Optical Fiber Technology 76 (2023) 103243

Contents lists available at ScienceDirect

Optical Fiber Technology

journal homepage: www.elsevier.com/locate/yofte

(INVITED)Reconfigurable topology testbeds: A new approach to optical system experiments

Emmanuel Akinrintoyo [a], Zehao Wang [b], Bob Lantz [c], Tingjun Chen [b], Dan Kilper [a]

[a] CONNECT Centre, Trinity College Dublin, 34, Westland Row, Dublin 2, Ireland
[b] Electrical and Computer Engineering, Duke University, Durham, NC 27708, United States
[c] COSM-IC Project, Columbia University, New York, NY 10027, United States

ARTICLE INFO

Keywords:
Networking testbeds
Digital twin
Machine learning (ML)
Stimulated Raman scattering (SRS)
Software-defined networking (SDN)
Optical physical layer control

ABSTRACT

Optical transmission systems provide high capacity, low latency and jitter, and high reliability for city-scale networks. Recirculating loop experiments have facilitated the study of signal propagation in long-haul optical transmission systems. However, they are unsuited for developing control and management software for city-scale optical networks with dozens or hundreds of reconfigurable optical add drop multiplexer (ROADM) units, diverse interconnection topologies, and dynamic traffic patterns. Large-scale testbeds can help, but may be inflexible and time- or cost-prohibitive. Reconfigurable testbeds such as COSMOS enable piece-wise emulation of a city-scale network by applying space and wavelength switching, dual-use software-defined networking (SDN) controllers, and comb sources, while digital twin models enable software emulation. Results from the development of a digital twin for COSMOS are presented for optical amplifiers and stimulated Raman scattering (SRS) including both analytical and machine learning (ML) models.

1. Introduction

Optical systems are commonly designed to maximise transmission reach and capacity, as demonstrated in record-setting experiments for submarine and long-haul networks [1,2]. As the scale and cost of early experiments increased, laboratory experiments quickly adopted experimental emulation methods, such as the widely used recirculating loop method [3]. Today, high capacity optical transmission systems are deployed in metro networks and are of interest in high capacity radio front/backhaul edge networks. High capacity remains important in these short reach systems, alongside new demands such as low latency and adaptation to traffic variations. High capacity systems thus extend from the network edge to the core with diverse performance and control requirements, yet an alternative experimental platform to the recirculating loop has not emerged.

The introduction of open and disaggregated optical transmission systems creates further experimental needs. Most of the research challenges associated with such systems involve their control and management. In proprietary systems, control complexity is addressed through proprietary engineering methods that are not exposed to network operators [4,5]. While some components of optical system control, such as the path computation element (PCE) [6], have been well studied, the actual control plane operations, such as lightpath provisioning and amplifier tuning, have received much less attention and documentation. Control can be complex due to factors such as non-linear impairments, component anomalies, and optical power dynamics from wavelength and polarisation-dependent effects in amplifiers and fibre spans. Although short distances reduce non-linear effects, metro systems might have a large number of switching node hops, and there is interest in increasing channel power to support higher spectral efficiency modulation formats. However, if the power is increased too high, then non-linear impairments can be significant even in short distance systems. The open line system approach, in which only the transceivers are disaggregated, and the engineering of the line systems is kept proprietary, is an example of a short-term solution to the lack of related research in this area [7,8]. Data-driven and machine learning (ML) methods have attracted attention as potential means to handle the added control complexity of more fully disaggregated systems [9,10,11,12]. Addressing the larger challenge of fully disaggregated systems requires experimental platforms for studying the interactions between new control and management systems and the physical transmission effects.

Optical networking experiments have typically consisted of three to six nodes connected over short distances (e.g., a single transmission span). These networking experiments are effective at studying network operation and management, but do not accumulate physical impairments, and therefore are ineffective at studying the interaction between the control system and the transmission physics. Noise or other impairments can be artificially introduced, but this does not capture the associated dynamics. As optical networks today can scale to hundreds or thousands of nodes and support signal transmission over dozens of nodes

https://doi.org/10.1016/j.yofte.2023.103243
Received 9 June 2022; Received in revised form 18 November 2022; Accepted 11 January 2023
Available online 23 January 2023

[13], new network control systems need to be evaluated under conditions where the associated transmission impairments are active. Recirculating loops have been an effective tool for studying such impairments in large scale systems, but generally are not compatible with optical networking experiments. Therefore, new methods are required for experimentation to study the physical effects at scale and how they interact with the novel software controls, ML algorithms, or control hardware innovations.

The trend towards using data-driven controls for improved management and greater automation also brings new requirements to experimental methods. Digital twin methods [14,15], for example, enable training and evaluation of ML algorithms on a virtual model of the system prior to implementation in the operational system itself. Such methods require detailed datasets of the system characteristics and performance over a wide range of operating conditions. Obtaining data from commercial systems is difficult both due to business practices and operating requirements, such as customer privacy regulations. On the other hand, testbeds can play an essential role in collecting datasets for experimentation, particularly if they include a field-deployed fibre plant and programmable optical hardware.

Recently, city-scale testbeds are being deployed with emphasis on experimentation with different networking technologies. For instance, Bristol is Open [16] in the United Kingdom was developed with a focus on 5th-generation wireless and optical networking experimentation, and the Japan-wide Orchestrated Smart/Sensor Environment (JOSE) testbed [17] in Japan has a focus on Internet-of-Things (IoT) services. These testbeds enable the investigation of the interoperability of different applications and devices in practical, deployed network scenarios. They can also be used to collect datasets to facilitate the development of ML algorithms and data-driven methods as described above. A key question is whether such testbeds can be used to address the needs of control and management experiments for optical networks at scale.

The rest of this paper is organised as follows: Section 2 reviews recirculating loop methods and their development in the context of developing optical networking experiments that are sensitive to physical impairments. Section 3 highlights how large-scale networking testbeds can be leveraged as emulation platforms that provide accurate impairment modelling and scale to enable the investigation of transmission and control complexity of optical communication systems. Section 4 discusses the challenges and trade-offs of optical physical layer control, and the need for a practical research and development platform for software-defined networking (SDN) and ML-based systems that can address control and management complexities at scale. Section 5 proposes a path forward to address these research challenges. The methods to emulate large-scale systems are discussed with a focus on the use of digital twin techniques. Section 6 presents an example of component-level erbium-doped fibre amplifier (EDFA) and stimulated Raman scattering (SRS) digital twin models of the city-scale NSF PAWR COSMOS optical-wireless testbed [18,19,20] which were implemented using both analytical and ML-based techniques. The performance of both techniques is evaluated and compared with the COSMOS experimental results. Lastly, Section 7 concludes the paper. Note that portions of this paper provide an extended version of a paper presented at the European Conference on Optical Communications (ECOC) 2021 [21].

2. Recirculating loop transmission experiments

Traditionally, experimental investigation of point-to-point optical systems has been conducted in laboratories using a recirculating loop. This comprises a link with amplified spans re-used in the loop to obtain long and ultra-long-haul distances (600 to 10,000 + km). A laboratory would need a prohibitively large amount of fibre and amplifiers to conduct such experiments. Hence, the use of a loop experiment was introduced as a method to conveniently study long-distance transmission physics without the cost and complexity of a full system [3].

A recirculating loop is an emulation tool that recirculates optical signals through a set of optically amplified fibre spans. The transmitter output signal is coupled into the loop for signal circulation through a switching mechanism that is activated for the duration required for the signal to fill the loop. The signal is then allowed to circulate through the loop multiple times. An optical tap is used to sample the signal with each successive circulation, and the resultant performance is measured. Thus, long-haul transmission distances can be emulated by recirculating a sample of optical data in a loop with just a few spans. Signal propagation over arbitrarily long distances can be emulated by keeping the signal circulating within the loop. While the optical signal experiences transmission over the equivalent distances, there are still differences between the emulated transmission and transmission within a full system.

Recirculating signals over the same fibre and components produce propagation effects that scale differently from a full-scale system. Random walk variations in polarisation effects and wavelength-dependent gain or loss (WDG or WDL) of a fibre span and components such as amplifiers, turn square root scaling into linear scaling with distance for related signal impairments. This error is mitigated by using a polarisation scrambler within the loop for the case of polarisation effects. Using a larger number of spans and amplifiers within the loop can further reduce the impact of such effects but at the cost of more hardware. However, one important limitation of recirculating loops is that all signals propagate together along the same effective route from end to end. Hence, the impact of dense wavelength- division-multiplexing (DWDM) transmission in mesh networks with diverse wavelength routing is not emulated [22].

ROADMs were introduced to provide optical path reconfigurability [25], but this leads to a signal path diversity that is not present in recirculating loop experiments. In addition, filtering effects due to the wavelength selective nature of the ROADMs can cause passband narrowing effects with distance. Although placing a ROADM within the loop helps to emulate these filtering effects [25], the random passband and signal wavelength variations of a full-scale system are again missed in a loop. On-off keying (OOK) modulation-based systems have the additional complication of requiring a dispersion map to manage the accumulated dispersion and the local residual dispersion in a transparent mesh network can vary [26]. Recirculating loops are constrained by the accumulated dispersion in the loop and do not allow for variations. The use of coherent transceivers significantly simplified system design by providing dispersion compensation within the transceiver and largely eliminated the need for dispersion maps. However, the active compensation controls within a coherent receiver are not compatible with recirculated loop data, and therefore are not used with loops. Instead, custom coherent receivers are used typically with offline signal processing [27,28].

ROADM-based networks that support mesh topologies introduce other complexities that are not well emulated through recirculating loops. The optical signals that propagate through the nodes in a ROADM network are not regenerated, and therefore accumulate impairments along the route. Due to the large multiplicity of possible configurations, the signal quality or quality of transmission (QoT) must be estimated at the time of provisioning to ensure that the engineering margins are not violated. A recirculating loop with an in-loop ROADM can emulate a path through a mesh network and test worst-case limits, but generally cannot replicate the range of signal and system configurations that might be found along an arbitrary transmission path. Often specific experiments are constructed to emulate and test specific types of impairments that may be problematic such as filter narrowing or neighbouring channel non-linear interference that might occur for different signal types [25]. Furthermore, the adding and dropping of a channel or group of channels may impact other channels in the transmission link since they share the same fibre. These effects, such as power excursions and crosstalk, have been the subject of numerous research studies [29,30,31,32,33,34]. In fact, the interaction of power fluctuations between the co-propagating channels can lead to feedback within a mesh network that can destabilise the system, thus requiring network-wide

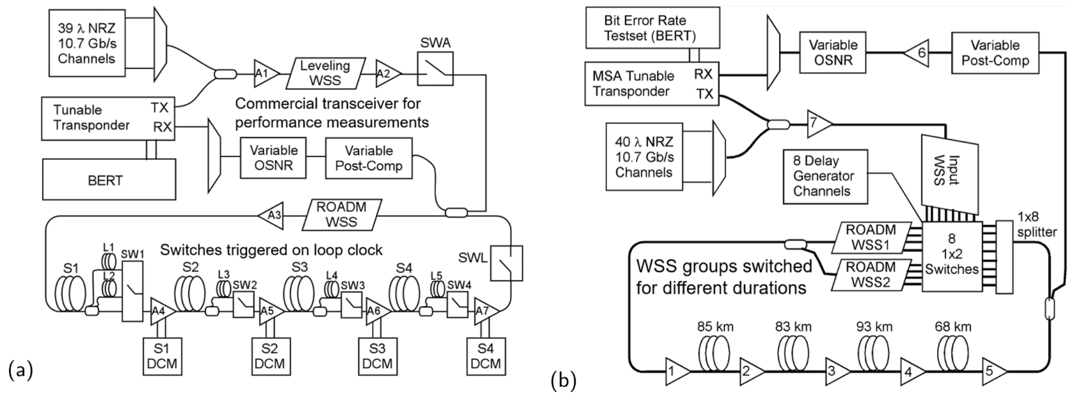

Fig. 1. (a) Experimental setup for a loop with different span lengths for each round trip [23]. (b) Recirculating loop that acts as 8 loops in one to allow 8 different groups of channels to propagate different numbers of round trips and study mesh network wavelength assignment [24].

