Attorney Docket No. NWISP013

PATENT APPLICATION

ROUTING MECHANISMS IN SYSTEMS HAVING MULTIPLE MULTI-PROCESSOR CLUSTERS

Inventors: David Brian Glasco
10337 Ember Glen Drive
Austin, Texas 78726
United States citizen

Carl Zeitler
11835 Brush Canyon Drive
Tomball, TX 77375
United States citizen

Rajesh Kota
12320 Alamaeda Trace Circle #1107
Austin, Texas 78727
India citizen

Guru Prasadh
11111 Callanish Park Drive
Austin, Texas 78750
United States citizen

Richard R. Oehler
8 Bonny Drive
Somers, New York 10589
United States citizen

Assignee: Newisys, Inc.
A Delaware corporation

BEYER WEAVER & THOMAS, LLP
P.O. Box 778
Berkeley, California 94704-0778
(510) 843-6200

PATENT

Attorney Docket No. NWISP013

ROUTING MECHANISMS IN SYSTEMS HAVING MULTIPLE MULTI-PROCESSOR CLUSTERS

BACKGROUND OF THE INVENTION

The present invention relates generally to multi-processor computer systems. More specifically, the present invention provides techniques for building computer systems having a plurality of multi-processor clusters.

A relatively new approach to the design of multi-processor systems replaces broadcast communication among processors with a point-to-point data transfer mechanism in which the processors communicate similarly to network nodes in a tightly-coupled computing system. That is, the processors are interconnected via a plurality of communication links and requests are transferred among the processors over the links according to routing tables associated with each processor. The intent is to increase the amount of information transmitted within a multi-processor platform per unit time.

One limitation associated with such an architecture is that the node ID address space associated with the point-to-point infrastructure is fixed, therefore allowing only a limited number of nodes to be interconnected. In addition, the infrastructure is flat, therefore allowing a single level of mapping for address spaces and routing functions. It is therefore desirable to provide techniques by which computer systems employing such an infrastructure as a basic building block are not so limited.

SUMMARY OF THE INVENTION

According to the present invention, a multi-processor system is provided in which a plurality of multi-processor clusters, each employing a point-to-point communication infrastructure with a fixed node ID space and flat request mapping functions, are interconnected using additional point-to-point links in such a manner as to enable more processors to be interconnected than would otherwise be possible with the local point-to-point architecture. The invention employs a mapping hierarchy to uniquely map various types of information from local, cluster-specific spaces to globally shared spaces.

Thus, the present invention provides an interconnection controller for use in a computer system having a plurality of processor clusters interconnected by a plurality of global links. Each cluster includes a plurality of local nodes and an instance of the interconnection controller interconnected by a plurality of local links. The interconnection controller includes circuitry which is operable to map locally generated transmissions directed to others of the clusters to the global links, and remotely generated transmissions directed to the local nodes to the local links. According to a specific embodiment, a computer system employing such an interconnection controller is also provided.

A further understanding of the nature and advantages of the present invention may be realized by reference to the remaining portions of the specification and the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Figs. 1A and 1B are diagrammatic representations depicting systems having multiple clusters.

Fig. 2 is a diagrammatic representation of an exemplary cluster having a plurality of processors for use with specific embodiments of the present invention.

Fig. 3 is a diagrammatic representation of an exemplary interconnection controller for facilitating various embodiments of the present invention.

Fig. 4 is a diagrammatic representation of a local processor for use with various embodiments of the present invention.

Fig. 5 is a diagrammatic representation of a memory mapping scheme according to a particular embodiment of the invention.

Fig. 6A is a simplified block diagram of a four cluster system for illustrating a specific embodiment of the invention.

Fig. 6B is a combined routing table including routing information for the four cluster system of Fig. 6A.

Figs. 7 and 8 are flowcharts illustrating transaction management in a multi-cluster system according to specific embodiments of the invention.

Fig. 9 is a diagrammatic representation of communications relating to an exemplary transaction in a multi-cluster system.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

Reference will now be made in detail to some specific embodiments of the invention including the best modes contemplated by the inventors for carrying out the invention. Examples of these specific embodiments are illustrated in the accompanying drawings. While the invention is described in conjunction with these specific embodiments, it will be understood that it is not intended to limit the invention to the described embodiments. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims. Multi-processor architectures having point-to-point communication among their processors are suitable for implementing specific embodiments of the present invention. In the following description, numerous specific details are set forth in order to provide a thorough understanding of the present invention. The present invention may be practiced without some or all of these specific details. Well known process operations have not been described in detail in order not to unnecessarily obscure the present invention. Furthermore, the present application's reference to a particular singular entity includes that possibility that the methods and apparatus of the present invention can be implemented using more than one entity, unless the context clearly dictates otherwise.

Fig. 1A is a diagrammatic representation of one example of a multiple cluster, multiple processor system which may employ the techniques of the present invention. Each processing cluster 101, 103, 105, and 107 includes a plurality of processors. The processing clusters 101, 103, 105, and 107 are connected to each other through point-to-point links 111a-f. The multiple processors in the multiple cluster architecture shown in Fig. 1A share a global memory space. In this example, the point-to-point links 111a-f are internal system connections that are used in place of a traditional front-side bus to connect the multiple

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.