# A new look at dynamic traffic scheduling in WDM networks with transceiver tuning latency

Andrew Brzezinski and Eytan Modiano
MIT, LIDS

A key feature of optical wavelength division multiplexed (WDM) networks is the inherently high degree of configurability available in the presence of tunable transceivers. More specifically, the flexibility of WDM networks lies in the ability to admit a large number of logical interconnections of nodes within the network. Indeed, this is a feature that distinguishes optical networks and drives a significant portion of the optical communication-theoretic research community.

One important area of current research has at its heart precisely this feature of network configurability; this is the problem of dynamic scheduling of logical topology reconfigurations under random traffic. This problem has been approached from several angles, including: utility-function-based optimization where a cost function is associated with different reconfiguration decisions given the current network configuration; algorithmic studies for load balancing of optical links; and analyses of the effects of very short and very long tuning latencies, respectively. In this work, we pursue the scheduling problem for dynamic reconfiguration in WDM networks with nonzero transceiver tuning latencies. In particular, we let the tuning latency explicitly take part in the network evolution by imposing service stoppage (corresponding to the delay associated with reconfiguration) following reconfiguration decisions. Our approach utilizes techniques from switching theory, dynamic programming, and system stability theory, and addresses the dynamic reconfiguration problem at a fundamental level.

To be precise, consider an optical WDM network consisting of N nodes, physically interconnected by optical fiber in an arbitrary topology. Each node is equipped with a single transceiver, and thus at any particular time may have at most one incoming and one outgoing logical link. Under this restriction, we assume that there exist sufficiently many wavelengths to allow any arbitrary logical interconnection of nodes. Each node is equipped with (N-1) virtual output queues (VOQ) in which data are held prior to transmission across the network. Time is assumed to be slotted, and data units are in the form of fixed-length packets, with the transmission time for one packet across the network equal to one slot time. The network allows a maximum of one packet to be transmitted across any logical link during a slot. At any time, the network may initiate a logical topology reconfiguration, under which existing logical links are torn down and new ones reestablished to form a new logical topology. Links affected by reconfiguration are forced to be idle for the reconfiguration time of D slots, while links that are not directly modified may continue to service traffic during reconfiguration.

Given the above network model, the reconfiguration problem is to determine a schedule that trades off additional delay imposed by logical topology reconfiguration against the delay imposed by continuing to service the current logical topology, where data may follow multiple-hop routes to their destinations. Assuming that exogenous arrivals to the network satisfy a strong law of large numbers (SLLN) property, with arrival rates that are strictly substochastic, we establish that the above network model is asymptotically stabilizable by employing a frame-based scheduling algorithm. This algorithm makes use of maximum weighted matchings to make reconfiguration decisions and imposes a fixed backoff (corresponding to the frame duration) in between decision epochs, during which time queues are serviced according to the logical topology selected by the matching. During the D slot reconfiguration phase, the algorithm does not service any queue in the system. The stability result is based on fluid limits of the network queueing system. Further, we develop a novel bias-based matching algorithm that is again asymptotically stabilizable, and whose delay performance improves on the frame-based scheme. In this algorithm, maximum weighted matchings are again employed, but the matchings are calculated at every time slot, with the current logical topology having its weight augmented by a fixed "bias". This stability result follows according to a Lyapunov drift argument on a sampled version of the underlying Markov chain describing the system.

Next, we note that the stabilizing algorithms never require data to follow a multiple-hop route, which we show to be suboptimal with respect to average data transmission delay when there are nonzero tuning latencies. Thus we consider the same matching algorithms but impose multiple hops on select data in order to improve the average data transmission delay. In particular, a random packet marking scheme is employed to designate particular packets for multiple-hop routes. We employ heuristic algorithms to search for the best multi-hopping marking schedule under the given matching algorithm with results showing significant improvements in average transmission delay over the initial stabilizing case.