US007039057B1

(12) **United States Patent**
Acharya et al.

(10) **Patent No.: US 7,039,057 B1**
(45) **Date of Patent: May 2, 2006**

(54) **ARRANGEMENT FOR CONVERTING ATM CELLS TO INFINIBAND PACKETS**

(75) Inventors: **Yatin Acharya**, Sunnyvale, CA (US); **Bahadir Erimli**, Campbell, CA (US)

(73) Assignee: **Advanced Micro Devices, Inc.**, Sunnyvale, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 859 days.

(21) Appl. No.: **09/907,586**

(22) Filed: **Jul. 19, 2001**

(51) **Int. Cl.**
H04L 12/28 (2006.01)

(52) **U.S. Cl.** **370/395.51**; 370/395.6; 370/466; 370/474; 370/401

(58) **Field of Classification Search** 370/395.1, 370/465, 395.52, 401, 466, 473, 474
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,275,494	B1 *	8/2001	Endo et al.	370/395.52
6,459,698	B1 *	10/2002	Acharya	370/392
6,522,667	B1 *	2/2003	Oda et al.	370/474
6,711,167	B1 *	3/2004	Ikeda et al.	370/395.1
6,788,706	B1 *	9/2004	Suzuki	370/474
6,799,220	B1 *	9/2004	Merritt et al.	709/238
6,826,196	B1 *	11/2004	Lawrence	370/466
2004/0128398	A1 *	7/2004	Pettey	709/249

OTHER PUBLICATIONS

William Stallings, "ISDN and Broadband ISDN with Frame Relay and ATM," 1998, Prentice-Hall, Inc, 4th Edition, pp. 438-450.*

Daniel Cassiday, InfiniBand™ Architecture Tutorial, Hot Chips, Aug. 2000, Sun Microsystems, 79 pages.

* cited by examiner

Primary Examiner—Ricky Q. Ngo
Assistant Examiner—Nittaya Juntima
(74) *Attorney, Agent, or Firm*—Manelli Denison & Selter PLLC; Leon R. Turkevich

(57) **ABSTRACT**

An ATM—InfiniBand™ router is configured for interfacing between an asynchronous transmission mode (ATM) network and an InfiniBand™ network, without a necessity of intermediate transport on a packet based network such as an Internet Protocol (IP) network. The router includes an ATM processor and a host channel adapter. The ATM processor is configured for generating ATM cells streams based on received InfiniBand™ packets, and recovering InfiniBand™ packet data from received ATM cells. The host channel adapter is configured for receiving the InfiniBand™ packets from the InfiniBand network and providing at least the payload data to the ATM processor, and outputting the recovered InfiniBand™ packet data onto the InfiniBand™ network. In addition, the ATM processor and the host channel adapter may be configured for mapping the ATM cells and the InfiniBand packets on prescribed virtual circuits and prescribed InfiniBand™ connections, respectively. Hence, the ATM—InfiniBand™ router operates as a call connection handler, enabling connections to be established across ATM and InfiniBand™ networks.

