

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
29 July 2004 (29.07.2004)

PCT

(10) International Publication Number
WO 2004/064413 A2

(51) International Patent Classification⁷:

H04Q

(21) International Application Number:

PCT/US2003/039609

(22) International Filing Date:

12 December 2003 (12.12.2003)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

10/340,390 10 January 2003 (10.01.2003) US

(71) Applicant: SRC COMPUTERS, INC. [US/US]; 4240 North Nevada Avenue, Colorado Springs, CO 80907 (US).

(72) Inventors: HUPPENTHAL, Jon, M.; 10015 Burgess Road, Colorado Springs, CO 80908 (US). SEEMAN, Thomas, R.; 5765 Huffman Court, Colorado Springs, CO 80919 (US). BURTON, Lee, A.; 966 Cedar Mountain Road, Divide, CO 80814 (US).

(74) Agents: BURTON, Carol, W. et al.; Hogan & Hartson LLP, 1200 17th Street, Suite 1500, Denver, CO 80202 (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

WO 2004/064413 A2

(54) Title: SWITCH/NETWORK ADAPTER PORT COUPLING A RECONFIGURABLE PROCESSING ELEMENT TO ONE OR MORE MICROPROCESSORS FOR USE WITH INTERLEAVED MEMORY CONTROLLERS

(57) Abstract: A switch/network adapter port ("SNAPTM") in a dual in-line memory module ("DIMM") or RambusTM in-line memory module ("RIMM") format for clustered computers employing multi-adaptive processor ("MAP²™", both trademarks of SRC Computers, Inc.) elements for use with interleaved memory controllers. Particularly disclosed is a microprocessor based computer system utilizing either a DIMM or RIMM physical format adapter port coupled to a reconfigurable processor element for the purpose of implementing a connection to an external switch, network, or other device. In a particular embodiment, connections may be provided to either the PCI, accelerated graphics port ("AGP") or system maintenance ("SM") bus for purposes of passing control information to the host microprocessor or other control chips. The field programmable gate array ("FPGA") based processing elements have the capability to alter data passing through it to and from an external interconnect fabric or device.

SWITCH/NETWORK ADAPTER PORT COUPLING A RECONFIGURABLE PROCESSING ELEMENT TO ONE OR MORE MICROPROCESSORS FOR USE WITH INTERLEAVED MEMORY CONTROLLERS

CROSS REFERENCE TO RELATED PATENT APPLICATIONS

5 The present invention claims priority from U.S. Patent Application Serial No. 10/340,390 filed January 10, 2003 having the same title as the present case, which contains subject matter which is related to U.S. Patent Application Serial No. 09/932,330 filed August 17, 2001 for "Switch/Network Adapter Port for Clustered Computers Employing a Chain of Multi-Adaptive
10 Processors in a Dual In-Line Memory Module Format" and U.S. Provisional Patent Application Serial No. 60/422,722 filed October 31, 2002 for "General Purpose Reconfigurable Computing Hardware and Software", all of which are assigned to SRC Computers, Inc., Colorado Springs, CO, the disclosures of which are specifically incorporated herein in their entirety by this reference.

15 BACKGROUND OF THE INVENTION

Field of the Invention.

The present invention relates, in general, to the field of computer systems and methods incorporating one or more reconfigurable processing elements. More particularly, the present invention relates to a switch/network adapter port ("SNAPTM") in a dual in-line memory module ("DIMM") or Rambus ("RIMM") format for a computing system employing multi-adaptive processing elements ("MAP[®]", both trademarks of SRC Computers, Inc.) for use with interleaved memory controllers in order to provide enhanced data transfer rates.

25 Relevant Background.

Among the most currently promising methods of creating large processor count, cost-effective computers involves the clustering together of a number of relatively low cost microprocessor based boards such as those commonly found in personal computers ("PCs"). These various boards are then operated using available clustering software to enable them to execute, in unison, to solve one or more large problems. During this problem solving

process, intermediate computational results are often shared between processor boards.

Utilizing currently available technology, this sharing must pass over the peripheral component interconnect ("PCI") bus, which is the highest 5 performance external interface bus, commonly found on today's PCs. While there are various versions of this bus available, all are limited to less than 1 GB/sec. bandwidth and, because of their location several levels of chips below the processor bus, they all exhibit a very high latency. In low cost PCs, this bus typically offers only on the order of 256 MB/sec. of bandwidth.

