

Development of the iPOINT Testbed for Optoelectronic Asynchronous Transfer Mode Networking

Abstract

Asynchronous Transfer Mode (ATM) technology will be used to build Internet backbones and within local-area networks to interconnect personal computers and workstations running network-intensive applications. Applications such as the World Wide Web (WWW), desktop video conferencing, distributed computing, network file systems, and remote access to supercomputing resources have dramatically increased the demand for bandwidth in current networks. Future networks must scale in bandwidth as the computational power of the desktop computer as well as the number of installed computer systems continues to grow exponentially.

The Illinois Pulsar-based Interconnection (iPOINT) testbed has been established to design and develop scalable ATM switch architectures, explore multicast scheduling algorithms to enable one-to-many data delivery, and investigate new techniques for data queuing. Within this testbed, Field Programmable Gate Array (FPGA) technologies have been used to prototype the core logic of the ATM switch, scheduler, and queue modules. In conjunction of the UIUC microelectronics laboratory, low-cost optoelectronic components have been developed and inserted into the testbed. As a part of the computer science research, an MPEG/JPEG video server and client has been developed and operates on this fully functional testbed.

John W Lockwood, Ph.D., Visiting Assistant Professor of Electrical and Computer Engineering
Sung Mo Kang, Ph.D., Department Head of Electrical and Computer Engineering
Stephen G. Bishop, Ph.D., Director of the Center for Compound Semiconductor Microelectronics
Haoran Duan, Ph.D. candidate in Electrical and Computer Engineering
Ashfaq Hossain, Ph.D. candidate in Computer Science
The University of Illinois at Urbana-Champaign
405 North Mathews Avenue
Urbana, IL 61801, USA
(217) 244-1565
(217) 244-8371 (fax)
<http://ipoint.vlsi.uiuc.edu> (WWW)
lockwood@ipoint.vlsi.uiuc.edu

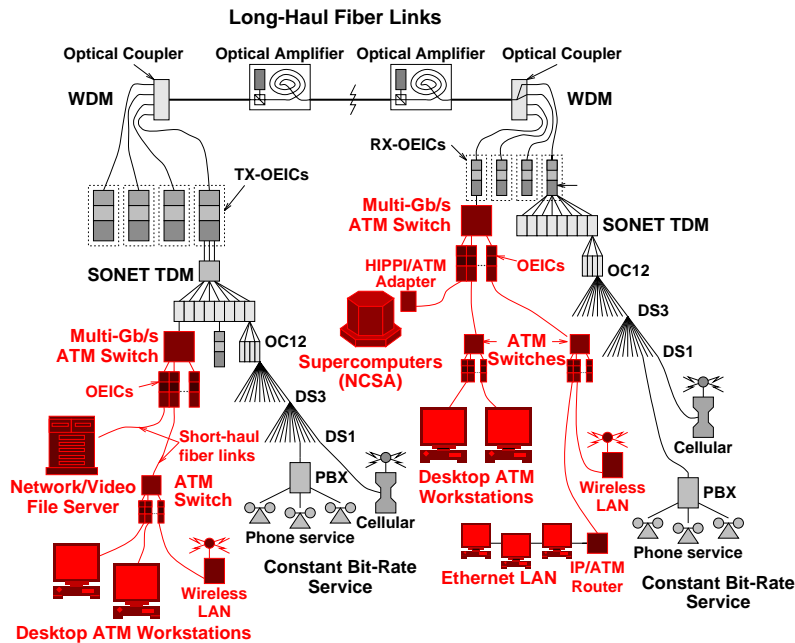


Figure 1: The role of ATM in high-speed networks

Broadband Integrated Networks

The implementation of a broadband communications network involves the integration of time-division multiplexors, ATM switches, optoelectronic components, and fiber-optics. As shown in Figure 1, the topology of the network is generally hierarchical in nature, with lower bit-rate links combined into faster bit-rate channels.

Constant-bit-rate communications, such as telephone service, generate data with arrival times that are strictly periodic. Time division multiplexing (TDM) can be used to hierarchically combine multiple connections into faster bit-rate data channels. Because the data arrival times can be predicted, intermediate switches can reserve time slots for each connection for the duration of the call. Existing TDM technology, such as SONET, can be used to switch constant-bit-rate data.