8 Claims, 3 Drawing Sheets

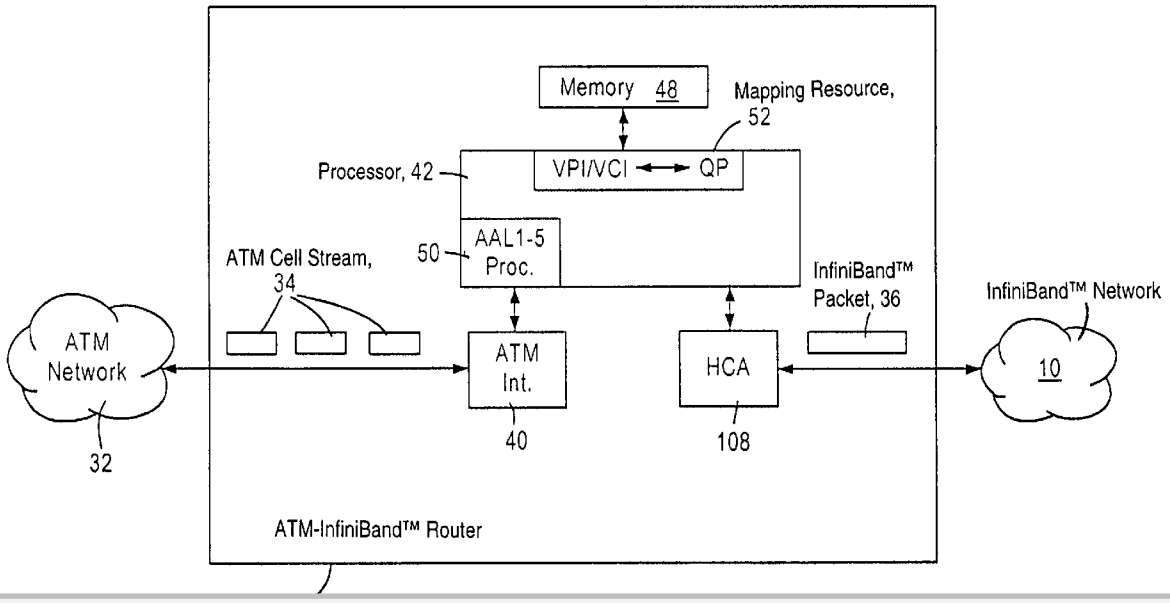

FIG. 1
(PRIOR ART)

