MRV | EMPOWERING THE OPTICAL EDGE

White Paper

Optical Software Defined Networking

Introduction

Service providers look for a dynamic, application-aware network infrastructure that suits today's cloud and mobility needs. Software Defined Networking (SDN) and Network Function Virtualization (NFV) have great potential to improve the flexibility and responsiveness of networks, decrease both operating and capital expenses, and add new and innovative features to many different types of communications networks. When applied to packet networks, the benefits are clear, as networks can be flexibly redefined and equipment repurposed quickly and easily to meet changing network needs. However, constructing a dynamic and flexible packet network on top of a rigid optical transport does not yield an optimal solution, and blocks the level of innovative services that could be introduced.

Up until recently there was no other choice, since optical transport couldn't be programed. Fixed wavelength transponders, fixed ROADMs, and fixed, non-optimal assignment of packet services over lightpaths were the only options. However, new advances in optical technology can allow network operators to program the optical layer. Using these programming capabilities, the optical transport layer can be abstracted to a set of shared, common resources, that can be used dynamically and on demand.
Standard bodies like ONF, ITU, and OIF have begun the work to define a common software interface to optical transport SDN, but meanwhile vendors and service providers can begin extending SDN principles to the optical arena to achieve a unified, optimized packet and optical network. These efforts need to be done with care to make sure that they do not result in unstable networks long term and are compatible with emerging standards.

In this white paper, MRV will explain our approach to Optical SDN, the current status of industry standards, and where we believe Optical SDN will add value to customers' networks, both today and in the future. We will introduce the components of Optical SDN that are designed into MRV's new generation of optical transport equipment and show how they can allow customers to leverage advanced, software-based architectures, in order to add flexibility and resource optimization to the optical transport.

Software Defined Networking

Software Defined Networking has evolved from an academic concept to a significant industry driver. With SDN, the control plane for a network is separated from the data plane, allowing much more flexible control of the network and the functions that it performs. This promises to be one of the most significant technology changes ever made to fundamental network design.

In traditional networking, the control plane and the data plane reside in the same hardware. This means that the controller – the software that makes the decisions about how packets are routed through the hardware – is on the same hardware that is actually moving the packets around. Sophisticated communications mechanisms exist to coordinate the local control software so that the network acts as a unified network, but local decision making is still controlled by the local controller.

In software defined networking, the controller is separated from the hardware and can be centralized, or run in a cloud of connected servers. By moving the controller off of the hardware, more flexible network applications can be designed and run enabling new services. For example, a network operator that needs to implement a temporary troubleshooting network in parallel with his existing network in a specific location can software define temporary routing tables for just the troubleshooting data in that area. Other examples include optimizing metro networks based on pre-knowledge of the applications that will be running and redesigning datacenter interconnections based on changing customer demands.

Network Function Virtualization is a technology concept that enables “virtualizing” specific network functions on generic hardware. In traditional networks, if a customer needs to add functionality such as an advanced firewall to their network, the network operator must install physical hardware at that location. With NFV, a network operator could launch the firewall functionality virtually as software running on a server either on the CPE or at the POP and immediately have a firewall in place.

The combination of SDN and NFV have the opportunity to greatly increase network flexibility and enable “network as a service” applications where the network can be designed for the applications that will be running and the needs of the end used. Optical SDN takes these concepts from the packet world and applies them to the underlying optical layer.

Optical SDN

Implementing an SDN-enabled packet network on top of a static optical transport network may result in suboptimal network operations and the inability to implement a full spectrum of possible innovation. For example, setting up dynamic bandwidth between remote locations may be limited by the wavelength routing between those locations. If the wavelength routing could also be dynamically controlled, then significantly more flexibility could be achieved. It makes sense, therefore, to extend SDN principles to the optical transport layer.

The main SDN building blocks of central control, programmability and abstraction should be a part of optical SDN implementation. The fundamental requirement for SDN-enabled optical transport is the ability of every element to be software programmable. Until recently, however, programmability at the optical layer was impossible and therefore optical resources were treated as dumb, fixed pipes. However, due to advanced optical technologies that were developed in recent years that allow us not to treat optical transport as fixed pipe resources, optical networks today can increasingly be software-controlled and programmable.

