Optical Physical Layer SDN: Enabling Physical Layer Programmability through Open Control Systems

Daniel C. Kilper and Yao Li
University of Arizona, College of Optical Sciences, 1630 E. University Blvd., Tucson, AZ 85721
dkilper@optics.arizona.edu, yaoli@optics.arizona.edu

Abstract: Software defined networking in the optical physical layer is complicated by transmission control used to both optimize performance and stabilize optical signals across multiple nodes. Different approaches are emerging to address these problems.
OCIS codes: (060.2330) Fiber optics communications; (060.4264) Networks, wavelength assignment

1. Introduction

Traffic in the Internet continues to increase and optical and electronic interfaces have been scaling to keep pace, following historical trends. Recently however, changes are occurring on several fronts that are motivating a new role for optical systems in the Internet. Perhaps the most significant change is the large diversity of applications for Internet or data network services, in particular the emergence of applications that call for capacity at optical data rates of 10 Gb/s and higher. Furthermore, cloud computing and the move to virtualize business functions in the cloud create a need for more than just bandwidth access, but instead networks as a service, connecting multiple enterprise locations with the quality of service and features important for specific business requirements. Another trend is specialized applications that also call for optical data rates. For example, big science applications are looking to move petabyte data sets across the Internet, which would require hours or even days at 100 Gb/s dedicated transmission rates. Other government and educational institutions have similar requirements as well as industry segments such as the entertainment industry. Moreover, 5G wireless, which is expect to be standardized in 2020, is calling for peak access rates for individual users of 1 Gb/s and peak access densities of 10 $Tb/s/km^2$, which corresponds to the capacity of a long haul optical system in a 1 km^2 area (www.ngmn.org). In this regime, optical systems are no longer fat pipes that can be treated like plumbing—carrying aggregated traffic of 10's of thousands of individual data flows. In fact, metro only traffic is growing at a much faster rate than long haul traffic and is estimated to have surpassed long haul traffic in total volume in North America. The need for more bandwidth at the edges of the networks, has brought attention to metro optical systems where data is less aggregated and subject to larger variations in location and time [1]. All of these trends point to a changing role for optical systems in which they are more responsive to changing traffic patterns and application demands. Applications or services that involve fewer users are also more cost sensitive. As a consequence, there is growing interest in optical systems that can operate dynamically and allow for sharing in order to be cost effective.

Optical systems are undergoing changes in order to meet the demands of this new environment. Integrated photonics, which involves the development of optical chips similar to micro-electronics, is rapidly changing the cost equation for optics. Several initiatives are underway to enable higher volume manufacturing of optics, overcoming packaging complexity and cost barriers, to put optical components on an integration cost curve (e.g. www.aimphotonics.com). Examples of integrated components have increasingly shown up in optical systems. The development of low cost transceivers has been accelerated by the massive quantities of data transported within data centers and this is beginning to translate over to optical transmission between data centers. Data center interconnect (DCI) networks, purpose built to connect data centers across a wide area, in particular are moving data center transceiver cost and density requirements to the transmission space, resulting in new low cost coherent transceivers. The move to coherent transmission also enables more advanced functionality for optical systems because a coherent transceiver is able to automatically compensate many transmission impairments such as group velocity dispersion, thus simplifying transmission design. In addition, the modulation flexibility of coherent transceivers adds a degree of freedom that can be useful in many applications.

Software control is another area in which optical systems are undergoing changes. Software defined networking (SDN) has emerged as a valuable tool to enable greater control and customization of network operation, starting with the higher network layers, Ethernet and IP switching. Protocols such as Openflow allow user customization of the switch routing tables and other capabilities. Extensions of Openflow have included capabilities for controlling optical transmission hardware. Network operating systems have emerged as a control framework to orchestrate unique or programmable functionality across all of the network layers [1]. The control and tuning of a wide range of optical system parameters and features have been investigated using the new SDN control capabilities for optical systems, including wavelength routing and transceiver bandwidth adaptation. In fact, much of the SDN work is a natural extension of decades of optical networking research. Moving these capabilities into commercial systems, however, is complicated by the operational control requirements of transmission in the physical layer. Quality of service requirements demand that physical layer operations on one signal do not impact the performance of other WDM signals. Getting the highest performance out of a system often requires careful balancing of impairments. While the priority has been capacity and performance, i.e. transmission distance, there has not been a willingness to give up performance or add cost in order to realize new functionality. Real time operations in a transmission system have also been fraught with examples of instability and crosstalk. Here we describe these technical challenges to physical layer SDN control, review past work to address these challenges, and discuss new directions for research.