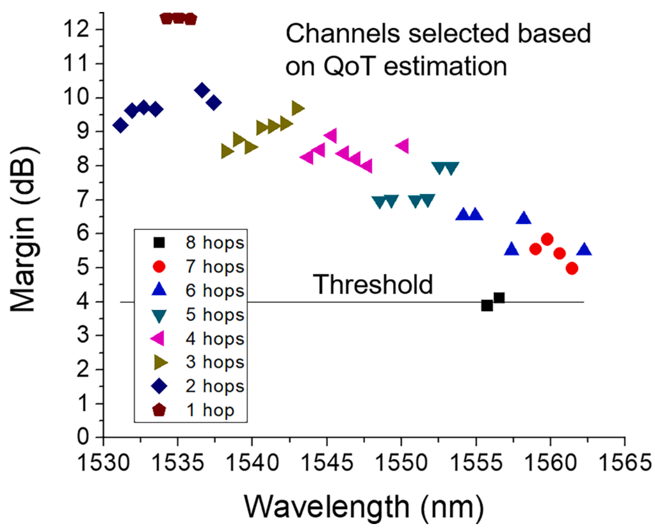

Fig. 2. Experimental results of wavelength assignment using QoT estimation for loop experiment in Fig. 1b. Each group of channels propagates one or up to 8 round trips (i.e., emulated mesh network hops) in the loop [24].

control [35] that cannot be emulated in a loop.

A number of recirculating loop methods were developed in an attempt to emulate mesh network effects. The experimental setup in Fig. 1a was used to investigate a dynamically reconfigurable span length circulating loop. The loop consists of four spans of lengths 77.5 km, 82.5 km, 90.2 km, and 67.2 km, respectively, with a split on the output of each span. Five additional fibre spools of lengths (L1–L5) 5 km, 10 km, 5 km, 2.3 km, and 5.3 km, respectively, can be introduced to the loop through switches that are triggered synchronously with the loop. In this way, the length of each span for each round trip can be varied to realise a wide range of different conditions. Constant power operation of the amplifiers compensated for the loss variations. This was effective in studying path-dependent dispersion variations [26].

In order to study the impact of channels propagating different distances along a path, the setup in Fig. 1b uses a ROADM in a recirculating loop experiment with eight different channel groups. Different channels could be assigned to different groups to travel various distances within the loop using eight independently triggered 1×2 loop switches. The channel groups travel through the same fibre, with each group travelling either once or up to eight times around the loop. QoT was used as a metric to perform wavelength assignment. It was used to choose the group of channels that should propagate through different distances in the loop. The best performing wavelengths were selected starting at the longest distance from the eight groups based on the QoT prediction (based on the required OSNR for a bit error rate of 10^{-3}). The results obtained are shown in Fig. 2. It can be seen that the lowest and highest margins were recorded for the groups of channels that propagate through eight hops and one hop, respectively. This approach could be used to study impairment-aware wavelength assignment. However, the complexity of the arrangement made it difficult to operate.

These mesh experiments show that further utility might be achieved with recirculating loops by emulating more complex mesh network effects. However, full-scale system experiments are still needed with the ability to incorporate the control system dynamics, both for the systems and in the transceivers. In fact, today, recirculating loops are primarily used for research hero experiments that push the limits of distance and capacity. Commercial development labs instead rely on a combination of full system experiments and simulation to develop and test optical transmission performance [8]. Research methods that can better emulate this production development environment are needed to progress research on such systems, including their control complexity.

3. Large-scale networking testbeds

Large-scale testbeds have the potential to address the transmission and control complexity needed in optical networking experiments. Some national-scale testbeds have been developed for networking research. In general, they require significant capital and time for planning, development, and deployment before becoming accessible for research. Remote programmability is particularly important at this scale, which can be problematic for optical systems. GENI [36] is an example of a U.S. national-scale testbed. Due to the complexity of operating such a network, commercial optical systems were used, which excluded most optical systems research. Extensions at specific nodes on GENI enabled local optical system experiments that could then communicate across the larger national network [37]. Integrating multiple optical network testbeds with a common control plane allows for larger-scale control experiments, although the physical layer effects are still missing. This federation of testbeds has been widely used to investigate control system complexity and scaling [37]. FABRIC [38] is a more recent U.S. national-scale testbed that again allows for a federation of optical testbeds at different end nodes (e.g., university campus labs and networks), but also potentially allows for experimentation on the commercial optical systems due to greater functionality in current commercial systems that might be exposed to the users.

With the emergence of smart city technologies, city-scale testbeds have attracted much attention in recent years. Similar to GENI, most city-scale testbeds do not have optical networking capabilities and often have heavy applications focus. The SmartSantander testbed [39], deployed in Santander, Spain, is one of the first city-scale testbeds in Europe that comprises over 10,000 IoT devices which consist of both fixed and mobile sensor nodes, gateway devices and near-field

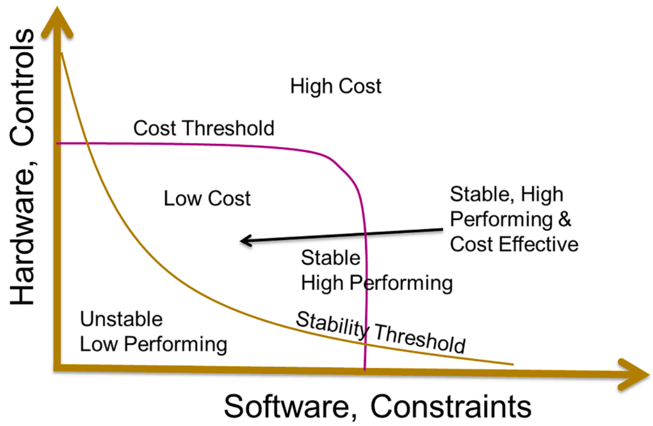

Fig. 3. Balance of software and hardware controls for the control and management optical communication systems. Practical systems require an open region between the cost and stability/performance thresholds: stable, high performing, and cost-effective.

communication (NFC) tags. SmartSantander has supported the research community in providing real-world datasets from its various sensors. Likewise, the CityLab testbed [40] located in Antwerp, Belgium, is a similar testbed that researchers are using for data collection to model real world traffic patterns. Its capabilities, such as the IEEE 802.15.4, IEEE 802.11, and sub-GHz wireless protocols, support multi-technology experimentation.

Smart city testbeds that use commercial networks allow for experiments using commercially available technologies. However, future technologies will run on networks with higher performance than current commercial networks. Therefore, using a research network with capabilities beyond those that are commercially available will enable experimentation on such future applications and technologies. This concept translates through all layers of the network, and thus, ultimately, a research optical network is needed to support the full range of future smart city technologies. Bristol is Open [16] in Bristol, UK, is a smart city testbed that pioneered the use of research network components, allowing for research on both applications and the underlying networks. This enabled researchers to consider the concept of a smart city operating system [41]. It also used academic lab testbed facilities as a hub connected to the larger smart city network and enabled more advanced research. The COSMOS testbed in New York City, U.S., uses this approach to combine a city-scale advanced wireless network with a fully disaggregated and programmable optical network with hub labs at Columbia University, Rutgers University, and City College of New York. COSMOS extends across the Harlem area and Manhattan with a multi-layered computing architecture that comprises a general-purpose cloud, radio cloud and radio/optical physical layer. Its architecture includes millimetre-wave (mmWave) and sub-6 GHz software-defined radios (SDRs), a disaggregated ROADM optical network with user-programmable SDN control, and a core and edge cloud. The optical network currently includes 8 Lumentum ROADM whitebox units, which will expand to more than 36 units when the testbed is fully deployed. The availability of such whitebox and SDN programmable devices has opened the door to building research testbeds at these large scales.

Open and programmable optical components such as whitebox ROADM units and Ethernet switches supporting coherent transceivers further enable these city-scale research optical networking testbeds. SDN and SDR capabilities enable full programmability so that these components can be repurposed for a wide range of optical and wireless experiments. With sufficient scale, such programmable testbeds can be used to study the control dynamics of optically amplified systems, overcoming the limitations of the recirculating loop experiments. While these tools put the large scale networks within reach, populating such networks with a full set of transceivers is not practical. Configuring these network elements to form different network topologies also presents a challenge. Similar to the recirculating loops, emulation methods might be used to address these issues.

4. Optical physical layer control challenges

Optical transmission systems carry digital signals, but their operation is complex and analogue in nature. The optical signal-to-noise ratio (OSNR), often in its generalised form that includes both accumulated optical amplifier noise and the impact of non-linear fibre propagation, must be managed throughout the system. Operation is complicated by the fact that the optical signals interact with each other through wavelength-dependent phenomena in both the amplifiers and transmission fibre. High-speed data pattern dependent effects are generally managed through compensation controls within the transceivers. However, due to time-dependent temperature and stress on the physical system components, signal interactions persist over time scales that range from microseconds to diurnal and seasonal periods [42]. Advanced physical layer operations such as optical switching, signal modulation adaptation, and wavelength provisioning and routing are complicated by this highly non-linear, multitemporal physical system. In fact, these complications have limited the use of optical switching in commercial systems to passive optical networks (PONs), which are short reach (single hop) access networks without active optical amplification. A ROADM unit includes optical switches, but these are used for flexible provisioning and operated over time scales of minutes to hours. Commercial systems using real-time optical switching have been successfully developed, but the scale and topologies had to be constrained to less than roughly 16 nodes in a ring topology [43]. Despite numerous research studies of real-time switching, including optical packet switching [44,45], the development of commercial systems continues to be challenging among other things due to the operational cost and complexities at scale [46].

4.1. Balancing cost, complexity and performance

The development of advanced functionality in optical networks requires balancing cost and complexity to achieve the desired performance. Performance includes both the traditional measures of transmission performance, such as the pre-forward error correction (pre-FEC) operation, as well as other measures such as system stability and resilience. Given that the challenges in optical physical layer control arise due to analogue signal transmission effects, performance metrics can generally be met simply by limiting the so-called transparent scale of the system – the number of transparent node hops or transmission distance. A fully opaque system using optical regeneration at every node does not suffer any of the control challenges discussed here. Indeed this opaque solution has been introduced on multiple occasions over the history of optical system development [47,48]. The primary motivation for large-scale, transparent optical networks is to avoid the high cost of such optical regeneration. Transceivers are by far the highest cost component, even in ROADM systems designed to minimise the use of transceivers except at network ingress and egress points. Thus, most research challenges associated with advanced functionality in optical networks fundamentally come down to the trade-off of maintaining the transparent, low-cost aspects of the network while managing the control complexity and system performance at scale.

