legacy PDH within the United Kingdom and Europe resulting in a need to support both formats. This differs from the U.S. long-haul market, which is predominantly SONET with inbuilt protection. PDH however, has only limited protection and is consequently susceptible to cable outages; protection in the optical domain therefore offers significant benefits. While the economic drivers of minimizing fiber build and reducing repeater costs applies, the ability to support multiple formats and to provide protection adds a new dimension. Network management in a multitechnology and multivendor environment is a key issue and the subject of this paper.

An important feature of the trial is the introduction of optical protection switching. In the optical domain, optical multiplex section linear trail protection provides both equipment and fiber protection. Either single-ended or dual-ended nonrevertive protection can be enabled by the management system, though the latter is preferable for operational reasons. Its principle attractions are that it minimizes both protection plant and the number of protection switching events. While this optical protection enhances the availability of unprotected client layer networks, care must be taken where more than one layer employs some form of protection. In the case of SDH VC-n subnetwork connection protection it is fast (~50 ms) and autonomous. This leads to a problem in that there is no communication between the optical network elements and the SDH terminals. Optical protection switching is not fast enough to prevent SDH protection from being invoked unless holdoff times are introduced. It is therefore necessary to have an escalation strategy and a number of these will be reviewed in integration testing prior to deployment.

The approaches adopted in the trial are wholly consistent with the architecture described in G.otn.[1] The management of optical network elements (ONEs) is independent of any supported transmission technologies.[2] This separation maximizes the potential for transparency. Network management is provided by means of a local/remote craft terminal or by means of an element management system (EMS). The remote EMS can be connected to every ONE via a datacommunications network. Alternatively, a gateway network element can be employed to pass messages to other ONEs within its domain via a data communications channel embedded in an optical supervisory channel. To obtain a network level viewpoint and correlation between PDH, SDH and optical layers requires the EMS to be connected to an Operational Support System (OSS) and is currently under study. The Network Operations Centre provides the focus for cross-platform management via the OSS. Within BT's current trial workstations in operations centers are remotely connected to the EMS, for national rollout the EMS would be connected to BT's OSS.

An OSS represents a significant investment for many network operators and is a key source of competitive advantage. As the number of underlying network technologies grows and the need to provide multiplatform management intensifies, the need for a well-designed OSS will become paramount.

1. ITU-T Draft Recommendation G.otn: "Architecture of optical transport networks."
2. A. McGuire *et al.*, "Applications of optical supervisory channels," EFOC&N '95, Brighton, United Kingdom.

TuP2 5:00pm

A novel architecture and medium access control (MAC) protocol for WDM networks

Eytan Modiano, Richard Barry, Eric Swanson, *MIT Lincoln Laboratory, 244 Wood St., Lexington, Massachusetts 02173*

The emergence of bursty bandwidth intensive applications such as medical imaging and supercomputer interconnection, has lead a recent wave of research on multiple access protocols for wavelength-division multiplexing (WDM)-based LANs and MANs.[1] Most of the proposed systems assume a synchronized and slotted system and many require multiple transceivers per optical terminal (OT),[1,2] contributing to their high cost. In this talk we will describe an architecture and an associated MAC protocol that eliminates the need for slotting and synchronization, uses one tunable transceiver per OT, yet results in efficient bandwidth utilization. We will also describe the potential application of this technology for providing bandwidth on demand in access networks and as a general purpose, high-performance, network technology.

Our system is novel in a number of ways. First, it uses an unslotted MAC protocol, yet results in high efficiency even in high-latency environments. While there are other unslotted MAC protocols, they are inefficient in high-latency environments.[3] A second novelty of our system is that it uses a centralized scheduler, which is able to schedule transmissions efficiently and overcome the effects of propagation delays. Lastly, our system is extended to MANs with a layered architecture that uses synchronization between LAN hubs, while terminals remain unsynchronized.