Data sources such as personal computers, workstations, supercomputers, and compressed video encoders, however, generate bursty traffic. Because the packet arrival times cannot be statically predicted, ATM switches and data buffers are required to dynamically schedule and buffer data from each of the incoming sources. By statistically multiplexing the data, bursts of traffic can be transmitted without dedicating time-slots or fibers for otherwise idle communications. By sharing the bandwidth, these variable-bit-rate applications can maximize the utilization of the network.

This paper focuses on the darkened components of Figure 1. It discusses our implementation of a scalable ATM switch and its corresponding queue modules; the implementation of application software for transmitting and receiving compressed video on workstations; and development of low-cost optoelectronic technology for short-haul fiber links.




Figure 2: Cell scheduling for input-buffered ATM switches

Configuration of the iPOINT Switch

The configuration of the iPOINT switch is illustrated in Figure 2. To increase the scalability of the iPOINT switch, queue modules are placed at the input of the switch fabric. Incoming ATM cells from each of the input fiber links are processed in parallel by each of the input modules. Processing functions include the ATM header CRC check and virtual path identifier (VPI) / virtual circuit identifier (VCI) lookup, translation, and destination mapping. Of the contending cells in the queue modules, the switch scheduler selects those cells such that no more than one cell is transmitted from an input to no more more than one output during any given cell interval. Once the cell scheduling is completed, the switch fabric transfers the cells to their appropriate output ports and transmits the data on the outgoing fiber links.

Silicon memory (SRAM or DRAM) is the building block of any practical ATM queue architecture. While optical links are used to interconnect multiple switches and endpoints of the network, purely optical technology does not provide an efficient means for the storage of large amounts of data. The short (424-bit) length of the ATM cell requires that cells can be inserted and removed from memory in a short time interval (tens of nanoseconds). As the rate of the incoming switch traffic increases, the amount of time to read or write a cell from memory (the access time of the SRAM or DRAM) decreases. By using input queues, rather than a single shared memory or output queues, the scalability of the switch can be achieved because each memory device need only read and write a single cell at the rate of a single link rather than at the aggregate rate of the switch or at a multiple of the link rate.

Multicast Cell Switching

The scheduling algorithm of the iPOINT switch supports multicast, thus allowing incoming data from one port to be simultaneously delivered to one or more output ports. Multicast is an essential feature for future communication paradigms, as it enables network-wide, multi-user distribution of data from a source to multiple endpoints throughout the network. Such a feature is critical for applications such as video distribution and conferencing.

The iPOINT switch scheduling unit supports atomic multicast. Rather than transmitting duplicate cells through the switch fabric and using multiple output buffers to store multicast ATM cells, the scheduling module of the iPOINT switch finds a single time slot such the cell can be read from an input queue and simultaneously delivered to multiple output ports.

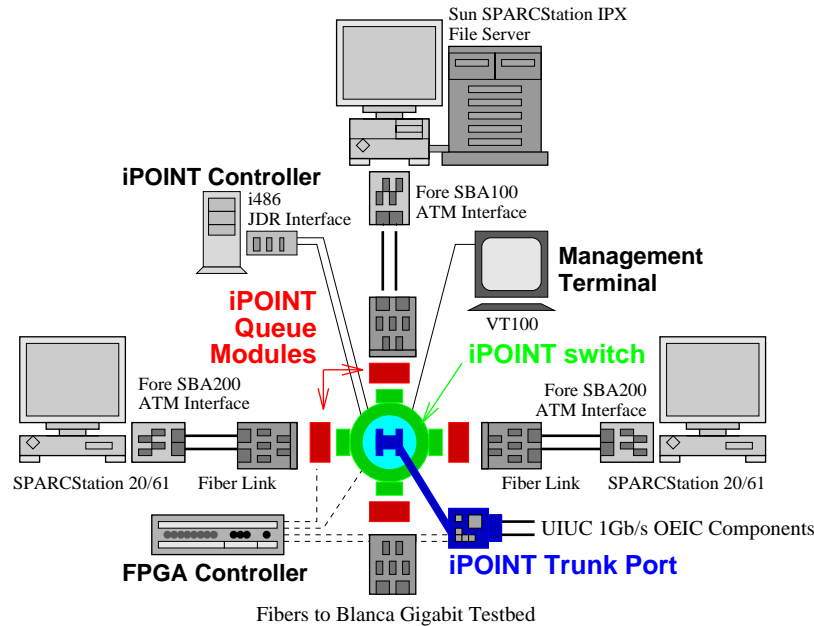


Figure 3: iPOINT testbed components

The Functional Testbed

The iPOINT testbed is a fully functional prototype system which currently serves as high-bandwidth local area network for a cluster of Sun SPARCStations and as a gateway to the Blanca/XUNET wide-area testbed. A diagram of the testbed is shown in Figure 3.