Fig. 3 illustrates the optical network control trade-off between cost and performance, considering the two dimensions of software and hardware controls. Software controls involve the use of intelligent algorithms such as machine learning or the introduction of operational constraints (e.g., slower switching transition times). Each of these incurs a cost in terms of development or higher system cost due to the added constraints. Similarly, hardware controls can be used, such as adding a dynamic gain equalising filter in each amplifier or increasing the amount of signal regeneration. The additional hardware directly drives

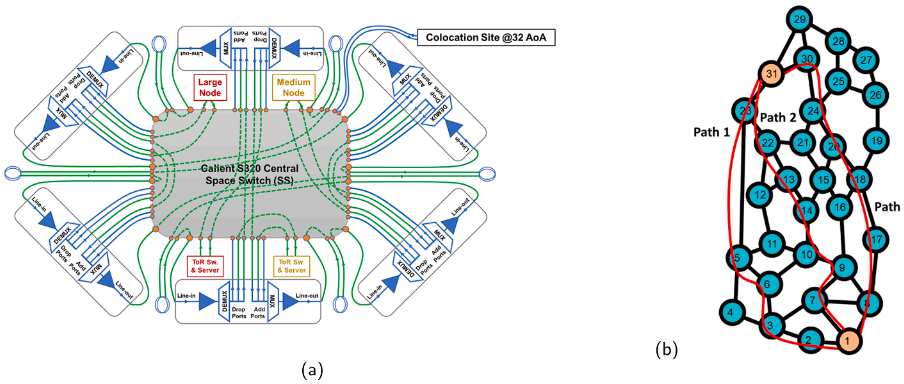

Fig. 4. (a) Linear uni-directional system topology realised using the Calient S320 space switch in COSMOS. (b) An illustration of how different paths can be emulated through a network.

up the network cost. Here a cost threshold is shown, above which the system becomes prohibitively expensive for the given application. Therefore, hardware and software controls need to be reduced in some combination to fall below this curve. Removing too many such controls, however, will eventually lead to performance degradations or issues, e.g., real-time switching might become unstable [49]. Thus a performance or stability threshold will also exist below which the system performs poorly or unreliably. The challenge then for optical network control research is to find solutions for which the performance curve falls below the cost threshold, allowing for high performing, cost-effective network. Since the performance or stability of the optical network is related to its transmission behaviour, experiments must be able to test such control systems at scale.

4.2. Applying SDN and ML-based control

The adoption of SDN has accelerated disaggregation by decoupling the data plane from the control plane, and it enables unified control and management [50]. At-scale experimental testbeds enable the study and development of effective SDN and ML-based control systems for these complex and disaggregated optical networks.

Transport application programming interface (*T*-API) [51] provides a standard northbound interface between the user (i.e., network operator/orchestration platform) and transport SDN controllers. This enables interoperability between SDN controllers from multiple sources. For example, the TeraFlow [52] project is developing methods to use machine learning with SDN control to integrate optical layer functionality with 5G and cloud operations, and supports the use of digital twin technologies for optical network performance prediction [53]. An operator might wish to combine it with the SYMPHONY [54] controller to provide compatibility with legacy networks. These SDN controllers and interfaces create an environment in which the provisioning and adaptation of optical signals can be controlled from customer or network operator user applications or automated cloud applications.

While monolithic SDN controllers can provide consistency, broad functionality, and one-stop shopping, they may limit flexibility and scalability. Frameworks such as application-based network operations (ABNO) [55] are being developed to offer a modular architecture. Microservices based controllers such as uABNO [56] and microONOS (μ-ONOS) [57] have also been proposed to allow for greater flexibility, fast deployment, and auto-scaling for optical networks than standard monolithic SDN controllers.

Running over commercial optical systems, on-demand provisioning operations can be carried out stably over the time scales of minutes. ML models can make dynamic configuration adjustments based on the field conditions to perform optical layer controls. Provisioning a new optical signal along a network path involves generating a QoT estimate in order to determine if the signal can reach the endpoint error-free and with a sufficient performance margin. The use of a QoT estimate, in fact, is an example of a digital twin method and recently, there has been an interest to incorporate machine learning in this digital twin framework. The motivation for using ML is to tighten QoT margins, particularly for disaggregated systems, and enable greater autonomic controls. Many recent studies of such methods have used synthetic data to train the models [58]. This can lead to inaccuracies that may differ significantly from the deployed systems, further emphasising the need for large scale networking testbeds.

5. A new generation of system experiments

Historically, testbeds have provided essential platforms for investigating optical systems technologies. With greater emphasis on system control, new testbed methods are needed at adequate scales that map to commercially deployed transmission systems. Emulation techniques are crucial in order to reduce the cost and complexity of such systems so that they can be readily used in research experiments. Recent techniques that have emerged in this context are discussed in the subsequent sections.

5.1. City-scale testbeds

City-scale testbeds provide an excellent platform in which to realise at-scale optical systems for control experiments. While in principle a large-scale testbed can be built up within a single laboratory environment, as is done with recirculating loop testbeds, the amount of equipment and the complexity of the system will typically exceed the resources of a single lab. Nevertheless, industrial labs can and have developed such at-scale systems [59], although such labs may have a product focus that involves the use of commercial systems without the full range of flexibility desired in research labs. City- scale testbeds, however, often bring together a consortium of partners and resources that can enable the necessary larger scale. Furthermore, they are often designed to be multi-user testbeds, enabling a much larger group of researchers to make use of the infrastructure. Sharing large-scale research infrastructure will be essential for stimulating research within the larger community.

The scale and complexity of city-scale testbeds creates support, usability, and maintenance challenges. For a multi-user testbed to be viable, dedicated staff is needed to maintain the equipment and work with users to overcome unique challenges that arise in many experiments when novel technologies are being investigated. Depending on the experiments, users may require large amounts of time and assistance in order to learn how to operate the testbed. Sandbox systems within the testbed can be used to prototype and debug experiments before testing in the full system. Significant effort is required to develop convenient user interfaces and sandboxes, and these tools will require continual upgrades to keep pace with technology.

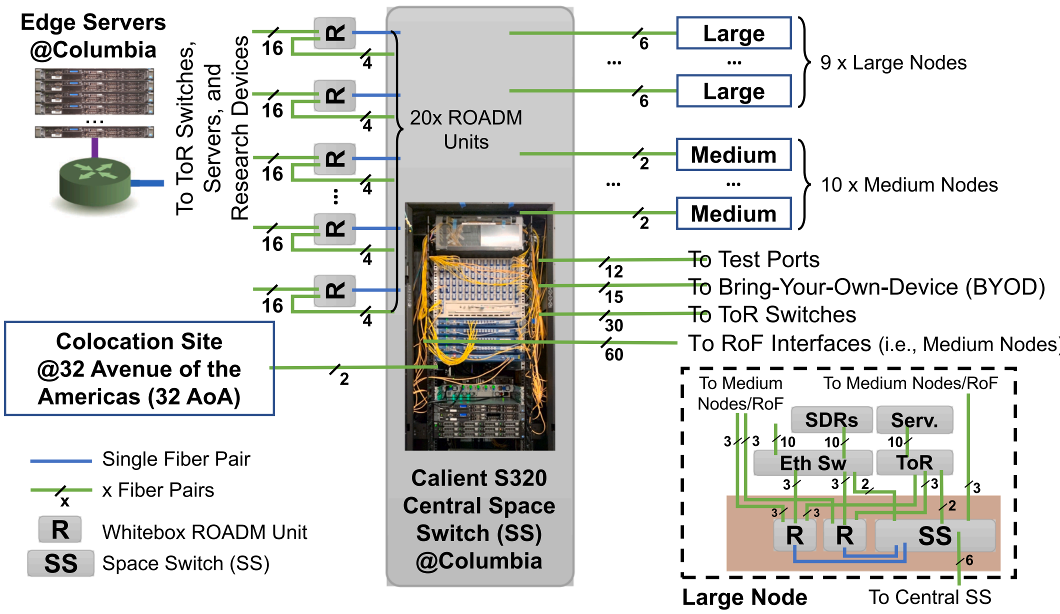

Fig. 5. The COSMOS optical testbed consisting of a Calient S320 space switch, whitebox ROADM units, edge servers and multiple nodes. A detailed description of the testbed is provided in [66].

5.2. Space switching

Outside fibre plant is expensive and difficult to manage. It often involves fixed topologies that can confine the range of configurations needed for experimental research. An important tool to overcome this limitation is the use of space or fibre switching on both the line-side and client-side of the networking equipment. A space switch is essentially a large programmable fibre patch panel. By connecting the fibre plant to such a switch, the interconnections between this fibre plant can be flexibly reconfigured for a wide range of experiments similar to the configuration shown in Fig. 4a while different paths through a network can be emulated as illustrated in Fig. 4b. Attaching spooled fibre of different lengths provides another degree of freedom for increasing the span lengths or introducing further topology variations. In this way, systematic studies of span length and distance dependence can be carried out, similar to transmission over multiple round trips in a loop. The optical hardware, such as whitebox ROADM units, optical transceivers, splitters and variable optical attenuators (VOAs), and other components, can also be attached to the space switch to allow for many different node architectures and component configurations. The same can be done for test equipment such as optical spectrum analysers and optical modulation analysers. Hardware (re)configuration through remotely controlled switches recovers some of the flexibility of the recirculating loops, while enabling control system experiments at scale.

Some challenges arise from using an optical spaceswitch to create a programmable testbed environment. For example, the switch has loss and transmission characteristics that can interfere with the transmission performance. Generally, these switches can have losses in the 1–3 dB range. In connecting a line amplifier to a transmission span, the amplifier output power needs to be increased by this amount to achieve the desired launch power. Line amplifiers often do not have an extra 1–3 dB of headroom to compensate for this loss. In addition, MEMS space switches can have input power limits. These issues are less of a concern for short-reach edge networks, as non-linear effects are less important and link budgets often have headroom.

5.3. Comb sources/shaped ASE channel loading

As mentioned previously, optical transceivers are typically the most expensive part of an optical system. Populating even dozens of nodes with the full system capacity to enable the full range of experiments can be prohibitively expensive. Spectrally shaped amplified spontaneous emission (ASE) noise can be used to create optical signal combs to emulate transmission signals. This method significantly reduces hardware and complexity. The ASE noise can serve as a substitute for the interfering data channels, as shown in [60,61]. An ASE noise source can be used for channel emulation by applying the source directly to a flexible grid wavelength selective switch (WSS) to limit the total bandwidth of the ASE source and flatten the spectral profile.

Numerous copies of this ASE comb can be made using optical splitters and amplifiers and sent to any add-drop port within the network. These emulated signals can then be routed through the network and managed with similar optical power and spectral density as the corresponding signal channels. This method can be used to emulate an arbitrary number of interfering optical signals. In addition, when performance measurements are taken on a particular signal, the ASE signal can be blocked at its ingress WSS and replaced with the output from the transceiver of interest, along with any number of nearest neighbours, depending on the needs of the experiment.

