High Performance & Optical Networks

By Jimmy Kyriannis
jimmy.kyriannis@nyu.edu

NYU's commitment to academic research has placed increasing demands on NYU-NET, the campus network, to facilitate high-speed access to a variety of resources, including: rich volumes of data sets, high performance compute grids,[1] national laboratories and external research institutions, partner universities, and dedicated research networks. Since 1997, the University's primary access to such destinations has been our 200 megabits per second (Mbps) connection to the Internet2 Abilene research network (commonly known simply as Internet2). Recently, however, a number of new research networks have been developed, and over the summer, NYU-NET connections were established to several of them. The means by which these connections have been established is new to our network and may present novel network capabilities in the future.

These new research networks have been established primarily to meet a need for access to resources via higher network performance than is conventionally possible with an Internet2 connection. Such resources may require that hundreds of Mbps be dedicated to data transfers to a university campus. In some cases, where very high bandwidth and/or low latency[2] is required, one Gigabit per second (1 Gbps) of bandwidth or more may be called for. (At a data rate of 1 Gbps, the contents of a full CD-worth of data will be transferred in just over five seconds.) Invariably, connections that take place over long distances require fiber optic cable links, which can be costly to implement, and fiber optic cabling between two facilities may not necessarily be in ample supply from intermediate service providers. Technologies such as Wave Division Multiplexing (WDM), however, can improve the utility and capacity of a single fiber optic cable link significantly.

Traditionally, a network communications signal is transmitted between two network devices as a wavelength of laser light over a pair of fiber optic strands (one fiber strand for receiving signal, and the other for transmitting), commonly referred to as a single "fiber optic cable." The signal "carrier" itself is maintained by rapidly modulating the transmitted light; one common example of this technique is to vary the intensity of the light in a manner which represents a series of binary 0's and 1's sent to the receiver. This approach, however, limits a fiber cable to a single network connection and doesn't offer significant capacity, especially over long distance fiber optic links. Wave Division Multiplexing (WDM), however, addresses this need by transmitting a "rainbow" of colors of light simultaneously on the fiber. Each color of the WDM rainbow is infrared light at a different wavelength, and is able to carry communications independent from all other colors within that rainbow by modulating each wavelength independently. Consequently, the application of WDM enables many network connections, such as connections to Internet Service Providers, external research institutions, video service providers, and telecommunications carriers, to take place simultaneously all over a single fiber optic cable.

Wave Division Multiplexing technology is available in two flavors: Dense WDM (DWDM) and Coarse WDM (CWDM). The primary distinction between the two lies in the characteristics of the optical signal used to provide the communications channels. Each channel corresponds to a different wavelength (color) used by the WDM system. In DWDM, the channels are densely packed to offer a greater number over a single fiber optic connection, while in CWDM many fewer channels are available. For example, DWDM systems available today can offer 128 simultaneous channels, while CWDM is usually limited to 8 or 16. CWDM

1. A compute grid is an interconnected group of computers that act as one. See "Global Grid Computing at NYU" by Richard Bonneau on p. 2 for more information about high performance compute grids.
2. Latency refers to the amount of time it takes a packet of data to move across a network connection.

