High Performance Interconnects

John Morrison
Los Alamos National Laboratory, Los Alamos, New Mexico, USA

Abstract

The High Performance Parallel Interface (HIPPI), Fibre Channel (FC), Scalable Coherent Interface (SCI), Futurebus+, Synchronous Optical Network (SONET), and Asynchronous Transfer Mode (ATM) are emerging standard interconnect technologies that are capable of delivering data at high data rates. Each of these technologies have different performance characteristics, approaches to error handling and architecture. An overview of these technologies and their applications to High Energy Physics experiments is discussed.

Introduction

The requirements for higher bandwidth interfaces are motivated by the exponential growth of processor performance and by increasing requirements for I/O. The computer bus, computer interface, and telecommunications standards bodies have chartered several projects directed to addressing the requirements for increased bandwidth. Several of these projects have reached the approved standards status and products are becoming available on the marketplace.

Futurebus+

Futurebus+ is a high performance bus that is capable of moving data at 100 MBytes/s and higher speeds. It is specified for high reliability with distributed arbitration, parity checking on both control and data lines, dual bus configurations, and support for live insertion/removal of modules. The bus also provides support for shared memory multiprocessor configurations. Futurebus+ provides support for real-time computing with its large number of priorities and accurate timing information [1]. Futurebus+ has been endorsed by the VMEbus International Trade Organization [2] and can be viewed as a successor to the VMEbus for higher performance applications.

Scalable Coherent Interface

The goal of the Scalable Coherent Interface (SCI) project was to develop a standard high performance bus that would overcome the limits of traditional buses for multiple processor, shared memory configurations and yet be affordable in single processor configurations. SCI achieves this goal through high speed (1000 MB/s and 1000 Mb/s) point-to-point links. Transactions over these links are packet oriented with each packet containing an address, command and data. These links can be interconnected via a switch fabric or in a ring. Shared memory, multiprocessor configurations are supported through a distributed cache coherence scheme. This distributed scheme scales as the number of processors is increased [3]. The SCI was recently approved as an IEEE standard this year. As components become available, SCI will enable sophisticated distributed computing architectures to be assembled to meet the specific requirements of high speed data acquisition.

High Performance Parallel Interface

The High Performance Parallel Interface (HIPPI) standard was developed by the ANSI X3T9.3 committee and the physical layer interface was standardized in the Fall of 1991 [4]. HIPPI grew out of a requirement to standardize and improve I/O bandwidth for high performance computing environments. The effort began in 1987 and targeted 100 MB/s and 200 MB/s performance goals. Another goal was to quickly develop a standard and not require custom silicon for its implementation. The result was a specification for a parallel copper implementation with a distance limitation of 25M. This limitation can be overcome with the use of a commercially available fiber optic based HIPPI-extender product. A number of HIPPI developers developed an Serial-HIPPI specification that also addresses this limitation. HIPPI is widely available today on high performance computing systems and workstation implementations are also beginning to appear. HIPPI is an easy standard to implement and is useful for those applications that require high performance over short distances.

Fibre Channel

The Fibre Channel standard is under development by the ANSI X3T9.3 committee. This is an ambitious effort to support a broad performance range of peripheral devices and networks using multiple physical media. Fibre channel has a rich set of options including support for striping across multiple interfaces. The intended data rates range from 100 to 800 Mb/s. Fibre Channel requires custom silicon for its implementation and this coupled with the protracted effort to stabilize the standard has hampered the development of commercial products. Recently a Fibre Channel forum was formed by IBM and HP to promote the standard and encourage implementations.

Synchronous Optical Network

The Synchronous Optical Network (SONET) standard is the international optical telecommunications transmission standard. It supports a hierarchy of rates that are integer multiples of the basic rate of 51.84 Mb/s. Not all rates in the hierarchy are likely to be implemented. The more popular rates are 155.52 Mb/s, 622.08 Mb/s, and 2488 Mb/s. SONET is specified for three distance ranges with a corresponding reduction in cost of implementation. SONET does not specify any addressing scheme and consequently another mechanism must be used to switch SONET "payload". ATM is the mechanism that will be used in BISDN.

Asynchronous Transfer Mode

Asynchronous Transfer Mode is the switching and multiplexing mechanism for BISDN. ATM is based upon fast switching of small (53 Byte) cells. The small cell size reduces latency and provides a flexible scheme for multiplexing and allocating bandwidth. Each cell contains a 5 byte header and 48 byte data area.

ATM is also being deployed as a LAN architecture. ATM's capabilities to handle the requirements of multimedia communication and the direct connectivity to wide-area ATM based networks are the primary reasons to consider ATM in a local environment. The media that is being used for this LANs today is multimode fiber running at 100 Mb/s. The data rates are expected to increase to 622 Mb/s. An industry forum has formed to deal with interoperability issues.

Conclusions

In the near future, a wide choice of standard high performance interface technologies will be available. These technologies were developed by organizations with different goals and yet there is considerable overlap in their applications [6]. For example, the SCI can viewed both as a bus and a LAN. We are witnessing an integration of computing and communications technology with the "network" becoming the computer.

Acknowledgements

The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36. This work was performed under auspices of the U.S. Department of Energy.

References

1. L. Sha, R. Rajkumar, and J.P. Lehoczky, "Real-Time Computing with Futurebus+", *IEEE Micro*, June 1991, pgs 30-33, 95-100.
2. J.A. Gallant, "Futurebus+", *EDN*, October 1990, pgs 81-98.
3. D.B. Gustavson, "The Scalable Coherent Interface and Related Standards Projects", *IEEE Micro*, February 1992, pgs 10-22.
4. American National Standard for Information Systems, X3.183-1991, "High Performance Parallel Interface. Mechanical, Electrical, and Signalling Protocol Specification (HIPPI-PH)".
5. Network Compatible ATM for Local Network Applications, Phase 1, Version 1, April 3, 1992.
6. W. Andrews, "High-performance networks challenge Ethernet", Computer Design, July 1992, pgs 77-92.