There have been several significant developments in recent years that combine to bring dynamics and flexibility to optical networking:

- Software-defined optical port speeds, protocols, and wavelengths
- Advanced modulation and detection schemes, especially at 100Gbps and higher speeds
- A flexible wavelength grid rather than a fixed 50GHz grid
- Flexible wavelength routing via dynamic ROADMs
- OTN grooming and switching

Software-defined Optical Port Speeds, Protocols, and Wavelengths

In order for new services to be provisioned or changed on an optical network via software without pre-knowledge of the service type, the physical ports that are attached to the network must be software defined – both in protocol and speed definition as well as at wavelength definition. While packet-based SDN sees the physical layer as merely a supporting layer for the logical network, with optical networking there is no such separation. The physical layer is the network. Therefore, it is the physical layer of optical networks that must be software defined in Optical SDN.

User-definable port speeds, protocols, and wavelengths have been available in optical network hardware for several years, with the level of flexibility varying from simple LAN/WAN Ethernet selection to fully flexible transponders covering a wide range of rates and protocols. Tunable lasers which enable flexible wavelength selection while not quite as inexpensive as fixed wavelength lasers, are now close enough in price to be a viable alternative when network flexibility is desired. However, it is only recently that the surrounding technologies and software infrastructures are being defined to allow protocol and wavelength tunability to become part of an overall software-defined flexible network design.

Advanced Modulation and Detection Schemes
For speeds up to 10Gbps, optical transmitters used a simple on-off keyed (OOK) encoding where the laser was effectively turned on and off very quickly to transmit bits through fiber. This basic encoding mechanism allowed the optical receivers to be relatively simple and cost efficient. When moving to data rates of 40Gbps and higher, however, OOK began to encounter fundamental physics issues that limited the ability of the signals to be cost-effectively transmitted over distance. At very high speed, the optoelectronics themselves have difficulties being turned on and off at very high rates, fiber issues like polarization mode dispersion begin to create significant signal distortion, and the fundamental bandwidth occupied by the OOK signal becomes larger than the bandwidth of a 50GHz ITU wavelength channel. For these reasons, some 40Gbs and nearly all 100Gbps initial implementations used more than one wavelength for transmission.

In order to increase the spectral efficiency and reduce the per-channel electronics speed within the optical transmitters and receivers, more advanced modulations were developed. These modulation schemes, adapted from radio and satellite transmission for use in the optical domain, use both amplitude and phase parameters of the light signal. In addition, advanced receiver technologies were developed that enabled these advanced modulation schemes to operate over longer distances. In particular, coherent detection and advanced digital signal processing have become the standard for high speed optical signal detection (100Gbps and above).

For 100Gbps transmission DP-QPSK modulation was chosen by the OIF as the recommended standard transmission modulation scheme and is the overwhelmingly prevalent 100Gbps scheme in service today. However, other modulation schemes have been demonstrated that can improve the spectral efficiency and the bitrate per clock cycle at the expense of decreasing the optical signal reach. Therefore, it is possible to design a transmitter with a modulation that is the optimal tradeoff between bandwidth and reach for a specific network span: higher bandwidth capacity over shorter distances or lower bandwidth capacity over very long distances.

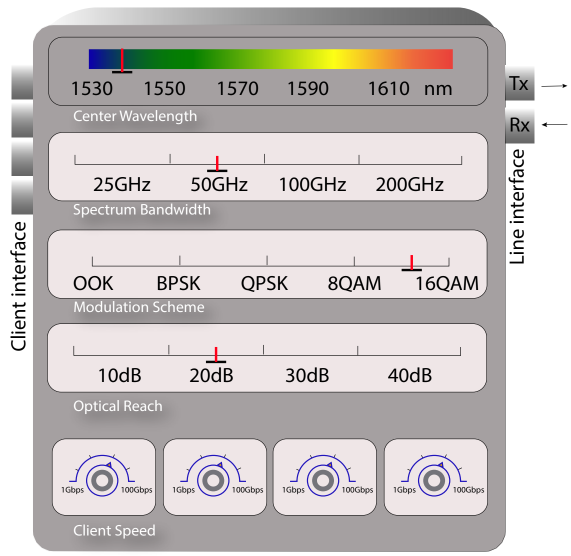

A model of a software-programmable transponder where transmission parameters can be flexibly adjusted based on network demand

The technology now exists to develop a software-programmable transponder where the modulation scheme can be defined in software. As the demands of the network change (e.g. a demand is dynamically reassigned from a metro interconnect to a long distance interconnect), the properties of the transponder (bit rate, optical reach, bandwidth requirement) can be adjusted based on the link length and physical properties of the channel.

Flexible Wavelength Grid
TWhen standards for wavelength division multiplexing were originally put in place by the ITU-T, a rigid grid of wavelengths was defined with strict 100GHz frequency spacing. This rigidity was required to enable interoperability between components vendors, equipment vendors, and network operators. Later the grid was enhanced with definitions at 50GHz as filter and wavelength-locking technology improved.