Field Programmable Gate Array (FPGA) devices were used to implement all of the logic functions within the testbed. By using reprogrammable logic, it has been possible to make iterative design changes to the logic without physically modifying the hardware. The FPGA controller allows an external workstation to download new design files into any of the individual FPGA devices.

At the center of the testbed is an 800 Mbps prototype switch. The switch is implemented with a Xilinx XC4013PG-223 FPGA device and has five ports. Four ports are used to connect to workstations at a rate of 100 Mbps. The fifth port is used as a trunk interface to the second-generation iPOINT switch (which is under development) and operates at 400 Mbps. A short-wavelength laser fabricated by the UIUC microelectronics center is used to transmit the data from the trunk port. Data is physically transmitted at 1 Gbps for compatibility with the second-generation iPOINT switch.

Each of port of the iPOINT switch is connected to an input queue module. The core logic of the queue module is implemented with a Xilinx XC4005PC-84 FPGA device. The input queue module verifies and calculates the cell header CRC; uses a table lookup operation to perform the VPI/VCI translation and determine to which outgoing port(s) the cell should be delivered; then buffers the cell until the switch scheduling algorithm finds in time slot in which the cell can be transmitted.

The iPOINT switch controller is used to dynamically create and modify Permanent Virtual Circuits (PVCs) and Switched Virtual Circuits (SVCs). Signals are sent from the switch controller to the FPGA device on the queue module to update translation table entries.

Optoelectronic Technologies

For long-haul fiber networks, the cost of the optoelectronic components (the laser, photo-detector, and modulators) can be amortized over a large number of users that are transmitting data on the link. Within local area networks, however, the cost of the optoelectronic components are proportional to the number of machines that are to be networked. For short-haul fiber links (< 500 m), it is possible to use lower-cost, short-wavelength lasers and photodetectors. A $0.85\mu\text{m}$ laser device has been fabricated by the UIUC microelectronics center and a GaAs-based integrated photoreceiver/amplifier has been developed for use in the 1 Gbps trunk port of iPOINT switch. Integrating optoelectronic devices are expected to play an important role in reducing the cost of optical communications in much the same way that the integrated circuit reduced the cost of building electronic circuits.

Experiments with ATM Video on Demand (VOD)

Video on Demand (VOD) allows network clients to access and display video information from a remote server. Within the iPOINT laboratory, we have developed a variable bit-rate video server application and video client programs to transmit either MPEG or motion JPEG video streams over an ATM network. Both applications run on Sun workstations and use the multithreading facilities available within Solaris operating system.

The Server awaits requests from the clients on the ATM network. When a request is received, the server reads the disk and transmits compressed video stream into the ATM network. The multicast features of the ATM network allow the data to be simultaneously transmitted to multiple clients. Feedback information from the client(s) allow the rate of data transmission to dynamically change as a function of network congestion or client processing overhead.

The client is responsible for receiving the compressed packets from the server, buffering the video packets, decompressing the data, displaying the video, and handling packet loss. Reception of the incoming packets and decompression of the encoded data are implemented as independent threads. As compressed frames arrive at the client, they are placed into a ring buffer by the receive thread. The decompress/display thread consumes the the compressed frames from the ring and displays the decompressed video. By storing a small portion of future video on the client, the display of the video is not affected by possible jitter in the network.

Conclusions

Asynchronous Transfer Mode networks can be scaled to meet the demands of future data traffic requirements. The ATM switch architecture discussed in this paper scales because each of the input queue modules can operate in parallel at the rate of an individual link rather than at a multiple of the link rate. A prototype ATM testbed has been developed and currently serves as the local area network for a cluster of Sun workstations and as a gateway to a wide area network. Video on demand applications have been developed and run on the workstations within this testbed. Using the multicast features of the switch, the VOD server discussed in this paper has been used to simultaneously distribute compressed video streams to multiple clients via the ATM network.