The interference effects of such ASE noise-based comb sources will differ from the effects of actual modulated signals, and several studies have examined these differences [61]. However, further investigation is needed. Effects such as optical power dynamics that are only sensitive to the mean optical power as a function of wavelength will be well emulated through the use of such comb sources. Furthermore, the exact spectral shape of the comb lines will likely differ from a modulated signal, particularly as the signal propagates over distance and experiences different forms of distortion and spectral broadening. Modelling spectral shape will be essential for studying wavelength filtering effects and related crosstalk [62].

5.4. Dual-use SDN controllers

The introduction of SDN controllers in optical communications has enabled external programmatic control of optical systems, often using open whitebox hardware [63]. While such controls are of research interest, they can also be exploited as a tool for testbed reconfiguration. This yields the flexibility that is essential for testbed experimentation. For example, in COSMOS, a Ryu-based SDN controller was developed to serve as the experimenter interface for setting up the optical system for

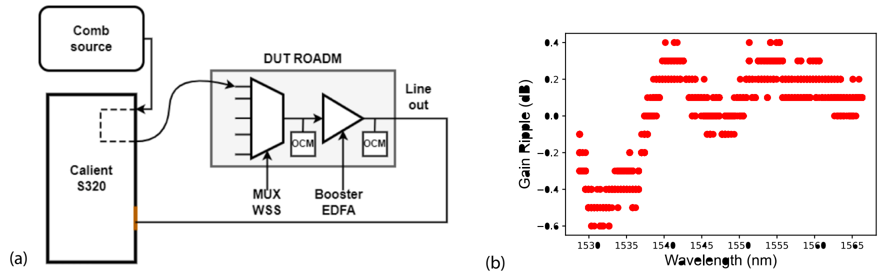

Fig. 6. (a) The data collection setup for a booster EDFA in the COSMOS testbed consists of a comb source, a Calient S320 space switch, two OCMs and a DUT ROADM. (b) 50 runs of the gain profile measurements using the booster EDFA within a Lumentum ROADM unit (rdm1-co1) deployed in the COSMOS testbed, under the fully loaded WDM spectrum with 95×50 GHz channels in the C-band.

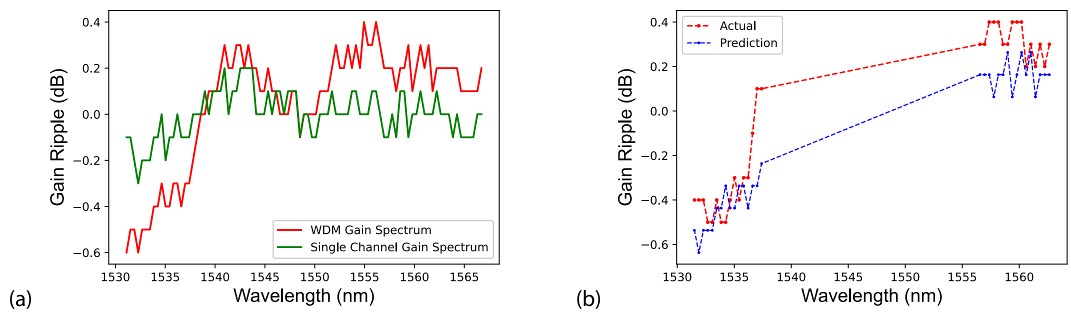

Fig. 7. (a) Gain ripple of the booster EDFA under single channel and WDM channel loading configurations 90 channels in the C-band. (b) Measured gain ripple compared with the predicted gain ripple using the developed CM model with 32 randomly selected channels.

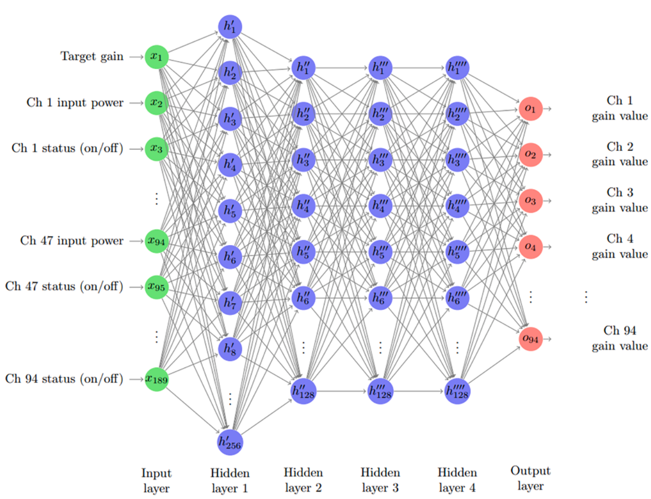

Fig. 8. A feed-forward DNN, consisting of an input layer, an output layer, and four hidden layers, predicts the EDFA wavelength-dependent gain, given a target gain, the input power and on/off status of each channel.

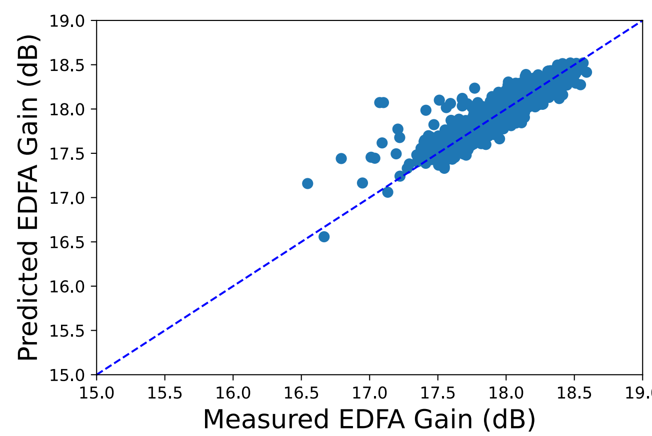

Fig. 9. ML-based prediction of the wavelength-dependent gain: The measured and predicted wavelength-dependent EDFA gain using the DNN model.

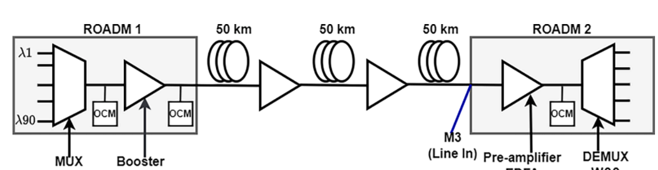

Fig. 10. An experimental setup consisting of two ROADMs with three 50 km spans and two in-line amplifiers. Monitoring is performed at the input of the second ROADM (M3).

different experiments [64]. This controller provides dual-use as a platform for SDN control research. Modular SDN software facilitates replacing or customising control algorithms for experimentation, and serves as a template that investigators can use to develop their own experiments with the testbed.

5.5. Digital twins

A digital twin is a software emulation of a specific physical system that enables control algorithms to be tested before deployment. A digital twin is constructed using data calibrated simulation models [65]. While digital twin models have been widely utilised in various fields such as aerospace engineering, smart manufacturing, and production engineering [15], they have not been widely adopted for optical transmission systems. Programmable testbeds can serve as platforms for data collection to build digital twins, and such a digital twin used with a testbed can be a tool for expanding the research capabilities of the testbed. For city-scale testbeds, the digital twin can allow researchers to

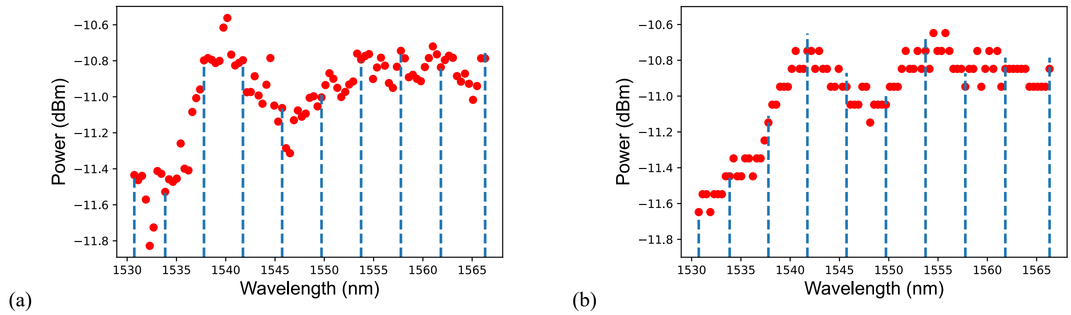

Fig. 11. Power excursions investigation with both the DNN-based (a) and CM (b) EDFA models, where the DNN-based EDFA model is able to better mitigate the effects of power excursions within the 50 GHz channel. The fully loaded spectrum of each model is depicted in red while the power levels of the ten surviving channels are the blue stems.

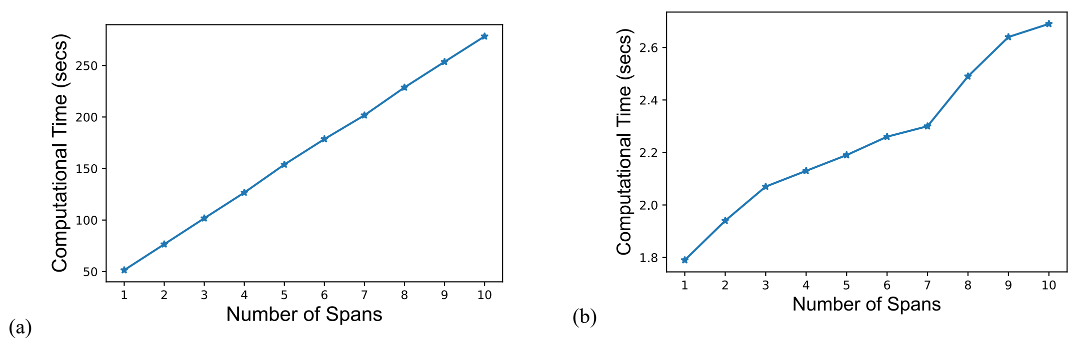

Fig. 12. Computation time for the DNN (a) and CM (b) EDFA models increases approximately linearly with the number of spans, but is much higher for the DNN model. The first span includes two EDFAs (a booster and preamp) with no in-line amplifier. One in-line amplifier is added for each subsequent span.

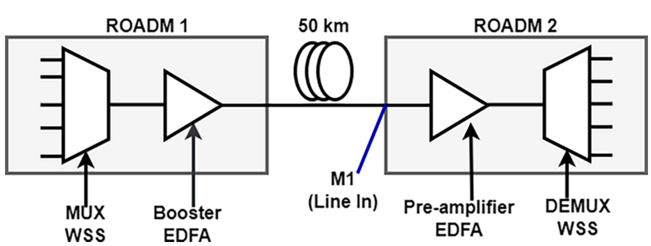

Fig. 13. An experimental setup consisting of two ROADMs with one 50 km span and no in-line amplifier. Monitoring is performed at the input of the second ROADM (M1).

Table 1
Spectral tilt comparisons between the measured tilt values from a COSMOS experiment and the predictions of the rectangular and general models for a 50 km fibre span with varying input power levels.

Input Power (dBm)	Tilt (dB)		
	Measured	RM	GM
0	1	1	1.2
−3	0.43	0.34	0.5
−6	0.11	0.01	0.17
−9	0.01	−0.15	0.07

perform rapid prototyping and testing of their technologies before dealing with the complexity of the full system at scale. Similarly to dual-use SDN control, digital twins can provide dual use: as a research tool to study network control, and as the subject of research on digital twin control methods.