10 These factors, both individually and collectively can significantly limit the overall effectiveness of the cluster and, if a faster interface could be found, the ability of clusters to solve large problems would be greatly enhanced. Unfortunately, designing a new, dedicated chip set that could provide such a port is not only very expensive, it would also have to be 15 customized for each type of clustering interconnect encountered. This would naturally lead to relatively low potential sale volumes for any one version of the chipset, thus rendering it cost ineffective.

With ever-increasing processor speeds, the need for high performance memory subsystems has also continued to increase. Since the development 20 of the Switch/Network Adapter Port system as disclosed in the aforementioned U.S. Patent Application Serial No. 09/932,330, the technology for high performance memory subsystems for the personal computer ("PC") market has come to include the use of interleaved memory.

In an interleaved memory system, two or more dual in-line memory 25 module ("DIMM") slots are accessed by the memory controller at the same time. When a by-two interleaving scheme is used, the width of the data bus is effectively doubled, thus doubling the bandwidth that is obtained to memory. A similar configuration can be established to form a by-four, or four way, interleaved system using four DIMM slots. This form of memory 30 controller is currently one of the more common high performance memories

found in the higher end server systems and is rapidly becoming available in more mainstream products.

SUMMARY OF THE INVENTION

As disclosed in a representative embodiment herein, SRC Computers, Inc. proprietary SNAP™ technology has been enhanced such that the signals from two or more DIMM (or RIMM) slots are routed to a common control chip. Physically, in a by-two configuration, two DIMM form factor switch/network adapter boards may be coupled together using ridged flex circuit construction to form a single assembly. One of the DIMM boards may also be populated with a control field programmable gate array ("FPGA") which may have the signals from both DIMM slots routed to it. The control chip then samples the data off of both slots using the independent clocks of the slots. The data from both slots is then used to form a data packet that is then sent to other parts of the system. In a similar manner, the technique disclosed herein may be utilized in conjunction with more than two DIMM slots, for example, four DIMM slots is an four-way interleaved system.

In accordance with the present invention, an interleaved memory system uses two or more memory channels running in lock-step. A connection is made to one of the DIMM slots and the signals derived are used in conjunction with the original set of SNAP™ board signals. In operation, this effectively doubles (or more) the width of the data bus into and out of the memory. This technique can be implemented in conjunction with the proper selection of a memory and input/output ("I/O") controller ("North Bridge") chip that supports interleaved memory.

Particularly disclosed herein is a computer system comprising at least one processor, a controller for coupling the processor to a control block and a memory bus, a plurality of memory module slots coupled to the memory bus, an adapter port associated with a subset of the plurality of memory module slots and a processor element coupled to the adapter port. In a preferred embodiment, the adapter port may be conveniently provided in a DIMM or RIMM form factor.

Also disclosed herein is a computer system comprising at least one processor, an interleaved controller for coupling the processor to a control block and a memory bus, a plurality of memory slots coupled to the memory bus, an adapter port associated with at least two of the memory slots and a 5 processor element coupled to each of the adapter ports.

BRIEF DESCRIPTION OF THE DRAWINGS

The aforementioned and other features and objects of the present invention and the manner of attaining them will become more apparent and the invention itself will be best understood by reference to the following 10 description of a preferred embodiment taken in conjunction with the accompanying drawings, wherein:

Fig. 1 is a functional block diagram of an exemplary embodiment of a switch/network adapter port for clustered computers employing a chain of multi-adaptive processors in a DIMM format to significantly enhance data 15 transfer rates over that otherwise available from the peripheral component interconnect ("PCI") bus;

Fig. 2A is a functional block diagram of an exemplary embodiment of a switch/network adapter port in accordance with the present invention illustrating a by-two configuration of interleaved DIMM slot form factor SNAP 20 elements coupled to a common control element;

Fig. 2B is a further functional block diagram of another exemplary embodiment of a switch/network adapter port in accordance with the present invention illustrating a by-four configuration of interleaved DIMM slot form factor SNAP elements coupled to a common control element;

25 Fig. 3 is a high level functional block diagram of an SRC-6E computer system available from SRC Computers, Inc. which may incorporate the switch/network adapter port elements of the preceding figures and in which each processor board is coupled to its own MAP® element;

Fig. 4 is a more detailed functional block diagram of a portion of the 30 SRC-6E computer system of the preceding figure illustrating one-half of a processor board in conjunction with one half of a MAP® board;

Fig. 5 is a further, more detailed functional block diagram of the MAP® element of the preceding figure illustrating the control FPGA, on-board memory and user FPGAs in particular;

Fig. 6 is a further, more detailed functional block diagram of the control 5 block, or chip, of the MAP® board illustrated in Figs. 4 and 5;

Fig. 7 is an additional, more detailed functional block diagram of a representative switch/network adapter port for use with the computer system illustrated in Figs. 3 and 4; and

Fig. 8 is an illustration showing the interaction between the MAP® 10 control block, switch/network adapter port and microprocessor in the operation of a computer system utilizing a switch/network adapter port in accordance with the present invention.