The latest WDM wavelength grid specification from the ITU-T, G.694.1, has defined a flexible grid with WDM channels spaced at 12.5GHz but with the ability to define an aggregate superchannel spectral width of N × 12.5 GHz. These superchannels can be flexibly defined to accommodate any combination of optical carriers, modulations, and data rates. In addition, the N × 12.5GHz spectral width of a flex-grid superchannel can be tuned, enabling rapid changes when the operating specifications of the superchannel are varied.

[1]– DP-QPSK, which stands for Dual Polarization – Quadrature Phase Shift Keying, uses two polarizations of light as well as four defined phases to modulate the 100Gbps signal. Two phases of light mean that the electronics for each channel run at half the total speed and QPSK modulation allows a very efficient bandwidth profile that fits easily within the 50GHz wavelength grid.

The flexible wavelength grid can be combined with flexible modulation schemes described above to enable a fully dynamic definition of the channel properties between locations. If spectral efficiency is of primary importance in network design, the assigned wavelength channel can be defined as 50GHz (2 x 12.5GHz channels) and the modulation scheme tuned accordingly. If span budget and distance is of primary importance, a wide channel (e.g. 200GHz) can be defined to accommodate a less bandwidth efficient but longer distance modulation scheme.

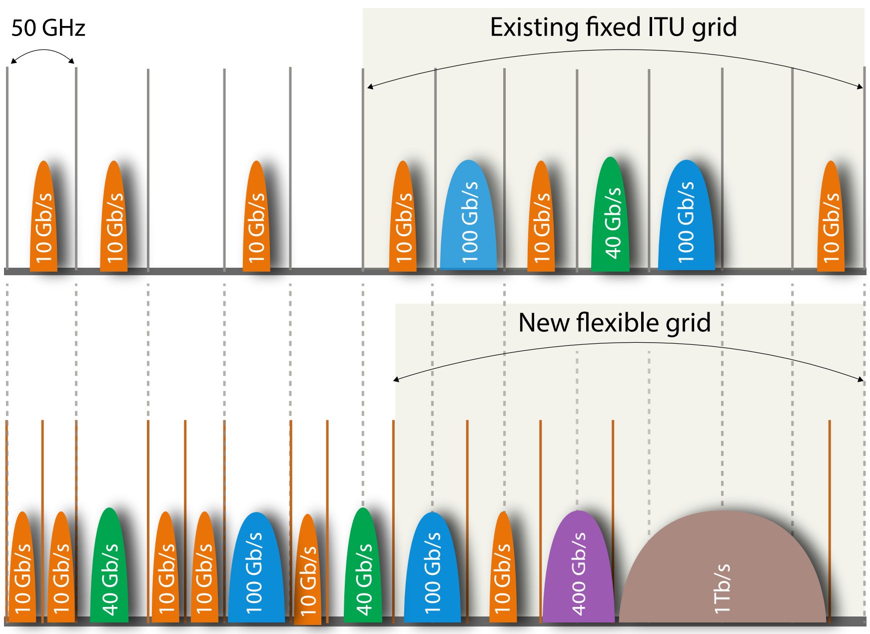

A flexible wavelength grid allows more efficient use of bandwidth for a wider variety of signal bandwidths

Dynamic ROADMs

Until recently, reconfigurable optical add drop multiplexers (ROADMs) had fixed add/drop pairs, with a 1:1 relationship between the transponder and the direction of transmission. Wavelengths could be add/dropped or passed through on demand, but there was little additional flexibility. Current generation ROADMs are much more technically advanced and allow a more much dynamic network design.

Colorless technology in ROADMs means that any wavelength can be assigned to any port, breaking the 1:1 wavelength to port relationship. This means that any wavelength on a truck can be dropped to any one of several client interfaces. Directionless technology further eliminates the 1:1 direction to port relationship, allowing any signal on any one of several trunks to be dropped to any one of several client interfaces. Finally, contentionless technology eliminates the possibility of wavelength blocking within the ROADM due to the same wavelength being utilized on several trunks. The combination of these technologies is referred to as CDC (colorless, directionless, and contentionless) and has enabled a fuly fle

OTN Grooming and Switching

But programmability is only the first step towards an SDN-enabled optical transport. As in SDN for packet networks, it is the SDN-controller and the applications above it that adds the intelligence to a data plane that is responsible for moving the information around the network in a flexible and structured way. In transport networks, the flexible and structured data plane is increasingly based on OTN framing. OTN establishes connection-oriented network similar to legacy protocols such as SONET/SDH, but in a much more packet-friendly way. OTN container sizes are built around standard packet protocols (1.25Gbps, 2.5Gbps, 40Gbps, 100Gbps) and can be flexibly combined to establish intermediate rates as well.