Optical transmission systems can make use of digital twin models for the purpose of QoT estimation for routing and wavelength assignment (RWA). QoT predictions in commercial systems typically use an offline tool developed through extensive system simulations [67] and development laboratory experiments [68]. These tools in fact can be considered a form of digital twin, although they are generally used as configuration tools. Metrics such as generalised optical signal-to-noise (GOSNR) ratio or Q factor and required OSNR (ROSNR) are evaluated through QoT estimation, and transmission margins applied based on these estimates. The transition to a full digital twin approach would involve real-time modelling and greater use of data collection to update the digital twin over time. Developing real-time QoT estimation for online optical system control has led to widespread research on ML-based QoT estimation. More advanced digital twin methods shift the focus from parameter-based to data-based modelling and can extend this approach to a wider range of system functions beyond RWA.

Key challenges in the use of digital twins for large scale testbeds include the collection and curation of data and the construction of simulation or emulation tools for running the digital twin network. One approach used for the COSMOS testbed (depicted in Fig. 5) makes use of the Mininet-Optical emulation platform. An essential feature of Mininet-Optical is that it is designed such that an SDN controller developed to run in Mininet-Optical should also be able to run on the hardware that is being emulated. As a result, a digital twin of COSMOS in Mininet-Optical can be used to develop network controllers that can then be run on the hardware deployed in COSMOS. The first demonstration of this capability was made recently, though further development is needed to include a wider range of functions [64]. Note that a controller for the COSMOS testbed will be different from a controller for an operational system because COSMOS itself is a hardware emulation of a full network of optical transmission systems. COSMOS control includes controls for the testbed's space switch, comb source, ROADM units, as well as the various pieces of test and measurement equipment. In this way, the testbed digital twin will differ from a digital twin under investigation for system control – the former providing a vehicle for the study of the

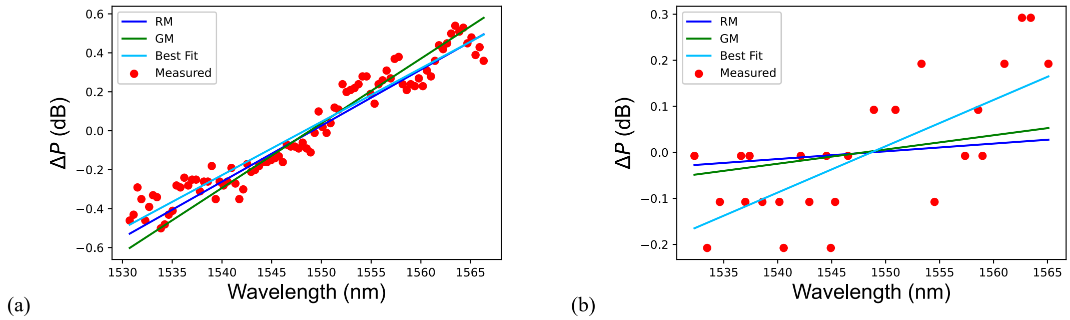

Fig. 14. Experimentally measured SRS tilt across a 50 km SMF-28 fibre for a 0 dBm per-channel launch power across 90 channels (a) and 26 randomly selected channels (b). Three lines are shown on the measured data plots to compare the predictions of the rectangular and general models with the line of best fit.

Table 2
Spectral tilt comparisons for random loading channel configurations between the measured tilt values from a COSMOS experiment and the predictions of both the general and rectangular models for a 50 km fibre span. Various combinations were used for each number of channels listed for the mean tilt values reported.

No. of Channels	Mean Tilt (dB)		
	Measured	GM	RM
26	0.29	0.10	0.05
35	0.41	0.24	0.17
41	0.46	0.33	0.25
48	0.56	0.45	0.36
56	0.63	0.61	0.50
63	0.72	0.69	0.57
74	0.85	0.82	0.69
84	0.94	1.02	0.87

latter.

In addition to running the control software in the digital twin environment, it is important that the software emulation can reproduce the physical effects of the hardware. This requires collecting data from the hardware and building appropriate models in the emulator. The remainder of this paper presents the development of such models within Mininet-Optical to create a digital twin of the physical COSMOS testbed.

6. COSMOS digital twin

A data collection process was conducted in COSMOS to build network element models, which were then deployed in Mininet-Optical [69]. Mininet-Optical is a software emulator for optical networks such as the COSMOS testbed, and it incorporates a steady-state transmission physics simulator. To build a digital twin of a physical testbed such as COSMOS, two general categories of data and models need to be considered: (i) end-to-end datasets, and (ii) component datasets. End-to-end datasets are generally collected in the process of experimentation, and include all the data associated with a given experiment or end-to-end system configuration. Component datasets, in contrast, are data collected specifically on individual components, such as amplifiers, fibre spans, and transceivers, to characterise their behaviour. These component datasets are particularly valuable for constructing digital twins as they enable virtual components to be created and then composed to form different experiments, as one would in a physical testbed. The end-to-end datasets can then provide an important source of validation of the digital twin models' accuracy at reproducing the physical testbed's behaviour. This work examines the development of amplifier and fibre span models for COSMOS in Mininet-Optical.

The testbed includes amplifiers, fibre spools, and space switches for which digital twin models can be constructed using analytical and ML models as considered in this work. Basic power measurements were performed for the switches and static losses throughout the system using the power monitors available in the network elements. The fibre spans are characterised using the GNPy model, fibre loss and SRS effects. The behaviour of an EDFA is highly configuration dependent and there is no optimal analytical model to describe the power dynamics. Recently, there have been numerous ML models developed to describe the power dynamics in the amplifiers. The main data collection for the digital twin considered in this work is related to characterising the amplifier gain ripple which is then input into the analytical models that account for the other network elements. The gain ripple is the difference between the mean gain and the actual gain in each channel.

6.1. Data collection

COSMOS' optical network includes Lumentum whitebox ROADM units, each of which has separate add WSS (MUX), drop WSS (DEMUX), as well as receive and transmit EDFAs, referred to as the preamp and booster EDFAs, respectively. A characterisation of the EDFAs in the ROADMs deployed in the COSMOS testbed was performed by measuring their gain profiles.

The approach taken here is to use the in-network tools (e.g., the optical channel monitors and optical power monitors built into the network elements) to make the measurements rather than relying on separate test equipment. The experimental data collection setup for the booster EDFA is depicted in Fig. 6a, which consists of a comb source that generates 95 channels with a 50 GHz spacing and a total channel power of + 5.8 dBm. The output of the comb source is fed into an add port of the device under test (DUT) ROADM. The booster EDFA of the ROADM was set to a target gain of 18 dB in constant gain mode with no gain tilt. The MUX WSS switches the on/off status of each channel and can also set an attenuation on each channel at 0.1 dB resolution. Channel power monitoring was performed using the built-in optical channel monitors (OCM) of the Lumentum ROADMs. This is a practical approach for field data collection since the setup described is used to emulate a deployed network.

The used booster EDFA is rdm1-co1 in the COSMOS testbed, as shown in Fig. 6b. The experimental process involved varying the channel configurations, including the on/off status and power level of each channel. The measured parameters include the input and output power spectra of the DUT booster EDFA. The gain profiles of the DUT EDFA were measured under six different conditions. These include single, double, odd, even, random, and WDM channel loadings. The single and double channel loadings involved turning on one and two channels each time, respectively. Odd-numbered and even-numbered channels were turned on for the odd and even channel loading cases. The channels to be loaded were randomly selected across the spectrum for the random loading measurements. For the fully loaded WDM gain profile, all the 95 channels were switched on.

6.2. EDFA modelling

Amplifier modelling is often done using numerical models [70], and simulation tools are commercially available. Recent analytical models were developed for online prediction of the amplifier gain dynamics under different channel loading conditions. These analytical models can serve as a valuable tool for evaluating the effectiveness of machine learning models. This work explores the use of both analytical and ML models for the characterisation of an optical amplifier. Both EDFA models were implemented in Mininet-Optical as user-selectable digital twins of the actual EDFAs deployed in the COSMOS testbed.

6.2.1. Analytical models

The Centre of Mass (CM) model [33] was implemented based on the collected booster EDFA gain profile data, described above, to predict the wavelength-dependent gain using the single and fully loaded gain spectrum. Fig. 7a shows the single and fully loaded gain spectra of the booster EDFA comprising a total number of 90 channels in the C-band. Ninety channels were considered for this modelling since the emulation platform (Mininet-Optical) was designed for 90 channels. More recently, Lumentum made 95-channel ROADMs available; hence, there is a difference that needs to be accounted for which will be updated in future Mininet-Optical releases. The CM model was used to predict the wavelength-dependent gain of different channel loading combinations and it achieved a mean absolute error (MAE) of 0.2 dB for 16 randomly selected channels across the spectrum. Likewise, the model achieved an MAE of 0.15 dB for a randomly selected 32-channel combination which is shown in Fig. 7b. These results are representative of the other cases which were considered.

While analytical models can serve as baseline models for prediction to compare with other advanced methods or as computationally efficient models, they can exhibit large errors for certain channel loading configurations [71]. Alternative techniques such as those based on ML models have been shown to provide a superior performance [71].

6.2.2. Machine learning (ML) models

ML enables computers to learn how to perform a specific task [72]. Supervised and unsupervised learning are the two categories of ML. Supervised learning involves using labelled data to learn the relationship between a given input and target output variable. This can then be used to predict the output based on other input data. Conversely, unsupervised learning uses unlabelled data to discover hidden information in a dataset. In this work, supervised learning is studied as a method to build a digital twin model for the amplifiers in a digital twin system using Mininet-Optical. Datasets from the measurements on an EDFA in the COSMOS testbed are used to build an ML-based EDFA model using deep neural networks (DNNs).

Previous work has investigated the use of ML methods in optical communications. For instance, the efficacy of utilising ML to make predictions for the wavelength-dependent gain of optical amplifiers was examined in [72,73,74]. The work in [71] compared the CM model with some ML models and found the latter to be superior at predicting the wavelength-dependent gain. Such trained models can be deployed as amplifier models in a simulator or emulator platform such as Mininet-Optical. Sufficient data is vital to train, validate and test DNN models; hence, 1652 gain spectra profiles comprising both fixed (80 %) and random (20 %) channel loading configurations were used to model the booster EDFA (as described in Section 6.1). This differs from the previous work that have modelled EDFAs since the known corner cases which can significantly impact an amplifier's behaviour are considered here. The dataset was split into three sets: the training, validation, and testing sets with a split ratio of 80 %–10 %–10 %. The validation set was used to optimise the hyperparameters of the DNN model.

The DNN architecture used to model the EDFA is shown in Fig. 8. It consists of a normalisation layer, input layer, four hidden layers, and an output layer, all of which are fully connected. The hidden layers comprise 256, 128, 128 and 128 neurons. The use of four hidden layers yielded an optimal performance with little training time. There are 189 and 94 neurons in the input and output layers, respectively. The inputs to the DNN include the target gain of the amplifier, input power of each channel and their corresponding status (on or off). The output layer predicts the wavelength-dependent gain of all 94 channels, one per output neuron. The DNN model was trained using the ReLU activation function with a learning rate of 0.01, and using 350 epochs with the Adam optimiser. In addition, the model's training was done using a customised MAE loss function that accounts for the fact that some of the channels in the random loading dataset are switched off.