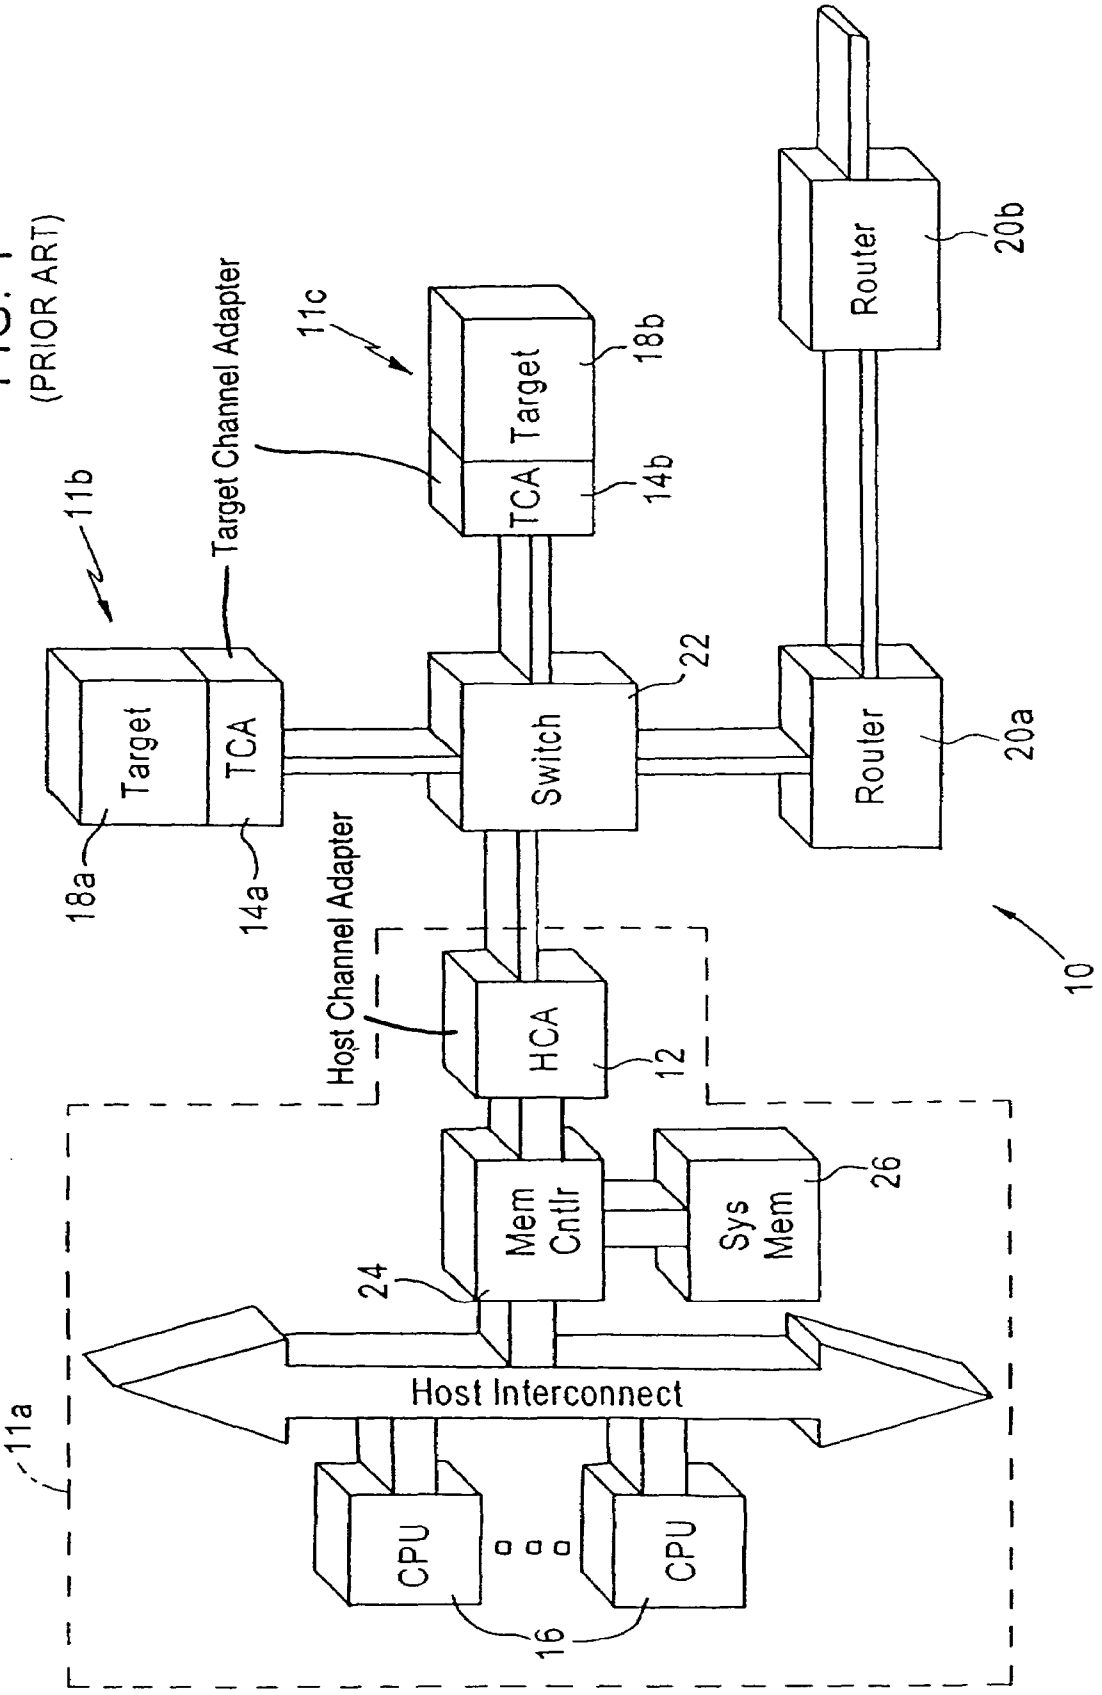

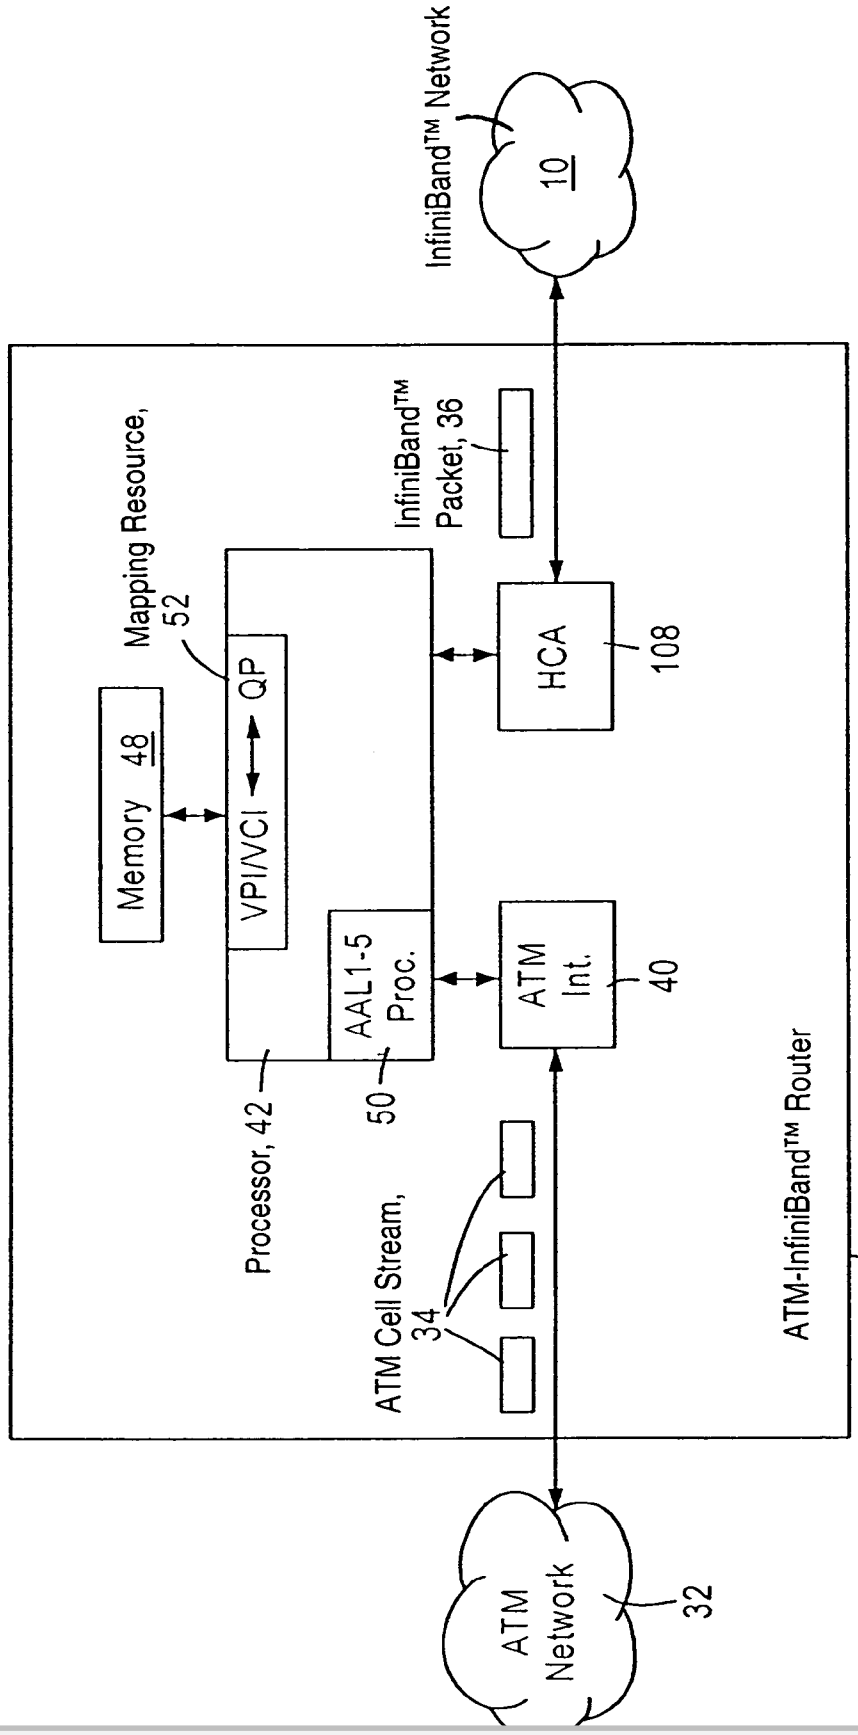

FIG. 2

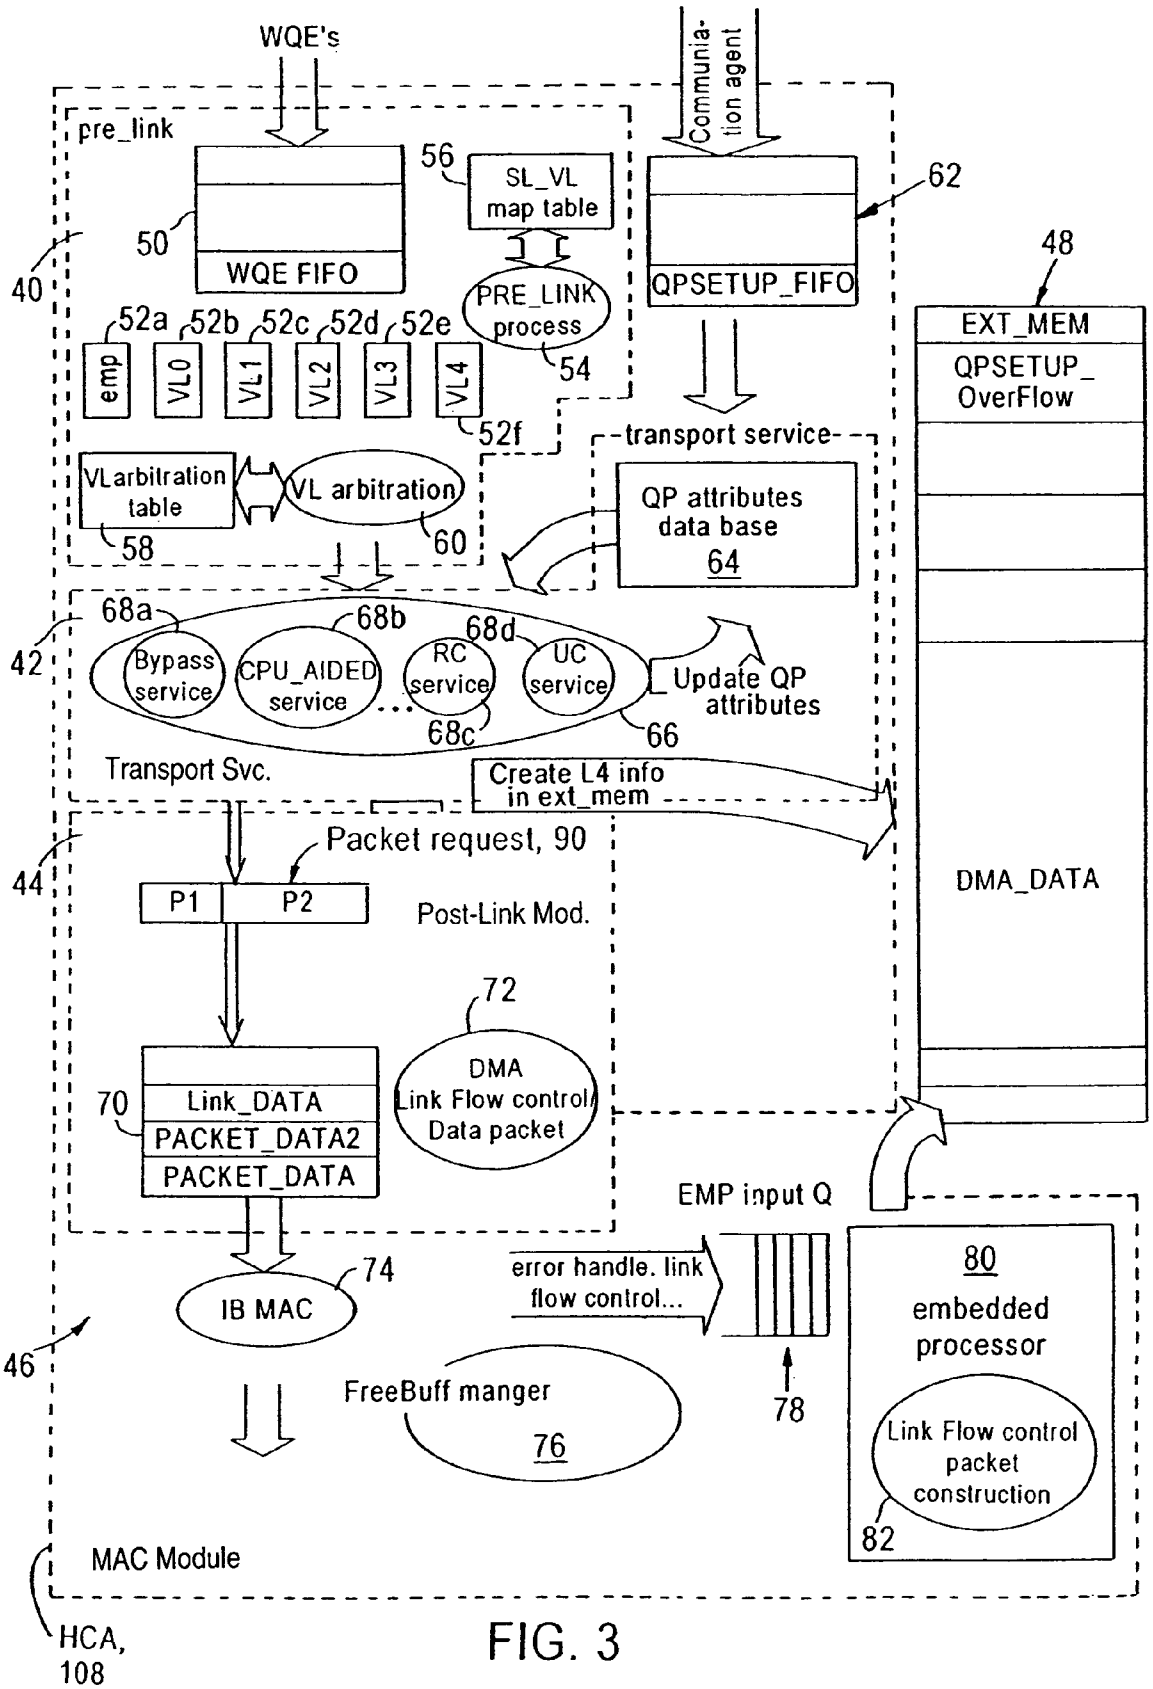

FIG. 3

ARRANGEMENT FOR CONVERTING ATM CELLS TO INFINIBAND PACKETS

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to an apparatus configured for forwarding data packets from a prescribed network, such as an Asynchronous Transmission Mode (ATM) network, to a destination node in an InfiniBand™ server system.

2. Background Art

Networking technology has encountered improvements in server architectures and design with a goal toward providing servers that are more robust and reliable in mission critical networking applications. In particular, the use of servers for responding to client requests has resulted in a necessity that servers have an extremely high reliability to ensure that the network remains operable. Hence, there has been a substantial concern about server reliability, accessibility, and serviceability.

In addition, processors used in servers have encountered substantial improvements, where the microprocessor speed and bandwidth have exceeded the capacity of the connected input/out (I/O) buses, limiting the server throughput to the bus capacity. Accordingly, different server standards have been proposed in an attempt to improve server performance in terms of addressing, processor clustering, and high-speed I/O.

These different proposed server standards led to the development of the InfiniBand™ Architecture Specification, (Release 1.0), adopted by the InfiniBand™ Trade Association. The InfiniBand™ Architecture Specification specifies a high-speed networking connection between central processing units, peripherals, and switches inside a server system. Hence, the term "InfiniBand™ network" refers to a network within a server system. The InfiniBand™ Architecture Specification specifies both I/O operations and interprocessor communications (IPC).