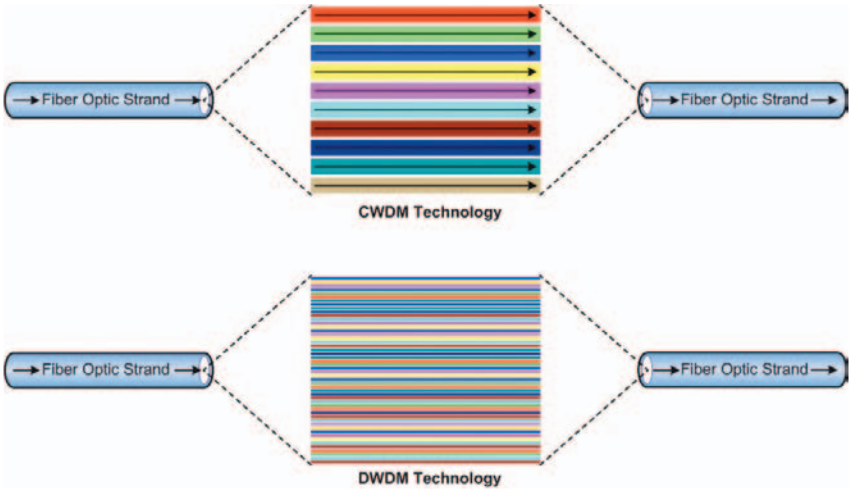

Figure 1: CWDM and DWDM technologies. DWDM makes much more efficient use of a fiber optic cable by transmitting many more channels (wavelengths of light) simultaneously over a single fiber strand.

is commonly used for inter-building or intra-city networks over short distances (one to two kilometers, typically) and with limited bandwidth capacity. DWDM systems, however, make more efficient use of fiber optic cabling and can deliver much more bandwidth. This is often achieved over very long distances (tens of kilometers or more) and can be accomplished as a network in the form of a ring, where a single break in the fiber cabling system will not disrupt communications; data will simply traverse the network to its destination via the opposite and unbroken path. As might be expected, DWDM systems are more costly to implement due to the complexity and intelligence of the devices involved.

DWDM networks are becoming a standard component of network support for research that has need for high performance communications, computing, and data storage. National networks that focus on supporting research and education (R&E) by interconnecting universities, research institutions, and government research laboratories are using DWDM as an enabler, providing institutions with high-bandwidth access to network-based resources through the reduced consumption of costly transcontinental fiber. Both National Lambda Rail (*www.nlr.net*) and Internet2 (*www.internet2.edu*) have built high-speed networks that span the United States and provide high-speed connectivity to Regional Optical Networks (RONs), such as NYSERNet (the RON for New York State), CENIC, and Merit.[3] The various RONs, in turn, support R&E networking within geographic regions of the U.S. by providing high-bandwidth connectivity to their subscriber institutions.

NYU and other New York State institutions, for example, traverse the NYSERNet network via its DWDM infrastructure. This connectivity offers high-bandwidth access to other NYSERNet subscribers and large-scale research networks such as Internet2, ESnet (the U.S. Department of Energy's Energy Sciences Network; *www.es.net*), and MANLAN (the Manhattan Landing network, where international research networks interconnect in NYC; *http://networks.internet2.edu/manlan/*).

Overall, DWDM technology is enabling service providers to offer an increasing variety of network services. The National Lambda Rail (NLR) currently has a DWDM service offering, called WaveNet, where a customer can purchase a "lambda" (i.e., a channel or wavelength of light) for dedicated communications between two NLR-accessible points within the United States. Internet2 recently announced its WaveCo service, which is a similar and competing lambda service offering. NLR's PacketNet service, on the other hand, is a competitor with Internet2. Both are national networks which offer IP (Internet Protocol) connectivity between member institutions, essentially creating networks parallel to the Internet and fairly similar in function, although the NLR and Internet2 networks are currently used predominantly by research institutions. Internet2 has been in production service since 1999, and is quite mature, with a very large number of connected participants. PacketNet, on the other hand, is less than a year old, and adoption of the network is not yet widespread.

In parallel with PacketNet, NLR offers an additional service, called FrameNet, which is an Ethernet network that spans the United States. Unlike PacketNet, which has intervening NLR routers that negotiate traffic between organizations, FrameNet has no intervening routers; institutional routers can communicate with one another directly, thereby enabling inexpensive yet direct communications between

3. Information about CENIC, the Corporation for Education Network Initiatives in California, is available at *www.cenic.org*. Details about the Merit Network are available at *www.merit.edu*.

institutions. Since FrameNet is a shared environment, however, the entire nation shares 10 Gbps of bandwidth, and performance issues may arise with its use.

