IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of: Morton et al.
U.S. Patent No.: 7,296,121
Issue Date: Nov. 13, 2007
Appl. Serial No.: 10/966,161
Filing Date: Oct. 15, 2004
Title: REDUCING PROBE TRAFFIC IN MULTIPROCESSOR SYSTEMS

Case Nos.: IPR2015-00159
IPR2015-00163

OPPOSITION DECLARATION OF DR. ROBERT HORST

1. I have reviewed the "Patent Owner Motion to Amend" in IPR2015-00159, the "Patent Owner Motion to Amend" in IPR2015-00163 and the "Declaration of Vojin Oklobdzija, Ph.D. in Support of Patent Owner's Motion to Amend," each filed on August 11, 2015. I also considered the references cited herein, including, for example: Michael John Sebastian Smith, APPLICATION-SPECIFIC INTEGRATED CIRCUITS (1997) ("Smith") (Ex. 1008); Deposition Transcript of Dr. Vojin G. Oklobdzija Vol. 1, November 23, 2015 (Ex. 1026); Deposition Transcript of Dr. Vojin G. Oklobdzija Vol. 2, November 24, 2015 (Ex. 1027); David E. Culler et al., PARALLEL COMPUTER ARCHITECTURE: A HARDWARE/SOFTWARE APPROACH (1st Ed.) (1998) (Ex. 1028); "InfiniBand Architecture Specification Volume 1 Release 1.0.a" (June 19, 2001) (Ex. 1029); and James Laudon and Daniel Lenoski, Proceedings of the 24th Annual International Symposium on Computer Architecture, "The SGI Origin: A ccNUMA Highly Scalable Server" (1997) (Ex. 1030). In my declaration, I am applying the

standards and legal principles that I applied when drafting the declaration entitled "Declaration of Dr. Robert Horst" dated October 28, 2014, which were outlined in paragraphs 8 and 40-61 of that document. Based on these principles and my expertise in the relevant technology, I provide the following description of prior art relevant to the amended claims.

I. The Culler Book and Laudon (Relevant to Claims 26-28)

2. The Culler Book includes a case study of the SGI Origin architecture. *See* Ex. 1028, p. 596. Similarly, Laudon is titled "The SGI Origin: A ccNUMA Highly Scalable Server" and describes the same SGI Origin architecture as described in the Culler Book's case study. Ex. 1030, p. 1. Both the Culler Book and Laudon share similar system diagrams, particularly with regard to the Hub chip, which acts as the probe filtering unit in the SGI Origin architecture and is shown identically in FIG. 8.21 of the Culler Book and FIG. 6 of Laudon. *See* Ex. 1028, p. 616; Ex. 1030, p. 245. Accordingly, it would have been obvious to a person of ordinary skill in the art to combine the teachings of the Culler Book and Laudon, as the combination of these references would have provided a more complete and confirmatory teaching of the SGI Origin architecture.

3. The Culler Book describes that "[t]he Origin system is composed of a number of processing nodes connected by switch-based interconnection network." Ex. 1028, p. 597. As shown in FIG. 3 of Laudon (reproduced below), "[t]he

interconnection network has a hypercube topology" is one form of the "scalable point-to-point interconnection network[s]" on which the directory-based coherence protocols of Chapter 8 of the Culler Book are based. Ex. 1028, pp. 553, 597, 615; Ex. 1030, p. 243; *see also* Ex. 1026, 132:24-134:1 (admitting that a switch-based network is a "point-to-point architecture"). According to the Culler Book, "[e]very processing node contains two MIPS R10000 processors, each with first- and second-level caches, a fraction of the total main memory of the machine, an I/O interface, and a single communication assist or cache coherence controller, called the Hub, that implements the coherence protocol." Ex. 1028, p. 597. Though the Culler Book describes each processing node as containing two processors, Laudon describes that "[e]ach node consists of one or two R10000 processors," and the Culler Book even "assumes for simplicity that each node contains only one processor" when discussing the cache coherence protocol. Ex. 1030, p. 241; Ex. 1028, p. 597. Accordingly, it would have been obvious to a POSITA to implement an SGI Origin machine in which each of multiple nodes has only a single processor, and the following section applies such a configuration to claims 26-34.

Point-to-Point Link Between Hub Chips (i.e., a Point-to-Point Architecture)

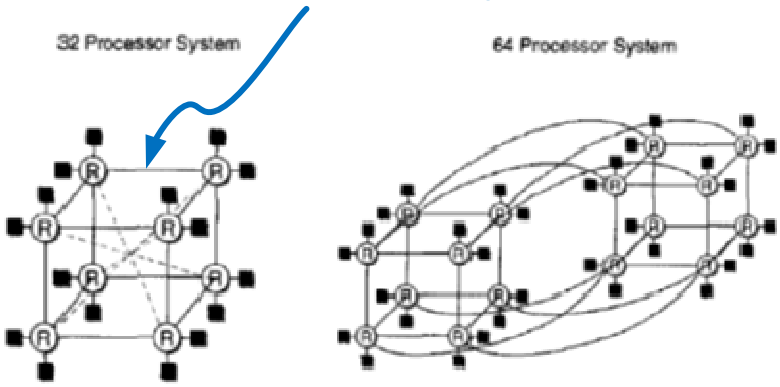

Figure 3 32P and 64P Bristled Hypercubes

Ex. 1030, p. 243 (showing point-to-point links between Hub chips, though, in the proposed combination, each "R" is a Hub chip and would only be associated with a single processor).