Fig. 9 depicts the measured and predicted EDFA gain values to show the performance of the trained ML model at predicting the EDFA's wavelength-dependent gain. In particular, 97 % of the predictions were within 0.2 dB accuracy, with the model achieving a mean absolute error of 0.01 dB. This is significantly lower and more accurate than the CM model's results for the cases considered, such as the 32-channel random configuration presented in the previous subsection.

The DNN has a performance which is quite similar to that which was used in [71] although the work employed the use of 90 neural networks, one for each channel. In addition, a significantly larger dataset was used for the model's training and validation. The model achieved an MAE of 0.06 dB for an input power dynamic range of $\pm$ 3 dB. Similarly, the CM model's performance is slightly higher in comparison to those reported in [71].

6.2.3. Analytical vs ML-based EDFA models

The CM and DNN EDFA models were implemented in Mininet-Optical's physics simulator to evaluate their accuracy in modelling power excursions. It is essential to consider cases when there are a few channels in an optically amplified system, as they can lead to more severe excursions [72]. The setup consists of a linear topology with two ROADM units, three 50 km fibre spans and two in-line amplifiers as depicted in Fig. 10. An equal spacing of 50 GHz was used for the 90 channels launched into the system with a per-channel launching power of 0 dBm. Then, 80 channels were removed at the first ROADM unit while the remaining 10 channels propagated through the system, emulating the steady state response following an upstream fibre cut, before any system adjustments are made to compensate for the resulting power excursion on the 10 channels. By using uniform loading of the 10 channels, the final gain ripple and power offset of the 10 channels should follow the original spectrum of the 90channel configuration.

The results are presented in Fig. 11, where the red curves represent the 90-channel spectrum for each EDFA model, and the blue stems represent the power levels of the 10 channels at the input of the second ROADM unit following the excursion event. The results show a significant difference in the prediction of the gain spectrum shape. It can be observed that the DNN-based EDFA model achieved a closer performance to what is expected in its prediction of the gain spectrum shape for each of the 10 channels in comparison to the CM model. Although the CM model exhibits lower accuracy in predicting cases such as shown here, its static prediction of the wavelength-dependent gain requires less computational time than the dynamic prediction based on the DNN model.

The two models were evaluated on a 3.3 GHz Intel(R) Xeon(R) E-2136 CPU and Nvidia Quadro P1000 GPU. A computational time of approximately 51.4 s was recorded for the DNN model with a 1-span topology similar to Fig. 10 which includes booster and preamp EDFAs but no in-line amplifier. The computational time increased to 101.2 s for the 3-span topology (Fig. 10) and 278.3 s for a 10-span topology which indicates a computational time of approximately 25 s for each amplifier in the setup. However, this differs significantly for the CM model, which executed in 2 s for the 3-span topology. Additional amplifiers did not cause a major increase in the computation, with a 10-span topology requiring 2.7 s. The computational time comparison between the two models is depicted in Fig. 12. The computational time is significant for the DNN model due to its inference time which is approximately 0.3 s; hence, this accumulates with the number of propagating signals in the transmission system.

6.3. Stimulated Raman scattering (SRS)

SRS is a non-linear effect that creates an optical power dependent tilt in the signal power spectrum after transmission through a fibre span [75]. In addition to the wavelength-dependent fibre loss, longer wavelength signals experience gain through Raman pumping by the shorter wavelength signals, which are correspondingly depleted. This leads to a reduction in the OSNR of the shorter wavelength signals, which is sometimes compensated using a pre-emphasis tilt of the optical signal powers at the input to the fibre span. Analytical models such as the triangular approximation model [75] have been proposed to calculate the SRS effect in optically amplified systems. The model assumes a linear relationship between the Raman gain and frequency resulting in the triangular approximation for the Raman gain, which was further extended for different fibre types [76]. The triangular gain model (TGM) assumes a rectangular input power spectrum for the uniform loading case. Hence, this form of the model is referred to as the rectangular model (RM) for this work. The TGM model is computationally simple if the signals are assumed to have a similar launch power and it is rewritten as:

$$\left(\frac{P_N}{P_1}\right)_{dB} = 2.17 g_c L_{eff} P_{tot}(N-1)\Delta V \tag{1}$$

In the expression, g_c represents g'/A_{eff} where g' is the experimentally measured Raman gain coefficient per frequency unit, L_{eff} represents the fibre's effective length, P_{tot} is the total channel launch power, N represents the current channel's index relative to the first active channel, and ΔV is the spacing between the channels. This model also assumes that the channels are uniformly distributed across the spectrum, and therefore some error will arise for non-uniform loading. However, the Raman tilt only becomes significant when the spectrum is heavily loaded, therefore the amount of non-uniformity that can lead to errors is small unless there are very large differences in channel power, which is not considered here.

Non-uniform loading was investigated using the model's general form equation proposed in [[75], Eq. (7)] which can be used for any given input spectrum. Thus, it is rewritten as:

$$S(z,\lambda) = \frac{S(0,\lambda)P_0 exp\{-\alpha z\}}{P_1 exp\{A'(\lambda_1-\lambda)\} + \cdots + P_N exp\{A'(\lambda_N-\lambda)\}} \tag{2}$$

In Eq. (2), $S(0,\lambda)$ is the input power distribution with a total input power ($P_1, \ldots, P_N$) with a total input power (P_0) and wavelengths ($\lambda_1, \cdots, \lambda_N$) while α is the linear wavelength-dependent loss of a fibre with length z. A' is $\beta P_0 L_{eff}$, where β is the fibre's Raman gain coefficient. Eq. (2) is referred to as the general model (GM) in this work.

6.3.1. Experimental setup

Experiments were conducted in Mininet-Optical and the COSMOS programmable testbed to investigate the power dependence of the SRS effect by varying the input power levels of 90 channels. The initial launch power of 0 dBm was decreased by a step size of 3 dB to −15 dBm for each channel. The experimental setup consists of two ROADM units separated by a 50 km single-mode fibre span with no in-line amplifier as illustrated in Fig. 13. 90 C-band channels with a spacing of 50 GHz were launched to the input of the first ROADM. The ROADM has a flat output spectrum. The power level of each channel was then monitored at the input of the second ROADM to examine the Raman-induced power tilt on the 90channels after the span. The end-to-end tilt is ascertained by applying a linear fit to the spectrum.

6.3.2. Experimental results

The tilt values' predictions of the rectangular model strongly correlate with the experimental measurements as reported in Table 1. It achieved high accuracy, as indicated by an MAE of 0.06 dB for the first three input power levels. The SRS effect is negligible at low power levels. Hence, the tilt observed in the spectrum is due to the wavelength-dependent loss of the fibre, which is small. The 0 dBm per channel launch power experimental measurement is shown in Fig. 14a with the predictions of both the rectangular and general models to compare with the line of best fit. The experiment was also conducted with a 25 km fibre span, and a similar accuracy level was achieved by the rectangular model.

Various non-uniform channel configurations were examined as presented in Table 2 with the rectangular and general models defined above. The largest tilt estimation error is 0.24 dB and 0.19 dB at 26 channels for the RM and GM models, respectively; corresponding to a small number of channels spread across the band (Fig. 14b). Although this error is large from a relative point of view, it is not significant because the absolute tilt is small for these cases and is likely impacted by the 0.1 dB measurement resolution. The general form of the model achieved a better prediction for the random loading configurations with an MAE of 0.09 dB for all the cases reported in comparison to the rectangular model's prediction with an MAE of 0.14 dB. The Raman tilt was more significant as the number of launched signals increased, reaching a value of 0.94 dB for the measured and a prediction of 1.02 dB for the 84 randomly loaded channels using the general model. This is in accordance with the observations reported in [76].

7. Conclusion

Optical transmission systems have evolved over the past decades to provide unprecedented capacity levels, fuelled in part by recirculating loop emulation methods that enable the study of transmission impairments. The transition to software-defined and disaggregated systems is motivating research on network control and flexibility. Various experimentation techniques have been investigated to study the interaction of such control systems and transmission impairments, but have not achieved the utility of recirculating loops. Recently, city-scale testbeds have emerged which offer the necessary scale. The addition of topology reconfiguration through space switching, comb source channel loading, dual use SDN control, and digital twins is presented as a promising network emulation platform that can allow for research on the interaction of transmission impairments with network controls. The construction of such a reconfigurable topology testbed is described using the COSMOS city-scale testbed with particular attention to the digital twin development. This work presents analytic and DNN digital twin models for an EDFA in the COSMOS testbed, and the results show a trade-off between accuracy and computation time. Characterisation of analytical models for the wavelength-dependent losses in the transmission fibre including SRS effects is also presented. These are key blocks to a full reconfigurable topology testbed to enable a new generation of optical network control experiments that include transmission impairments at scale.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

This publication has emanated from research supported in part by Science Foundation Ireland under Grant #13/RC/20 77_P2 and NSF grants CNS-1827923, OAC-2029295, and CNS-2112562, a Google Research Scholar Award, and an IBM Academic Award.

References

[1] N.S. Bergano, J. Aspell, C. Davidson, P. Trischitta, B. Nyman, F. Kerfoot, Bit error rate measurements of 14000 km 5 Gbit/s fibre amplifier transmission system using circulating loop, Electronics Letters 27 (21) (1991) 1889–1890.

[2] G. Rademacher, R. S. Luís, B. J. Puttnam, R. Ryf, S. Van Der Heide, T. A. Eriksson, N. K. Fontaine, H. Chen, R.-J. Essiambre, Y. Awaji, et al., 172 Tb/s C+ L band transmission over 2040 km strongly coupled 3-core fiber, in: Optical Fiber Communication Conference, Optical Society of America, 2020, pp. Th4C–5.

[3] N.S. Bergano, C. Davidson, Circulating loop transmission experiments for the study of long-haul transmission systems using erbium-doped fiber amplifiers, Journal of Lightwave Technology 13 (5) (1995) 879–888.

[4] D. Kilper, K. Bergman, V.W. Chan, I. Monga, G. Porter, K. Rauschenbach, Optical networks come of age, Optics and Photonics News 25 (9) (2014) 50–57.

[5] D. C. Kilper, S. Zhu, J. Yu, Physical layer control for disaggregated optical systems, in: Asia Communications and Photonics Conference, Optical Society of America, 2018, pp. Su1J–1.

[6] V. López, B. Huiszoon, J. Fernández-Palacios, O. González de Dios, J. Aracil, Path computation element in telecom networks: Recent developments and standardization activities, in: 2010 14th Conference on Optical Network Design and Modeling (ONDM), 2010, pp. 1–6. doi:10.1109/ONDM.2010.5431606.

[7] E. Riccardi, P. Gunning, Ó.G. de Dios, M. Quagliotti, V. López, A. Lord, An operator view on the introduction of white boxes into optical networks, Journal of Lightwave Technology 36 (15) (2018) 3062–3072, https://doi.org/10.1109/JLT.2018.2815266.

[8] M. Filer, M. Cantono, A. Ferrari, G. Grammel, G. Galimberti, V. Curri, Multi-vendor experimental validation of an open source QoT estimator for optical networks, Journal of Lightwave Technology 36 (15) (2018) 3073–3082.