While OTN grooming and switching is not absolutely required to build an optical SDN-based network, having OTN adds another layer of flexibility to the overall network design. Circuits that are sub-wavelength in size can be managed at the transport layer, giving network operators a much more service-oriented transport network. At the edge of the network, this capability may be used merely for service adaptation (turning a variety of protocols into one standard protocol) or multiplexing, but in the metro and core, OTN switching is increasingly seen as a necessary component to the optical transport network.

Optical SDN Standards Bodies and Industry Groups

There are several industry groups and standards bodies working on various aspects of optical SDN, including the Open Networking Foundation (ONF), the Optical Internetworking Forum (OIF), the International Telecommunications Union (ITU), and the Internet Engineering Task Force (IETF). Each are addressing a component or series of components of the overall system and the work is expected to be coordinated and complementary.

The Open Networking Foundation (ONF), for example, launched the Optical Transport working group with a charter to address SDN and OpenFlow Standard-based control capabilities for optical transport networks. The workgroup concentrates on three use cases: Photonic Enterprise Network (private L0 networks connecting data centers and Greenfield scenarios), Service Provider Network Virtualization/Data Center Interconnection (virtualization of network resources to multiple clients, virtual network slices offered to Data Center Operators), and Service Provider Packet/ Optical Integration (provider visibility into and orchestration of IP and optical networks where the optical network may include multilayer transport elements).

The ONF OTWG current work includes defining a target reference architecture for controlling optical transport networks incorporating the OpenFlow Standard, and identifying and creating OpenFlow protocol extensions.

The Optical Internetworking Forum (OIF) Carrier Working Group has published a carrier requirements document addressing the transport network part of an SDN network. The document supports the evolution of transport networks towards SDN architectures, summarizing what is already available and highlighting new features and functionalities to support the deployment of SDN applications, services and technologies. The next step for the OIF is to complete the SDN framework document that will provide future SDN related specifications and Implementation Agreements and also give structure to future OIF interoperability demonstrations.

Through the work of these and other industry groups and standards bodies a rich and diverse group of vendor implementations is expected to be developed that can be integrated into a complete SDN network design. MRV, as a member of the OIF and ONF, will be a strong participant in the standards definitions and a valuable contributor to the SDN/NFV vendor community.

Optical SDN and MRV

At MRV, we like to say that we have developed the SDN bricks necessary to build a strong SDN house and that the house is being built to the requirements communicated by our customers. This is true for Optical SDN as well.

On the packet side, MRV has a strong foundation on which to build our SDN house, including a common Linux-based operating system that can be virtualized to run on any platform inside or outside of the network. MRV is a participating member of several industry groups that are actively working to define SDN and has developed software and hardware architectures that are specifically engineered for SDN-based networks.

On the optical side, MRV has already implemented a hardware architecture that supports the flexibility required for Optical SDN. The DMR series of transponders from MRV support fully programmable signal rates and formats, allowing a single circuit pack to be used for applications ranging from low-speed Ethernet services to high-speed OTN services.

Additional flexibility in the DMR transponders allows the same hardware to be used , for example, as either a protected client/line interface, a set of two independent transponders, or a broadcast/tap interface. All optical hardware in the MRV product line supports optics that can be flexibly tuned via software to any wavelength in the ITU-T spectrum. Flexible ROADMs are available that have flex-grid hardware at the core, allowing future upgrades to network designs as more complex modulation schemes and higher bandwidth signals are introduced.

MRV's OptiDriver Product Line

On the software controller side, MRV offer our open-architecture-based Pro-Vision™ service management platform and a set of open interfaces for integration with third-party controllers. MRV is committed to having an open control interface to allow integration into increasingly powerful and flexible network controllers. MRV is also implementing a standards-based GMPLS control plane that will interact with an SDN controller to allow, e.g., dynamic, constraint-based optical transport network programming. This can enable applications such as fast network restore combined with disaster recovery to re-establish work/protect paths in the case of network failure.

By combining our expertise in packet and optical networking and in software-defined networking in particular, MRV expects to be a leader in the emerging field of Optical SDN. For more information about MRV and our SDN portfolio, please contact your local MRV representative.

MRV operates worldwide sales and service offices across four continents.
Contact us at info@mrv.com
www.mrv.com