WAVELENGTH ALLOCATION STRATEGIES IN OPTICALLY SWITCHED NETWORKS FOR AVIONICS

Casey B. Reardon, John D. Profumo, Alan D. George
High-performance Computing and Simulation (HCS) Research Laboratory
Department of Electrical and Computer Engineering, University of Florida
Gainesville, FL 32611-6200

1. Introduction

For advantages in performance, scalability, protocol transparency, cost-effectiveness, etc., WDM-based optical networks are being evaluated for use in emerging avionics systems. One area where optical technologies are rapidly developing is optical switching. Optical switches offer the potential performance, scalability, and flexibility that advanced avionics networks will demand in the future. This paper investigates via simulation modeling an optical switching architecture for networking on advanced avionics platforms. Two strategies for wavelength allocation are presented, and their performance characteristics are compared. Additional simulation experiments analyze the effects of varying two architectural parameters within each allocation strategy.

2. Background Information

The optical switching architecture evaluated in this paper is based on the OSMOSIS architecture developed by IBM for high-performance computing systems [1]. Each point-to-point connection includes an optical data path and a separate electronic control path used to request and reserve optical connections. Transmission requests are made to the switch through the control path. The switch arbiter responds to the request indicating when the optical path is available, and reserves the path in the optical backplane. Data is then transmitted over the optical connection for the number of time slots allotted by the switch arbiter. At the end of each time slot, the optical outputs are reconfigured, if needed, for data transmission over the next time slot. A broadcast-and-select design is used for the optical switching. Each optical input is split and distributed to all outputs. Each output then chooses the desired optical input. The choice of optical devices in this stage will heavily impact switching speeds.

An OSMOSIS-based switching architecture was modeled in the Library for Integrated Optical Networks (LION). Developed at the University of Florida, LION was created within a discrete-event simulation environment called MLDesigner from MLDesign Technologies. LION provides users with the flexibility to design systems with a broad range of optical devices and varying high-level components. New and legacy network protocols can be layered on top of optical components to realize and evaluate almost any system design.

3. Case Study based on OSMOSIS

In this case study, two wavelength allocation approaches are considered. The first approach, *fixed-destination* assignment, allocates each destination node a single wavelength. Optical transmitters are thus responsible for tuning to the correct wavelength of the destination for each message. The effects of tuning delays are minimized in this approach by performing the tuning while the switch arbiter performs the scheduling. The second approach, *fixed-transmitter* assignment, allocates a fixed wavelength to each transmitter. Detectors are now responsible for tuning to the correct wavelength. While tuning delays cannot be overlapped with scheduling as before, efficient multicasting is possible, as multiple outputs can select a single input during any time slot.

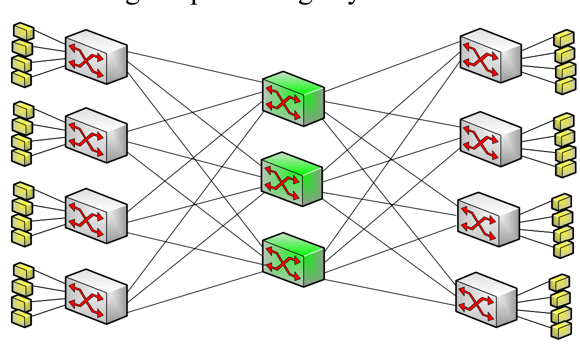

Figure 1. Proposed LAN Topology

An experimental setup designed to represent a future avionics platform is used in our simulative experiments. This setup is comprised of 97 nodes, which generate a combination of bursty, random, and continuous traffic, totaling 200 MB/s. Message sizes are uniformly distributed between 1,000 and 30,000 bits. The traffic pattern of this setup mimics a centralized avionics architecture, where most information passes through a central processing system. Results are obtained from one second of traffic simulated in this system.

The layout of the network implementation is illustrated in Figure 1. The various end-nodes of this network are connected to one of eight perimeter switches, which are interconnected by three backplane switches. Each perimeter switch is designed to accommodate up to 32 end-nodes.

4. Results

The mean packet latencies measured in our simulative experiments are presented in Tables 1 and 2. For each approach, four different timeslot periods are considered. Also, three values for the maximum number of consecutive slots allotted to a single node at once are used. In each case, 100 nanoseconds of every time slot is used for optical switching. Laser transmitters operate at 2.5 Gb/s.

Table 1. Fixed-Destination Mean Latency (μs)

Maximum Slot Allotment	Timeslot Period (ns)			
	300	**500**	**1,000**	**2,000**
7	29.8	23.2	21.3	22.8
10	29.5	22.0	21.4	22.4
15	29.2	22.9	21.0	22.8

Table 2. Fixed-Transmitter Mean Latency (μs)

Maximum Slot Allotment	Timeslot Period (ns)			
	300	**500**	**1,000**	**2,000**
7	31.0	24.0	22.1	23.5
10	30.8	23.3	22.2	23.5
15	30.6	23.6	21.7	23.2

The results in Table 1 and Table 2 show that the fixed-destination wavelength protocol offers slightly better performance than the fixed-transmitter protocol in these experiments. The average difference is approximately 1 us, which is also the optical tuning delay used in our models. The primary reason for the difference is the ability to overlap scheduling and optical device tuning. Networks with broadcast and multicast traffic would see increased benefits from the fixed-transmitter strategy, where multiple nodes can simultaneously receive the same transmission.

Within each wavelength assignment protocol, the most obvious trend observed is the reduction of packet latency as the timeslot period increases. Since the time reserved in each slot for optical switching is constant (100ns), a smaller fraction of overall time is used performing optical switching with longer timeslots. When the timeslot period is increased to 2,000 ns, performance begins to decrease. One factor is the underutilization of allotted resources with small messages, which only occupy a fraction of one timeslot. Additionally, a minimum of two timeslots are needed for scheduling any transaction. Thus, a longer timeslot period increases the minimum scheduling delay every packet observes.

Small gains in network performance were also observed by increasing the maximum slot allotment parameter. While a higher consecutive timeslot allotment decreases the average message latency, the theoretical maximum queueing delay in the round-robin scheduler increases.

5. Conclusions

In this paper, an architecture for an optically-switched avionics network is presented. The performance of this architecture was analyzed using simulative experiments for two wavelength allocation strategies. Additionally, the effects of varying the timeslot period and maximum number of consecutively allotted timeslots were analyzed. Results showed slightly lower packet latencies with a fixed-destination wavelength protocol. Larger timeslot periods, up to 1,000 ns, also improved performance, although further increases would not be beneficial. The optimal parameters for any platform will always depend upon the nature of the network traffic. Our modeling tools and approach allow us to evaluate such design decisions for a wide range of network scenarios.

6. References

[1] Hemenway, R., and R. Grzybowski, "Optical Packet-Switched Interconnect for Supercomputer Applications," *Journal of Optical Networking*, Vol. 3, No. 1, Dec. 2004.