[9] A. Sgambelluri, J.-L. Izquierdo-Zaragoza, A. Giorgetti, L. Gifre, L. Velasco, F. Paolucci, N. Sambo, F. Fresi, P. Castoldi, A.C. Piat, R. Morro, E. Riccardi, A. D'Errico, F. Cugini, Fully disaggregated ROADM white box with NETCONF/YANG control, telemetry, and machine learning-based monitoring, in, Optical Fiber Communications Conference and Exposition (OFC) 2018 (2018) 1–3.

[10] F. Musumeci, C. Rottondi, A. Nag, I. Macaluso, D. Zibar, M. Ruffini, M. Tornatore, An overview on application of machine learning techniques in optical networks, IEEE Communications Surveys Tutorials 21 (2) (2019) 1383–1408, https://doi.org/10.1109/COMST.2018.2880039.

[11] E. Seve, J. Pesic, Y. Pointurier, Accurate QoT estimation by means of a reduction of EDFA characteristics uncertainties with machine learning, in, International Conference on Optical Network Design and Modeling (ONDM) 2020 (2020) 1–3, https://doi.org/10.23919/ONDM48393.2020.9133020.

[12] M. Sena, R. Emmerich, B. Shariati, J.K. Fischer, R. Freund, Link tomography for amplifier gain profile estimation and failure detection in C+L-band open line systems, in, Optical Fiber Communications Conference and Exhibition (OFC) 2022 (2022) 1–3.

[13] A. Lord, The future of optical transport: Architectures and technologies from an operator perspective, in: Optical Fiber Communication Conference, Optica Publishing Group, 2022, pp. W4F–W11.

[14] DTC Innovation Forum, Digital twin computing (white paper), , Accessed: 2022-03-29 (2019). URL https://iowngf.org/white-papers/.

[15] F. Tao, H. Zhang, A. Liu, A.Y.C. Nee, Digital twin in industry: State-of-the-art, IEEE Transactions on Industrial Informatics 15 (4) (2019) 2405–2415, https://doi.org/10.1109/TII.2018.2873186.

[16] D. Simeonidou, Bristol is open, in: 5G Radio Technology Seminar. Exploring Technical Challenges in the Emerging 5G Ecosystem, 2015, pp. 1–32. doi:10.1049/ic.2015.0035.

[17] Y. Teranishi, Y. Saito, S. Murono, N. Nishinaga, JOSE: An open testbed for field trials of large-scale IoT services, Journal of the National Institute of Information and Communications Technology 62 (2) (2016) 151–159.

[18] J. Yu, T. Chen, C. Gutterman, S. Zhu, G. Zussman, I. Seskar, D. Kilper, COSMOS: Optical architecture and prototyping, in,, Optical Fiber Communications Conference and Exhibition (OFC), IEEE 2019 (2019) 1–3.

[19] C. Gutterman, A. Minakhmetov, J. Yu, M. Sherman, T. Chen, S. Zhu, I. Seskar, D. Raychaudhuri, D. Kilper, G. Zussman, Programmable optical x-haul network in the COSMOS testbed, in: In: 2019 IEEE 27th International Conference on Network Protocols (ICNP), IEEE, 2019, pp. 1–2.

[20] D. Raychaudhuri, I. Seskar, G. Zussman, T. Korakis, D. Kilper, T. Chen, J. Kolodziejski, M. Sherman, Z. Kostic, X. Gu, et al., Challenge: COSMOS: A city-scale programmable testbed for experimentation with advanced wireless, in: Proceedings of the 26th Annual International Conference on Mobile Computing and Networking, 2020, pp. 1–13.

[21] D. Kilper, J. Yu, S. Santaniello, Optical networking in smart city and wireless future networks platforms, in: In: 2021 European Conference on Optical Communication (ECOC), IEEE, 2021, pp. 1–4.

[22] S. Chandrasekhar, D. Kilper, Using testbeds for optically-transparent mesh network experimentation, in: LEOS 2006–19th Annual Meeting of the IEEE Lasers and Electro-Optics Society, IEEE, 2006, pp. 771–772.

[23] D. C. Kilper, S. Chandrasekhar, F. Smyth, L. P. Barry, Dynamic circulating-loop methods for transmission experiments in optically transparent networks, in: 2008 10th Anniversary International Conference on Transparent Optical Networks, Vol. 1, 2008, pp. 99–102. doi:10.1109/ICTON.2008.4598380.

[24] D. C. Kilper, D. Bayart, S. Chandrasekhar, A. Morea, S. K. Korotky, F. Leplingard, Mesh network transport experiments using a distributed-distance circulating loop, in: 2008 34th European Conference on Optical Communication, 2008, pp. 1–2. doi:10.1109/ECOC.2008.4729352.

[25] M. D. Feuer, D. C. Kilper, S. L. Woodward, ROADMs and their system applications, in: Optical Fiber Telecommunications VB, Elsevier, 2008, pp. 293–343.

[26] D. C. Kilper, S. Chandrasekhar, E. Burrows, L. L. Buhl, J. Centanni, Local dispersion map deviations in metro-regional transmission investigated using a dynamically re-configurable re-circulating loop, in: OFC/NFOEC 2007 - 2007 Conference on Optical Fiber Communication and the National Fiber Optic Engineers Conference, 2007, pp.1–3. doi:10.1109/OFC.2007.4348705.

[27] A. Gnauck, P. Winzer, G. Raybon, M. Schnecker, P. Pupalaikis, 10 X 224-Gb/s WDM transmission of 56-Gbaud PDM-QPSK signals over 1890 km of fiber, IEEE Photonics Technology Letters 22 (13) (2010) 954–956, https://doi.org/10.1109/LPT.2010.2048100.

[28] M. Kong, K. Wang, J. Ding, J. Zhang, W. Li, J. Shi, F. Wang, L. Zhao, C. Liu, Y. Wang, W. Zhou, J. Yu, 640-Gbps/carrier WDM transmission over 6,400 km based on PS-16QAM at 106 Gbaud employing advanced DSP, Journal of Lightwave Technology 39 (1) (2021) 55–63, https://doi.org/10.1109/JLT.2020.3024771.

[29] M.R. Phillips, D.M. Ott, Crosstalk due to optical fiber nonlinearities in WDM CATV lightwave systems, Journal of Lightwave Technology 17 (10) (1999) 1782.

[30] J. Junio, D.C. Kilper, V.W. Chan, Channel power excursions from single-step channel provisioning, Journal of Optical Communications and Networking 4 (9) (2012) A1–A7.

[31] A.S. Ahsan, C. Browning, M.S. Wang, K. Bergman, D.C. Kilper, L.P. Barry, Excursion-free dynamic wavelength switching in amplified optical networks, Journal of Optical Communications and Networking 7 (9) (2015) 898–905.

[32] Y. Huang, W. Samoud, C. L. Gutterman, C. Ware, M. Lourdiane, G. Zussman, P. Samadi, K. Bergman, A machine learning approach for dynamic optical channel add/drop strategies that minimize EDFA power excursions, in: ECOC 2016; 42nd European Conference on Optical Communication, VDE, 2016, pp. 1–3.

[33] K. Ishii, J. Kurumida, S. Namiki, Experimental investigation of gain offset behavior of feedforward-controlled WDM AGC EDFA under various dynamic wavelength allocations, IEEE Photonics Journal 8 (1) (2016) 1–13.

[34] M. Freire-Hermelo, D. Sengupta, A. Lavignotte, C. Tremblay, C. Lepers, Reinforcement learning for compensating power excursions in amplified wdm systems, Journal of Lightwave Technology 39 (21) (2021) 6805–6813.

[35] D. C. Kilper, C. A. White, Amplifier issues for physical layer network control, in: Optically Amplified WDM Networks, Elsevier, 2011, pp. 221–251.

[36] M. Berman, J.S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, I. Seskar, GENI: A federated testbed for innovative network experiments, Computer Networks 61 (2014) 5–23.

[37] L. Liu, W.-R. Peng, R. Casellas, T. Tsuritani, I. Morita, R. Martínez, R. Muñoz, S. J. Ben Yoo, Experimental demonstration of OpenFlow-based dynamic restoration in elastic optical networks on GENI testbed, in: 2014 The European Conference on Optical Communica-tion (ECOC), 2014, pp. 1–3. doi:10.1109/ECOC.2014.6964202.

[38] I. Baldin, A. Nikolich, J. Griffioen, I.I.S. Monga, K.-C. Wang, T. Lehman, P. Ruth, Fabric: A national-scale programmable experimental network infrastructure, IEEE Internet Computing 23 (6) (2019) 38–47, https://doi.org/10.1109/MIC.2019.2958545.

[39] P. Sotres, J.R. Santana, L. Sánchez, J. Lanza, L. Muñoz, Practical lessonsfromthedeploymentandmanagementofasmartcityInternetof-Things