2. Physical Layer Control

Wavelength division multiplexed (WDM) optical systems were first developed for high-speed, long-haul transmission carrying highly aggregated traffic. They were appropriately referred to as the "plumbing of the Internet," as optical circuits (light paths) or wavelengths were not intended to be disturbed after the initial provisioning. Manual tuning and adjustments were required for provisioning optical circuits with these early-generation WDM systems. Today WDM systems are interconnected in large mesh networks. The configuration of these systems at provisioning time is largely automated and the system designs allow for flexible routing and wavelength assignment. Specifically, Colorless, Directionless, Contentionless (CDC) Reconfigurable Optical Add/Drop Multiplexers (ROADMs) offer increased flexibility. However, provisioning optical circuits through these networks still requires extensive simulation and testing of engineering rules. Controls are applied slowly and methodically to avoid disrupting existing traffic, and signals are often left to soak for long periods in order to ensure proper operation before handling live traffic [2]. The software interface for optical systems is a combination of network and element management tools with other custom and proprietary tools developed by the system vendors for the network operators, increasingly adapted to include SDN interfaces.

Recently different groups have tackled the various aspects of network programmability including programmable devices, resource abstraction, and control protocols. A sliceable, bandwidth-variable transponder (SBVT) has been used as a programmable multi-wavelength source, generating subcarriers with asymmetric channel programmed by an SDN control plane extended to support optical functionality [3]. SDN control has also been applied to bandwidth flexible spatial mode division multiplexing (SDM) networks [4], with a programmable architecture on demand [5], including sliceable and bandwidth variable spatial super-channels.

These results show the tremendous progress in developing programmable optical infrastructure and SDN/NFV control functionality for optical systems and devices working with the higher network layers. This prior work, however, is largely independent of optical transmission system research, which addresses the physics of optical transmission and the design of optical transmission systems [6]. For small networks, as are often studied in SDN research, most transmission effects are negligible or can be easily avoided through choice of configuration. In commercial systems, however, all of the complexity and engineering in transmission system design must be taken into account so that the system can be reliable and perform to specifications [2]. A programmable function that causes the system to become unstable of course is not viable. Even one that disrupts traffic or degrades performance would not be acceptable. Large international research projects have addressed control software for on-demand functionality in optical systems, but have not addressed the physical layer system control: how to tune the amplifiers and switches to maintain error free performance. While many SDN studies include path computation elements (PCE) or quality of transmission (QoT) estimators in their control operations, they do not implement and stress the system in optical hardware in a way to would test the effectiveness and reliability of the optical transmission related control aspects, or they rely on software emulation instead [7].

The lack of a hardware transmission research component to SDN studies is in part due to limited experimental methods available to study such systems. Optical transmission is traditionally studied using recirculating loop experiments. However, recirculating loops do not use bi-directional connections and in general do not allow for transmission control because the 'recirculation' provides a feedback path that can impact control. Our group has studied the application of recirculating loops in transmission control [8,9]. For example, we developed a technique to study the impact of channel reconfiguration in a recirculating loop by keeping the total number of channels constant to minimize the impact on the constant gain control of the amplifiers. The complexity and limitations of these techniques make them problematic for most applications.

Two early landmark papers first defined the optical power control stability problem for optical systems [10, 11]. They showed that when the optical powers of different channels interact through the optical amplifiers, fiber nonlinearities, and power control elements (e.g. wavelength selective switches) instability can arise. This was first experimentally observed for constant power controlled amplifiers and later for constant gain controlled amplifiers [12]. Note that these phenomena are related to, but independent from optical power transients. Optical power transients caused by the optical amplifier response can be largely mitigated through fast feedforward control and non-linear feedback to the EDFA pump power [13]. Optical power dynamics describe the channel power interactions that occur from a wide variety of sources including the amplifiers and are addressed through individual channel power corrections in, for example, the wavelength selective switch. We showed that synchronizing optical power control adjustments across ROADM nodes based on channel power interactions alleviate power dynamics, but requires time consuming step by step control [12]. Recently several groups have investigated techniques to address power excursions and the potential instabilities that can occur. These include machine learning algorithms for wavelength selection [14] and a wavelength assignment algorithm based on a model of the channel power interactions [15]. We have recently studied hardware based stabilization methods using fast tunable lasers, which provide a new level of control that can be used for WDM reconfiguration [16].