In the LAN our system uses a broadcast star architecture with 32 wavelengths each operating at 1–10 Gbit/s per sec., and a simple master/slave scheduler as shown in Fig. 1. All OTs send their requests to the scheduler on a dedicated control wavelength, λ_C. The scheduler, located at the star, schedules the requests and informs the OTs on a separate wavelength, $\lambda_{C'}$, of their turn to transmit. Upon receiving their assignments, OTs immediately tune

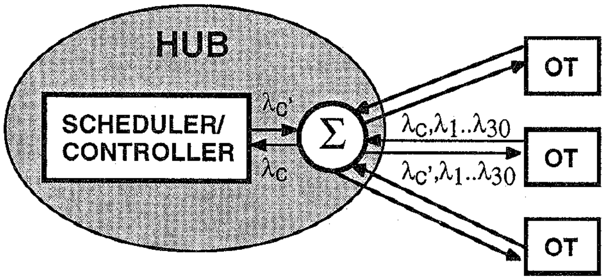

TuP2 Fig. 1. Scheduler-based LAN.

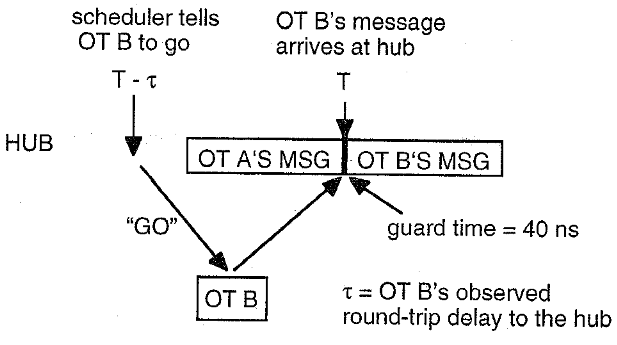

TuP2 Fig. 2. Use of ranging to overcome propagation delays.

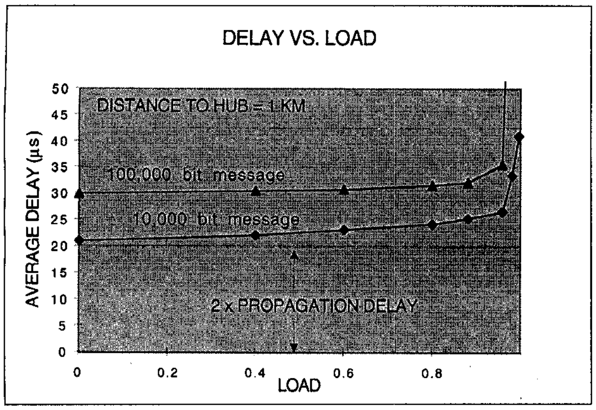

TuP2 **Fig. 3.** Total system delay vs. load.

to their assigned wavelength and transmit. Hence OTs do not need to maintain any synchronization or timing information. By measuring the amount of time that OTs take to respond to the assignments, the scheduler is able to obtain an estimate of each OT's roundtrip delay to the hub. This delay information is then used by the scheduler to overcome the effects of propagation delays, as shown in Fig. 2.

The above approach is not easily extendible to the MAN because of the increase in both propagation delays and the number of users. We, therefore, extended the system to the MAN with a layered architecture, where LANs are interconnected using lightpaths through the MAN. The MAN architecture can be either passive (e.g., wavelength router) or configurable (e.g., ring with frequency selective switches). In order to efficiently schedule transmissions in the MAN, LAN hubs are synchronized to a common clock, however, OTs remain unsynchronized. With some modifications, this architecture can also be used to provide bandwidth-on-demand in access networks.

The system uses simple scheduling algorithms that can be implemented in real time. Unicast traffic is scheduled using first-come-first-serve input queues and a window selection policy to eliminate head-of-line blocking, and multicast traffic is scheduled using a random algorithm.[4] Simulations show that the total system delay is relatively low even at high traffic loads (see Fig. 3). In the talk, we will discuss the details of our novel system, contrast its performance with other work, and describe its potential application to general purpose IP networks and as an access network technology.

1. B. Mukherjee, IEEE Network (May 1992).
2. N. Mehravari, IEEE J. Lightwave Technol. (April 1990).
3. Fouad A. Tobagi, IEEE Trans. Commun. (April 1990).
4. Eytan Modiano, "Unscheduled Multicasts in WDM Broadcast-and-Select Networks," submitted to INFOCOM '98.