infrastructure: The SmartSantander testbed case, IEEE Access 5 (2017) 14309–14322, https://doi.org/10.1109/ACCESS.2017.2723659.
[40] J. Struye, B. Braem, S. Latré, J. Marquez-Barja, The CityLab testbed—Large-scale multi-technology wireless experimentation in a city environment: Neural network-based interference prediction in a smart city, in: IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS), IEEE, 2018, pp. 529–534.
[41] A. Ersoy, Smart cities as a mechanism towards a broader understanding of infrastructure interdependencies, Regional Studies, Regional Science 4 (1) (2017) 26–31.
[42] G.C. Papen, R.E. Blahut, Lightwave communications, Cambridge University Press, 2019.
[43] C. Kiss Kalló, J. Shields, C. O'Malley, R. Carley, V. Lopez, J. P. Fernández-Palacios, Cost-effective sub-wavelength solution for data centre location in scaled next-generation networks, in: 2012 16th International Conference on Optical Network Design and Modelling (ONDM), 2012, pp. 1–6. doi:10.1109/ONDM.2012.6210262.
[44] M.J. O'Mahony, D. Simeonidou, D.K. Hunter, A. Tzanakaki, The application of optical packet switching in future communication networks, IEEE Communications magazine 39 (3) (2001) 128–135.
[45] Y. Yoshida, A. Maruta, K.-I. Kitayama, M. Nishihara, T. Tanaka, T. Takahara, J. C. Rasmussen, N. Yoshikane, T. Tsuritani, I. Morita, et al., SDN-based network orchestration of variable-capacity optical packet switching network over programmable flexi-grid elastic optical path network, Journal of Lightwave Technology 33 (3) (2015) 609–617.
[46] J. Mata, I. de Miguel, R.J. Duran, N. Merayo, S.K. Singh, A. Jukan, M. Chamania, Artificial intelligence (AI) methods in optical networks: A comprehensive survey, Optical switching and networking 28 (2018) 43–57.
[47] D.J. Lambert, C.H. Joyner, J. Rossi, F.A. Kish, R. Nagarajan, S. Grubb, M.F. Van Leeuwen, M. Kato, J.L. Pleumeekers, A. Mathur, et al., Large-scale photonic integrated circuits used for ultra long haul transmission, in: LEOS 2007-IEEE Lasers and Electro-Optics Society Annual Meeting Conference Proceedings, IEEE, 2007, pp. 778–779.
[48] R. Nagarajan, M. Kato, S. Corzine, P. Evans, C. Joyner, R. Schneider, F. Kish, D. Welch, Monolithic, multi-channel DWDM transmitter photonic integrated circuits, in,, IEEE 21st International Semiconductor Laser Conference, IEEE 2008 (2008) 5–6.
[49] A. Teixeira, L. Costa, G. Franzl, S. Azodolmolky, I. Tomkos, K. Vlachos, S. Zsigmond, T. Cinkler, G. Tosi-Beleffi, P. Gravey, et al., An integrated view on monitoring and compensation for dynamic optical networks: from management to physical layer, Photonic Network Communications 18 (2) (2009) 191–210.
[50] Y. Li, D.C. Kilper, Optical physical layer SDN, Journal of Optical Communications and Networking 10 (1) (2018) A110–A121.
[51] V. Lopez, R. Vilalta, V. Uceda, A. Mayoral, R. Casellas, R. Martinez, R. Munoz, J. P. Fernandez Palacios, Transport API: A solution for SDN in carriers networks, in: ECOC 2016; 42nd European Conference on Optical Communication, 2016, pp. 1–3.
[52] R. Vilalta, R. Muñoz, R. Casellas, R. Martínez, V. López, O.G. de Dios, A. Pastor, G. P. Katsikas, F. Klaedtke, P. Monti, A. Mozo, T. Zinner, H. Øverby, S. Gonzalez-Diaz, H. Lønsethagen, J.-M. Pulido, D. King, in: TeraFlow: Secured autonomic Traffic Management for a Tera of SDN Flows, Summit (EuCNC/6G Summit), 2021, pp. 377–382, https://doi.org/10.1109/EuCNC/6GSummit51104.2021.9482469.
[53] R. Vilalta, R. Casellas, L. Gifre, R. Muñoz, R. Martínez, A. Pastor, D. López, J. Fernández-Palacios, Architecture to deploy and operate a digital twin optical network, in, Optical Fiber Communications Conference and Exhibition (OFC) 2022 (2022) 1–3.
[54] V.D. Chemalamarri, P. Nanda, K.F. Navarro, SYMPHONY - a controller architecture for hybrid software defined networks, in, Fourth European Workshop on Software Defined Networks 2015 (2015) 55–60, https://doi.org/10.1109/EWSDN.2015.61.
[55] A. Aguado, V. López, J. Marhuenda, Ó.G. de Dios, J.P. Fernández Palacios, Abno,, A feasible SDN approach for multivendor IP and optical networks, Journal of Optical Communications and Networking 7 (2) (2015) A356–A362.
[56] R. Vilalta, J. L. de la Cruz, A. M. López-de Lerma, V. López, R. Martínez, R. Casellas, R. Muñoz, uABNO: A cloud-native architecture for optical SDN controllers, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC), 2020, pp. 1–3.
[57] ONOS open network operating system (ONOS), https://docs. onosproject.org/, accessed: 2022-04-02 (2022).
[58] J. Pesic, Missing pieces currently preventing effective application of machine learning to QoT estimation in the field, in, Optical Fiber Communications Conference and Exhibition (OFC) 2021 (2021) 1–3.
[59] A. Ferrari, K. Balasubramanian, M. Filer, Y. Yin, E. Le Rouzic, J. Kundrát, G. Grammel, G. Galimberti, V. Curri, Assessment on the in-field lightpath QoT computation including connector loss uncertainties, Journal of Optical Communications and Networking 13 (2) (2021) A156–A164.
[60] D.J. Elson, L. Galdino, R. Maher, R.I. Killey, B.C. Thomsen, P. Bayvel, High spectral density transmission emulation using amplified spontaneous emission noise, Opt. Lett. 41 (1) (2016) 68–71, https://doi.org/10.1364/OL.41.000068.
[61] D.J. Elson, G. Saavedra, K. Shi, D. Semrau, L. Galdino, R. Killey, B.C. Thomsen, P. Bayvel, Investigation of bandwidth loading in optical fibre transmission using amplified spontaneous emission noise, Optics express 25 (16) (2017) 19529–19537.
[62] X. Liu, Challenges and opportunities in future high-capacity optical transmission systems, in: Optically Amplified WDM Networks, Elsevier, 2011, pp. 47–82.
[63] A.S. Thyagaturu, A. Mercian, M.P. McGarry, M. Reisslein, W. Kellerer, Software defined optical networks (SDONs): A comprehensive survey, IEEE Communications Surveys & Tutorials 18 (4) (2016) 2738–2786.
[64] J. Yu, C. Gutterman, A. Minakhmetov, M. Sherman, T. Chen, S. Zhu, G. Zussman, I. Seskar, D. Kilper, Dual use SDN controller for management and experimentation in a field deployed testbed, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC), IEEE 2020 (2020) 1–3.
[65] D. Wang, Z. Zhang, M. Zhang, M. Fu, J. Li, S. Cai, C. Zhang, X. Chen, The role of digital twin in optical communication: fault management, hardware configuration, and transmission simulation, IEEE Communications Magazine 59 (1) (2021) 133–139, https://doi.org/10.1109/MCOM.001.2000727.
[66] T. Chen, J. Yu, A. Minakhmetov, C. Gutterman, M. Sherman, S. Zhu, S. Santaniello, A. Biswas, I. Seskar, G. Zussman, D. Kilper, A software-defined programmable testbed for beyond 5g optical-wireless experimentation at city-scale, IEEE Network 36 (2) (2022) 90–99, https://doi.org/10.1109/MNET.006.2100605.
[67] D.A. Fishman, D.L. Correa, E.H. Goode, T.L. Downs, A.Y. Ho, A. Hale, P. Hofmann, B. Basch, S. Gringeri, The rollout of optical networking: LambdaXtreme® national network deployment, Bell Labs Technical Journal 11 (2) (2006) 55–63, https://doi.org/10.1002/bltj.20161.
[68] B. Lavigne, F. Leplingard, L. Lorcy, E. Balmefrezol, J. C. Antona, T. Zami, D. Bayart, Method for the determination of a quality-of-transmission estimator along the lightpaths of partially transparent networks, in: 33rd European Conference and Exhibition of Optical Communication, 2007, pp. 1–2. doi:10.1049/ic:20070307.
[69] A.A. Díaz-Montiel, B. Lantz, J. Yu, D. Kilper, M. Ruffini, Real-time QoT estimation through SDN control plane monitoring evaluated in mininet-optical, IEEE Photonics Technology Letters 33 (18) (2021) 1050–1053.
[70] Y. Liu, X. Liu, L. Liu, Y. Zhang, M. Cai, L. Yi, W. Hu, Q. Zhuge, Modeling EDFA gain: Approaches and challenges, in: Photonics, Vol. 8, Multidisciplinary Digital Publishing Institute, 2021, p. 417.
[71] J. Yu, S. Zhu, C.L. Gutterman, G. Zussman, D.C. Kilper, Machine-learning-based EDFA gain estimation, Journal of Optical Communications and Networking 13 (4) (2021) B83–B91.
[72] W. Mo, C.L. Gutterman, Y. Li, S. Zhu, G. Zussman, D.C. Kilper, Deep-neural-network-based wavelength selection and switching in ROADM systems, Journal of optical communications and networking 10 (10) (2018) D1–D11.
[73] S. Zhu, C.L. Gutterman, W. Mo, Y. Li, G. Zussman, D.C. Kilper, Machine learning based prediction of erbium-doped fiber WDM line amplifier gain spectra, in: In: 2018 European Conference on Optical Communication (ECOC), IEEE, 2018, pp. 1–3.
[74] S. Zhu, C. Gutterman, A.D. Montiel, J. Yu, M. Ruffini, G. Zussman, D. Kilper, Hybrid machine learning EDFA model, in: Optical Fiber Communication Conference, Optical Society of America, 2020, p. T4B.
[75] M. Zirngibl, Analytical model of Raman gain effects in massive wavelength division multiplexed transmission systems, Electronics Letters 34 (8) (1998) 789–790.
[76] S. Bigo, S. Gauchard, A. Bertaina, J.-P. Hamaide, Experimental investigation of stimulated Raman scattering limitation on WDM transmission over various types of fiber infrastructures, IEEE Photonics Technology Letters 11 (6) (1999) 671–673.

Emmanuel Akinrintoyo (Student Member, IEEE) received a B. Eng. degree in Computer and Communications engineering from Technological University Dublin (TUD), Ireland in 2021. He is currently pursuing a master's degree in Electronic Information Engineering at Trinity College Dublin (TCD), Ireland as an E3 scholar. His research interest lies in the application of machine learning and artificial intelligence to achieve greater efficiency in domains such as optical networks control.

Zehao Wang (Student Member, IEEE) received the B.Eng. degree in information science from Zhejiang University, China, in 2019, and the M.S. degree in electrical and computer engineering from Carnegie Mellon University in 2020. He is currently pursuing the Ph.D. degree with Department of Electrical and Computer Engineering at Duke University. His research interest is in the area of next-generation wireless and optical networks. He was a recipient of Zhejiang State Scholarship and MobiCom'21 travel grant.

Bob Lantz is technical lead for Mininet-Optical, part of the COSM-IC project at Columbia University. In addition to starting the Mininet network emulator project, he helped to develop the OpenFlow protocol specification at Stanford University, as well as SDN systems including OpenFlow, ONOS (Open Network Operating System) and FAUCET, working for Arista Networks, DOCOMO Communication Laboratories, Open Networking Foundation, and Google. He received the ACM SIGCOMM Software Systems Award and the ACM SIGCOMM Test of Time Award, and holds a Ph.D. in Electrical Engineering from Stanford University.

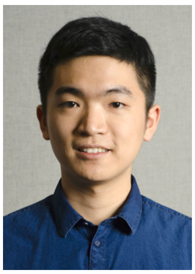

Tingjun Chen received the Ph.D. degree in electrical engineering from Columbia University in 2020, and the B.Eng. degree in electronic engineering from Tsinghua University in 2014. He is an Assistant Professor in the Department of Electrical and Computer Engineering at Duke University and was a postdoctoral associate at Yale University from 2020–2021. His research interests are in the area of networking and communications with a specific focus on next-generation wireless networks and Internet-of-Things systems. He received the Google Research Scholars Award, the IBM Academic Award, the Facebook Fellowship, the Wei Family Private Foundation Fellowship, the Columbia Engineering Morton B. Friedman Memorial Prize for Excellence, and the Columbia University Eli Jury Award and Armstrong Memorial Award. He also received ACM CoNEXT'16 Best Paper Award and the ACM MobiHoc'19 Best Paper Award finalist, and his Ph.D. thesis received the ACM SIGMOBILE Dissertation Award Runner-up.

Dan Kilper is Professor of Future Communication Networks at Trinity College Dublin and Director of the CONNECT Center. Recently he was a research professor in the College of Optical Sciences and Electrical and Computer Engineering at the University of Arizona, Tucson. He holds an adjunct faculty position at the Columbia University Data Science Institute and was a faculty appointee at NIST in 2020. He received M.S. (1992) and Ph.D. (1996) degrees in Physics from the University of Michigan. From 2000–2013, he was a member of technical staff at Bell Labs, and he co-founded Palo Verde Networks, a startup developing optical technologies for AI-controlled software-defined optical networks. He holds thirteen patents and authored six book chapters and more than one hundred and sixty peer-reviewed publications. His research is aimed at solving fundamental and real-world problems in communication networks, addressing interdisciplinary challenges for smart cities, sustainability, and digital equity.