robert Ricci for helping us improve the paper and the anonymous reviewers for their constructive feedback and suggestions. This work was supported in part by ACE, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA. The authors of this work are also supported by NSF Award #2444537.

References

[1] Extending gnpy. https://gnpy.readthedocs.io/en/master/extending.html. Accessed: 2024-09-19.

[2] Govind P. Agrawal. *Fiber-Optic Communication Systems*. Wiley, 5 edition, 2021.

[3] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 2623–2631, New York, NY, USA, 2019. Association for Computing Machinery.

[4] Vivek Alwayn. *Optical Network Design and Implementation*. Cisco Press, 2004.

[5] Cisco. Practical aspects of raman amplifier. https://www.cisco.com/c/en/us/support/docs/optical-networking/ons-15454-m6-multiservice-transport-platform-mstp/212834-practical-aspects-of-raman-amplifier.pdf, 2022. Accessed: 2025-04-22.

[6] Vittorio Curri. Gnpy model of the physical layer for open and disaggregated optical networking

invited

. *J. Opt. Commun. Netw.*, 14(6):C92–C104, Jun 2022.

[7] Andrea D'Amico, Elliot London, Bertrand Le Guyader, Florian Frank, Esther Le Rouzic, Erwan Pincemin, Nicolas Brochier, and Vittorio Curri. Experimental validation of gnpy in a multi-vendor flex-grid flex-rate wdm optical transport scenario. *J. Opt. Commun. Netw.*, 14(3):79–88, Mar 2022.

[8] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *Proceedings of the ACM SIGCOMM 2015 Conference*, 2015.

[9] Andrea D'Amico, Vittorio Gatto, Antonino Nespola, Giacomo Borraccini, Yanchao Jiang, Pierluigi Poggiolini, Esther Le Rouzic, Arturo Mayoral López de Lerma, Gert Grammel, Roberto Manzotti, and Vittorio Curri. Gnpy experimental validation in a c+l multiband optical multiplex section. In *2024 24th International Conference on Transparent Optical Networks (ICTON)*, pages 1–4, 2024.

[10] René-Jean Essiambre, Gerhard Kramer, Peter J. Winzer, Gerard J. Foschini, and Bernhard Goebel. Capacity limits of optical fiber networks. *Journal of Lightwave Technology*, 28(4):662–701, 2010.

[11] Qirui Fan, Guodong Zhou, Tao Gui, Chao Lu, and A. Pak Tao Lau. Advancing theoretical understanding and practical performance of signal processing for nonlinear optical communications through machine learning. *Nature Communications*, 11:3694, 2020.

[12] Alessio Ferrari, Mark Filer, Karthikeyan Balasubramanian, Yawei Yin, Esther Le Rouzic, Jan Kundrát, Gert Grammel, Gabriele Galimberti, and Vittorio Curri. Gnpy: an open source application for physical layer aware open optical networks. *J. Opt. Commun. Netw.*, 12(6):C31–C40, Jun 2020.

[13] Alessio Ferrari, Mark Filer, Karthikeyan Balasubramanian, Yawei Yin, Esther Le Rouzic, Jan Kundrát, Gert Grammel, Gabriele Galimberti, and Vittorio Curri. Gnpy: an open source application for physical layer aware open optical networks. *J. Opt. Commun. Netw.*, 12(6):C31–C40, Jun 2020.

[14] Alessio Ferrari, Mark Filer, Esther Le Rouzic, Jan Kundrát, Bruno Correia, Karthikeyan Balasubramanian, Yawei Yin, Gert Grammel, Gabriele Galimberti, and Vittorio Curri. GNPy: An Open Source Planning Tool for Open Optical Networks. https://github.com/Telecominfraproject/oopt-gnpy.

[15] Manya Ghobadi and Ratul Mahajan. Optical layer failures in a large backbone. IMC'16, November 2016. Best dataset award.

[16] G. Scott Glaesemann. Optical fiber mechanical reliability. Technical report, Corning Incorporated, July 2017. White Paper WP8002.

[17] Corning Incorporated. Corning® leaf® optical fiber product information. https://www.corning.com/media/worldwide/coc/documents/Fiber/LEAF%20optical%20fiber.pdf, 2023. Accessed: 2025-04-24.

[18] ITU-T. Generic characteristics of raman amplifiers and raman amplified subsystems. ITU-T Recommendation G.665, International Telecommunication Union, January 2005.

[19] Xin Jin, Yiran Li, Da Wei, Siming Li, Jie Gao, Lei Xu, Guangzhi Li, Wei Xu, and Jennifer Rexford. Optimizing bulk transfers with software-defined optical wan. In *Proceedings of the 2016 ACM SIGCOMM Conference*, pages 87–100, 2016.