TuP3 5:15pm

Batch wavelength assignment in all-optical central-switch networks

Won S. Yoon, Richard A. Barry,* *Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 77 Massachusetts Ave., Room 35-303, Cambridge, Massachusetts 02139; E-mail: won@mit.edu*

There have been numerous studies of dynamic wavelength assignment in all-optical networks that sequentially establish calls one-by-one. We believe we present the first known study of assigning wavelengths to batches of calls at a time. One motivation for this is an optical backbone network supporting logical topologies (each one a set of connections used by a higher-layer network), which arrive, depart, or reconfigure unexpectedly.

We propose a framework for studying batch assignment and apply it to a central-switch network where N stations are interconnected by bi-directional links to a wavelength-selective switch without wavelength changers. The batches are characterized by: 1. a maximum link load, *i.e.*, no more than L directed calls share a link, and 2. a minimum batch size, *i.e.*, no fewer than B calls may be requested at a time. Further, we assume calls disconnect individually at random. There are no other traffic assumptions.

Let $W_{NB}(B)$ be the minimum number of wavelengths needed for no blocking for a fixed L. Note that $W_{NB}(B)$ is strictly nonincreasing in B with a maximum at $B = 1$ (corresponding to sequential assignment) and a minimum at $B = NL$ (corresponding to static assignment).[1,2] The latter follows because there can be at most NL calls in this network at any time.

We study two scenarios: 1. strict-sense nonblocking (SSNB) operation, where there always exists a valid set of wavelengths for the next batch regardless of previous assignments, and 2. nonblocking operation using any greedy algorithm. The class of greedy methods, defined as "never assigning a new wavelength unless there is no other choice," includes First-Fit[3,4] and Most-Used[1] when $B = 1$.

For SSNB, it was previously known that $W_{SSNB}(B = 1) = 2L - 1$.[5] We prove that in fact $2L - 1$ wavelengths are needed for all B up to $B_{cutoff}^{SSNB} = NL - 2L + 2$. For larger B, the number of wavelengths steadily decreases to a minimum of $W_{SSNB} = L$ at $B = NL$ (Fig. 1).

For greedy algorithms, $W_{SSNB}(B)$ is upper-bounded by $W_{SSNB}(B)$ and lower bounded by the curve shown in Fig. 2 for several values of N and L. For small L/N (upper plot), the upper and lower bound are nearly identical, and both are flat over a large range of B. As L/N becomes large (lower plots), however, greedy batch methods may require fewer wavelengths than SSNB.

Thus, in large central-switch networks with small link loads (small L/N), we conclude that unless batch sizes are very large ($\approx NL$), batch assignment performs no better than sequential for SSNB as well as for greedy batch algorithms. However, in small networks with large link loads (large L/N), SSNB batch methods require fewer wavelengths for moderate sized batches. Furthermore, greedy batch algorithms can provide at most a moderate improvement. It is an open question if there exist nonblocking algorithms, which require fewer wavelengths than greedy batch methods.

We have also studied two other call-departure scenarios of no disconnections and disconnections only in the same batches, and we find

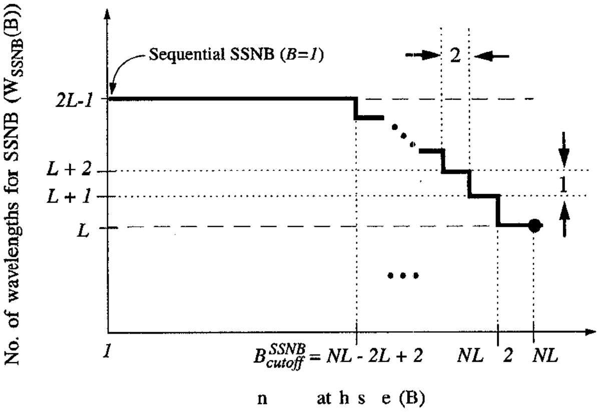

TuP3 **Fig. 1.** Number of wavelengths needed for strict-sense nonblocking.