Finally, Internet2 has recently embarked on the development of its second-generation network, called NewNet,[4] which also relies heavily on DWDM. Though NewNet is different from the other services, a hallmark of all these networks is the provision of high-bandwidth network access for institutions at speeds as high as 10 Gbps. Through the adoption of hybrid network technology, NewNet represents a new approach to data networking. Pioneered by such research projects as Internet2's HOPI,[5] hybrid networking makes use of DWDM and IP technologies simultaneously, and in some cases, extends DWDM capabilities to individual users of the service. In a hybrid network, IP can be used as it conventionally is on the Internet today: email, web services, and network applications all function as they normally do. However, in such a network environment, IP can also be used to signal the optical network to create a point-to-point lambda to a particular destination on an on-demand basis. Similarly, when no longer needed, that lambda can be deactivated and the network resources made available again for the next application.

The applications that might make use of hybrid networks are those which require direct fiber optic communications between two endpoints, such as a processor interconnect[6] between two compute grid nodes, a SAN (Storage Area Network) connection between disk storage systems in different locations, or delivery of an uncompressed high-definition digital video feed. It is noteworthy, however, that hybrid networking technology is still very much experimental and it is not yet clear whether hybrid will be the next evolutionary step in networking. In a hybrid environment, ultimately the applications drive the need; if applications emerge that cannot be adequately supported via IP technology, then a hybrid may present the solution. Hybrid networks, however, are very complex systems that require a great deal of planning and effort to implement, and the underlying control software and networking protocols are still fairly young technologies. Only a handful of hybrid networks have been deployed thus far, and NewNet will be the world's largest endeavor to date. As ITS is a participant in the Internet2 NewNet design process, we will be following its progress and collaborating with other institutions with similar hybrid interests.

This past summer, high-speed network connections have been established from NYU-NET to a number of these new research networks. As described in the article "High Energy Physics & High Performance Technology Coming Together at NYU" on p. 7, a 1 Gbps link has been established to the USLHCNet network.[7] USLHCNet was designed to support the computing needs of physicists participating in research activities at the CERN Physics Laboratory in Geneva, Switzerland, particularly the ATLAS and CMS projects, through the establishment of a high-speed network supporting the transmittal of the vast amounts of data generated by those activities. As a result of this networking initiative, NYU now has 1 Gbps of dedicated network access to CERN.

Through a similar collaboration with the California Institute of Technology and CERN, ITS has also established a high-speed connection to the UltraLight network. UltraLight is a National Science Foundation-funded project designed to advance hybrid networks in support of the high-bandwidth needs of experimental physicists. The network is supported by much of the same staff and infrastructure that operates USLHCNet. The designers of UltraLight hope to use hybrid communications to further technologies such as compute grids, disk-to-disk communications for storage of large volumes of data across large distances, and applications that require brief periods of high-speed connectivity on an on-demand basis. UltraLight provides an environment for the development of systems, networking components, communications protocols, and software, all involved in hybrid network communications. Additionally, 1 Gbps connections have been established to the NLR FrameNet and PacketNet networks. Though still fledgling networks, they offer the potential for high-bandwidth connectivity to institutions at speeds that exceed our current performance levels using Internet2.

For the past few years, CWDM technology has been used at NYU to provide access to these and other network resources external to the University, where fiber optic cabling availability has been limited. This semester, that infrastructure is being upgraded to DWDM technology. The expectation is that this new service will support higher levels of external bandwidth, increased service reliability, greater capacity, and academic endeavors requiring access to these and future research networks.

For information about high performance computing at NYU, see *www.nyu.edu/its/supercomputing/* or send email to hpc@nyu.edu.

Jimmy Kyriannis is Senior Technology Architect in ITS Communications and Computing Services. High performance networking and computing are among his focus disciplines.

---

4. A technical overview of the NewNet architecture is described in: *www.internet2.edu/files/Internet2-New-Network-Tech-v0.9.pdf.*
5. Additional information on HOPI is available at *http://networks.internet2.edu/hopi/*.
6. A processor interconnect is a connection between computers that enables their CPUs (Central Processing Units) to directly communicate with one another. This is often used in grid systems, where multiple computers are expected to act as one.
7. *www.uslhcnet.com*