[20] Abhishek Vijaya Kumar, Bill Owens, Nikolaj Bjørner, Binbin Guan, Yawei Yin, Paramvir Bahl, and Rachee Singh. CHISEL: An optical slice of the wide-area network. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 859–875, 2024.

[21] Hongqiang Harry Liu, Yibo Zhu, Jitu Padhye, Jiaxin Cao, Sri Tallapragada, Nuno P. Lopes, Andrey Rybalchenko, Guohan Lu, and Lihua Yuan. Crystalnet: Faithfully emulating large production networks. In *Proceedings of the 26th Symposium on Operating Systems Principles*, SOSP '17, page 599–613, New York, NY, USA, 2017. Association for Computing Machinery.

[22] Toru Mano, Yue-Kai Huang, Giacomo Borraccini, Ezra Ip, Andrea D'Amico, Zehao Wang, Hideki Nishizawa, Gil Zussman, Tingjun Chen, Ting Wang, Koji Asahi, Daniel Kilper, Vittorio Curri, and Koichi Takasugi. Modeling the input power dependency of transceiver ber-onsr for qot estimation. In *2024 Optical Fiber Communications Conference and Exhibition (OFC)*, pages 1–3, 2024.

[23] Congcong Miao, Minggang Chen, Arpit Gupta, Zili Meng, Lianjin Ye, Jingyu Xiao, Jie Chen, Zekun He, Xulong Luo, Jilong Wang, and Heng Yu. Detecting ephemeral optical events with OpTel. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, pages 339–353, Renton, WA, April 2022. USENIX Association.

[24] Congcong Miao, Zhizhen Zhong, Ying Zhang, Kunling He, Fangchao Li, Minggang Chen, Yiren Zhao, Xiang Li, Zekun He, Xianneng Zou, and Jilong Wang. FlexWAN: Software Hardware Co-design for Cost-Effective and Resilient Optical Backbones. In *Proceedings of the ACM SIGCOMM 2023 Conference*, page 319–332, 2023.

[25] Adva Optical Networking. Adva fsp3000r7 datasheet, 2014.

[26] Adva Optical Networking. Adva fsp3000c datasheet, 2016.

[27] Optuna. Optuna: A next-generation hyperparameter optimization framework. 2019.

[28] Rüdiger Paschotta. Tutorial: Fiber amplifiers. https://www.rp-photonics.com/tutorial_fiber_amplifiers9.html. Accessed: 2025-04-24.

[29] P. Poggiolini, G. Bosco, A. Carena, V. Curri, Y. Jiang, and F. Forghieri. The gn-model of fiber non-linear propagation and its applications. *J. Lightwave Technol.*, 32(4):694–721, Feb 2014.

[30] Talha Rahman, Antonio Napoli, Danish Rafique, Bernhard Spinnler, Maxim Kuschnerov, Iveth Lobato, Benoit Clouet, Marc Bohn, Chigo Okonkwo, and H. Waardt. On the mitigation of optical filtering penalties originating from roadm cascade. *Photonics Technology Letters, IEEE*, 26:154–157, 01 2014.

[31] C. E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.

[32] Jane M Simmons. *Optical Network Design and Planning*. Springer, 2014.

[33] Rachee Singh, Nikolaj Bjørner, Sharon Shoham, Yawei Yin, John Arnold, and Jamie Gaudette. Cost-effective capacity provisioning in wide area networks with Shoofly. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, page 534–546, 2021.

[34] Rachee Singh, Manya Ghobadi, Klaus-Tycho Foerster, Mark Filer, and Phillipa Gill. RADWAN: Rate Adaptive Wide Area Network. In *Proceedings of the ACM SIGCOMM 2018 Conference*, August 2018.

[35] Tong Yu and Hong Zhu. Hyper-parameter optimization: A review of algorithms and applications, 2020.

[36] Yao Zhang, Min Zhang, Yuchen Song, Yan Shi, Chunyu Zhang, Cheng Ju, Bingli Guo, Shanguo Huang, and Danshi Wang. Building a digital twin for large-scale and dynamic c+l-band optical networks. *J. Opt. Commun. Netw.*, 15(12):985–998, Dec 2023.

[37] Ying Zhang, Nathan Hu, Carl Verge, and Scott O'Brien. Cross-layer diagnosis of optical backbone failures. In *Proceedings of the 22nd ACM Internet Measurement Conference*, pages 419–432, 2022.

[38] Zhizhen Zhong, Manya Ghobadi, Alaa Khaddaj, Jonathan Leach, Yiting Xia, and Ying Zhang. ARROW: Restoration-Aware Traffic Engineering. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 560—-579, 2021.