4. Moreover, an SGI Origin system may contain "up to 512 nodes." Ex. 1028, p. 612. The following discussion assumes a combination of the Culler Book and Laudon with four single-processor nodes. However, the proposed combination would be equally applicable for systems containing more or less than four nodes. Following is an adaptation of FIG. 8.15 of the Culler Book that illustrates the proposed combination and annotates the relevant components.

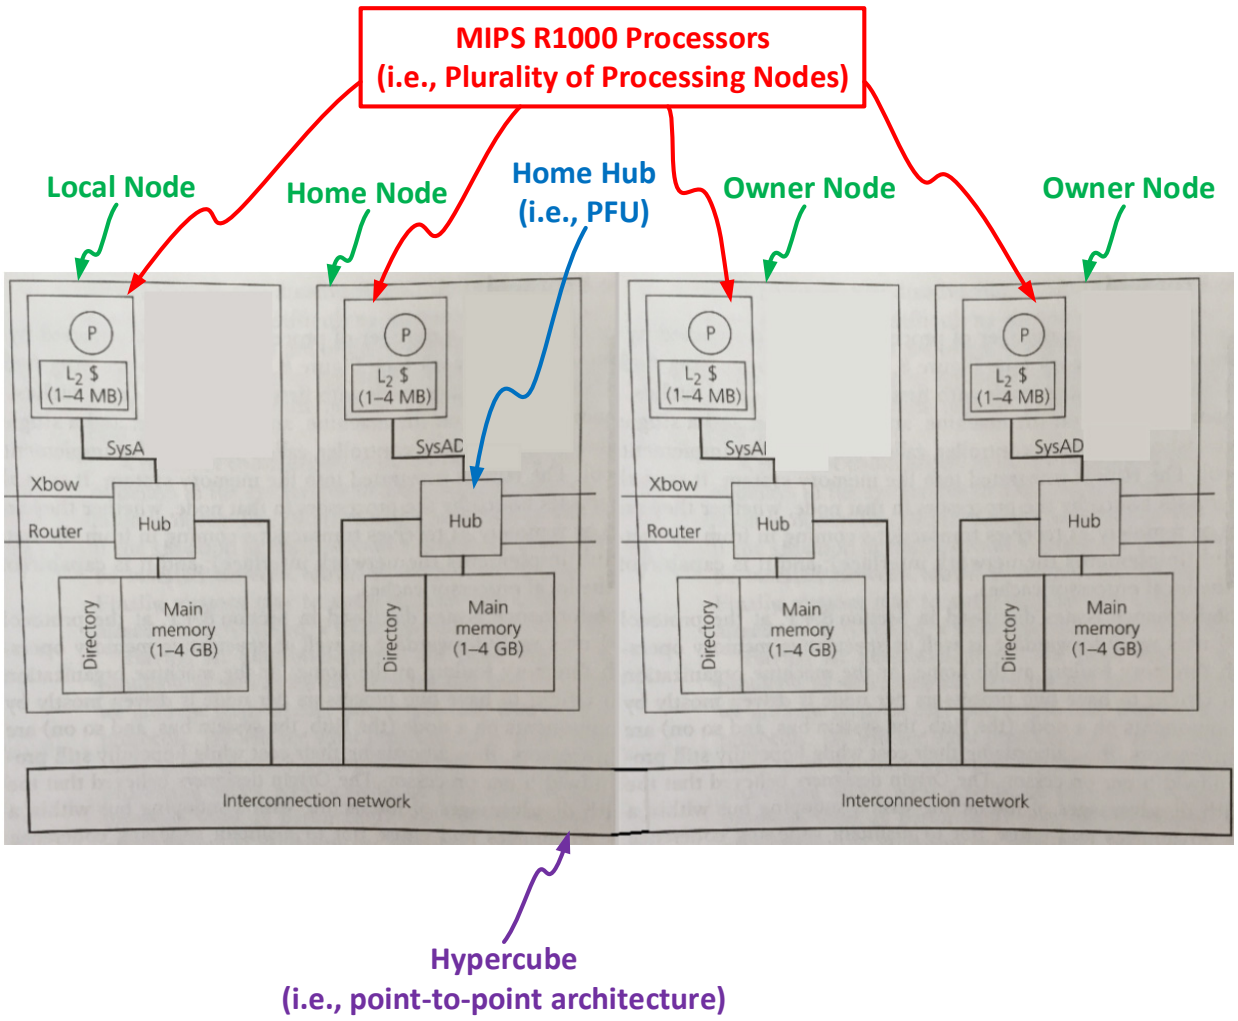

5. As described in the Culler Book, each of the processors in the SGI Origin architecture is associated with "first- and second-level caches." Ex. 1028, pp. 597, 612. The caches of an SGI Origin machine "use[] the same MESI states as used in Chapter 5" of the Culler Book, which are "modified (M) or dirty, exclusive-clean (E), shared (S), and invalid (I)." Ex. 1028, pp. 598, 299. In the directory referenced by the Hub chip, these MESI states are represented by one of seven directory states. Ex. 1028, p. 598. For example, the "shared" directory state indicates "zero or more read-only cached copies whose whereabouts are indicated

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.