3. Open Control Systems

Just as with most problems in optical transmission research, solutions are a trade-off of cost and performance [6]. For example, an opaque transmission system with no optical switches can solve all problems related to software control—essentially reducing optical systems to electronic systems with point to point optical interconnects. This approach, however, is very expensive and not scalable in the wide area, as has been proven out by the marketplace. Therefore, a balance is needed with many potential solutions depending on the network application, preferred performance metrics, and technological approach. An effective way to investigate this balance is to open optical system control to research so that solutions can be developed for different applications.

Two directions can be identified for introducing SDN control into the optical physical layer. The first we identify as the bare metal whitebox approach, which is to remove all system control with the exception of the 'drivers' within the network elements that operate the element functions. The system control is then entirely implemented in an SDN controller that performs all operational functions in the system. Internal node controls may also be supported such as amplifier constant gain control or channel power leveling running on a wavelength selective switch. All node coordination operations in the system are carried out through the SDN controller. This approach has been used in a new industry multi-source agreement (MSA) initiative openROADM (www.openroadm.org).

The second approach is a brightbox model in which the optical system includes a control or operating system that facilitates the SDN control of the system, managing aspects of the physical layer control complexity. Effectively the proprietary software control system can be turned into an SDN controller and make use of abstraction to communicate and work with other SDN controllers or network operating systems.

The bare metal approach has the advantage in that each ROADM node is independent and can be sourced from different vendors on a component-by-component basis. However, one loses the ability to implement multi-node control operations using in-band communication, which can enable faster and more scalable control. This drawback could be overcome through standardization of in-band communication methods or other node to node capabilities. In the current MSA implementation, however, all multi-node coordination is through Ethernet interfaces to a centralized SDN controller. Open questions include what level of performance can be achieved, or put in other words, what advantages are given up in moving to a bare metal approach over what scale of network? To what extent additional information such as OSNR estimates obtained from optical performance monitoring could facilitate decisions in SDN based physical layer control is also still an open question.

The authors acknowledge support from NSF #CNS1650669, #EEC0812072, and DOE grant DE-SC0015867.

5. References

[1] G. Parulkar, T. Tofigh, M. DeLeenheer, in Proc. Optical Fiber Communication Conference, 2015, paper W1G.4
[2] S. A. Taegar, D. A Fishman, D. L. Correa, in Proc. Optical Fiber Communication Conference, 2006, paper NTuF4.
[3] N. Sambo, G. Meloni, F. Paolucci, et. al., in European Conference on Optical Communication, 2014, paper We.3.2.3.
[4] N. Amaya, et. al., in 39th European Conference and Exhibition on Optical Communication, 2013, pp. 1-3.
[5] B. Rahimzadeh Rofoee, et. al., Journal of Optical Communications and Networking, vol. 5, no. 9, pp. 1096-1110, Sep. 2013.
[6] M. D. Feuer, D. C. Kilper, and S. Woodward, in *Optical Fiber Communications V B*, eds. I. P. Kaminow, T. Li, and A. E. Willner, (Academic Press, Burlington, MA) 2008.
[7] S. Azodolmolky, et. al. IEEE J. Opt. Commun. Netw., vol. 3, no. 4, pp. 290–299, 2011.
[8] A. Morea, et. al., International Conference on Transparent Optical Networks (ICTON) 2009, Tu.C2.4.
[9] F. Smyth, D. C. Kilper, S. Chandrasekhar, and L. P. Barry, *J. Lightwave Technol.*, vol. 27, no. 21, pp. 4686-4697, 2009.
[10] L. Pavel, *Automatica* 40, pp. 1361–1370 (2004).
[11] D. Gorinevsky, and G. Farber, Journal of Lightwave Technology, vol. 22, no. 10, pp. 2245-2255, Oct. 2004
[12] D. C. Kilper, C. A. White, in *Optical Amplifiers for Advanced Communication Systems and Networks*, eds. J. Zyskind, A. Srivastava, (Academic Press, Burlington, MA) 2011.
[13] C. Tian, and S. Kinoshita, Journal of Lightwave Technology, vol. 21, no. 8, pp. 1728-1734, Aug. 2003
[14] Y. Huang, et al., European Conference on Optical Communication, 2016
[15] K. Ishii, J. Kurumida and S. Namiki, IEEE Photonics Journal, vol. 8, no. 1, pp.1-13, Feb 2016
[16] A. S. Ahsan, et. al., Journal of Optical Communications and Networking, vol. 7, no. 9, 898-905, 2015