A particular feature of InfiniBand™ Architecture Specification is the proposed implementation in hardware of the transport layer services present in existing networking protocols, such as TCP/IP based protocols. The hardware-based implementation of transport layer services provides the advantage of reducing processing requirements of the central processing unit (i.e., "offloading"), hence offloading the operating system of the server system.

The InfiniBand™ Architecture Specification describes a network architecture, illustrated in FIG. **1**. The network **10** includes nodes **11** including channel adapters **12** or **14**; for example, the nodes **11** include processor nodes **16**, peripherals **18** such as Ethernet bridges or storage devices, routers **20**, and InfiniBand™switches **22**. Channel adapters operate as interface devices for respective server subsystems. For example, host channel adapters (HCAs) **12** are used to provide processor nodes **16** with an interface connection to the InfiniBand™ network **10**, and target channel adapters (TCAs) **14** are used to provide the peripherals **18** with an interface connection to the InfiniBand™ network. Host channel adapters **12** may be connected to a memory controller **24** as illustrated in FIG. **1**. Host channel adapters **12** implement the transport layer using a virtual interface referred to as the "verbs" layer that defines in the manner in which the processor **16** and the operating system communicate with the associated HCA **12**: verbs are data structures (e.g., commands) used by application software to communi-

respective devices **18** according to the respective device protocol (e.g., PCI, SCSI, etc.).

The InfiniBand™ Architecture Specification requires that a packet to be sent via an HCA **12** undergoes transport layer service, followed by link layer service. Examples of operations performed during transport layer service include constructing a transport layer header, generating a packet sequence number, validating service type, etc. Examples of operations performed during link layer service include service layer and virtual layer mapping (SL-VL mapping), link layer flow control packet generation, link layer transmission credit checking, etc.

However, arbitrary hardware implementations may result in substantially costly hardware designs. In particular, problems may arise when attempting to deploy an InfiniBand™ network to send and receive data between other networks. For example, wide area networks (WAN) often rely on ATM switching technology to handle information transfer within a single network or between networks. ATM is a connection oriented, cell based switching technology that uses 53-byte cells to transport information. However, conventional approaches to connecting the InfiniBand™ network to a Metropolitan Area Network (MAN) or a Wide Area Network (WAN) may require an Asynchronous Transmission Mode (ATM) edge device that recovers an IP packet from ATM cells using a prescribed adaptation layer processing. The recovered IP packet then would be sent via an IP network to an IP router having a presence on the InfiniBand™ network. Such an arrangement can be substantially costly for smaller business consumers. In addition, disadvantages arise from converting between a connection-based transport protocols (e.g., ATM and InfiniBand™) and packet-based protocols such as Internet Protocol.

SUMMARY OF THE INVENTION

There is a need for an arrangement that enables an InfiniBand™ network to be connected to a Metropolitan Area Network, or a Wide Area Network utilizing connection-based transport protocols, in an efficient and economical manner.

There also is a need for an arrangement that enables an InfiniBand™network to be connected to a Metropolitan Area Network or a Wide Area Network by maintaining a connection-based transport protocol.

These and other needs are attained by the present invention, where an ATM—InfiniBand™ router is configured for interfacing between an asynchronous transmission mode (ATM) network and an InfiniBand™ network, without a necessity of intermediate transport on a packet based network such as an Internet Protocol (IP) network. The router includes an ATM processor and a host channel adapter. The ATM processor is configured for generating ATM cells streams based on received InfiniBand™ packets, and recovering InfiniBand™ packet data from received ATM cells. The host channel adapter is configured for receiving the InfiniBand™ packets from the InfiniBand network and providing at least the payload data to the ATM processor, and outputting the recovered InfiniBand™ packet data onto the InfiniBand™ network. In addition, the ATM processor and the host channel adapter may be configured for mapping the ATM cells and the InfiniBand packets on prescribed virtual circuits and prescribed InfiniBand™ connections, respectively. Hence, the ATM—InfiniBand™ router oper-

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.