DESCRIPTION OF A REPRESENTATIVE EMBODIMENT

With reference now to Fig. 1, a functional block diagram of an 15 exemplary embodiment of a computer system 100 is shown comprising a switch/network adapter port for clustered computers employing a chain of multi-adaptive processors in a DIMM format to significantly enhance data transfer rates over that otherwise available from the peripheral component interconnect ("PCI") bus.

20 In the particular embodiment illustrated, the computer system 100 includes one or more processors 102₀ and 102₁ which are coupled to an associated PC memory and I/O controller 104. In operation, the controller 104 sends and receives control information from a PCI control block 106. It should be noted that in alternative implementations of the present invention, 25 the control block 106 may also be an accelerated graphics port ("AGP") or system maintenance ("SM") control block. The PCI control block 106 is coupled to one or more PCI card slots 108 by means of a relatively low bandwidth PCI bus 110 which allows data transfers at a rate of substantially 256 MB/sec. In the alternative embodiments of the present invention 30 mentioned above, the card slots 108 may alternatively comprise accelerated graphics port ("AGP") or system maintenance ("SM") bus connections.

The controller 104 is also conventionally coupled to a number of DIMM slots 114 by means of a much higher bandwidth DIMM bus 116 capable of data transfer rates of substantially 2.1 GB/sec. or greater. In accordance with a particular implementation of the present invention, a DIMM MAP® element 112 is associated with, or physically located within, one of the DIMM slots 114. Control information to or from the DIMM MAP® element 112 is provided by means of a connection 118 interconnecting the PCI bus 110 and the DIMM MAP® element 112. The DIMM MAP® element 112 then may be coupled to another clustered computer MAP® element by means of a cluster 10 interconnect fabric connection 120 connected to MAP® chain ports as will be more fully described hereinafter. It should be noted that, the DIMM MAP® element 12 may also comprise a Rambus DIMM ("RIMM") MAP® element.

Since the DIMM memory located within the DIMM slots 114 comprises the primary storage location for the PC microprocessor(s) 102₀, 102₁, it is 15 designed to be electrically very "close" to the processor bus and thus exhibit very low latency. As noted previously, it is not uncommon for the latency associated with the DIMM to be on the order of only 25% of that of the PCI bus 110. By, in essence, harnessing this bandwidth as an interconnect between computer systems 100, greatly increased cluster performance may 20 be realized.

To this end, by placing the DIMM MAP® element 112 in one of the PC's DIMM slots 114, its control chip (as will be more fully described hereinafter) could accept the normal memory "read" and "write" transactions and convert them to a format used by an interconnect switch or network. As 25 will also be more fully described hereinafter, each MAP® element 112 includes chain ports to enable it to be coupled to other MAP® elements 112. Through the utilization of the chain port to connect to the external clustering fabric over connection 120, data packets can then be sent to remote nodes where they can be received by an identical board. In this particular 30 application, the DIMM MAP® element 112 would extract the data from the packet and store it until needed by the receiving processor 102.

This technique results in the provision of data transfer rates several times higher than that of any currently available PC interface such as the PCI bus 110. However, the electrical protocol of the DIMMs is such that once the data arrives at the receiver, there is no way for a DIMM module within the

5 DIMM slots 114 to signal the microprocessor 102 that it has arrived, and without this capability, the efforts of the processors 102 would have to be synchronized through the use of a continued polling of the DIMM MAP® elements 112 to determine if data has arrived. Such a technique would totally consume the microprocessor 102 and much of its bus bandwidth thus stalling

10 all other bus agents.

To avoid this situation, the DIMM MAP® element 112 may be further provided with the connection 118 to allow it to communicate with the existing PCI bus 110 which could then generate communications packets and send them via the PCI bus 110 to the processor 102. Since these packets would

15 account for but a very small percentage of the total data moved, the low bandwidth effects of the PCI bus 110 are minimized and conventional PCI interrupt signals could also be utilized to inform the processor 102 that data has arrived. In accordance with another implementation of the present invention, the system maintenance ("SM") bus (not shown) could also be

20 used to signal the processor 102. The SM bus is a serial current mode bus that conventionally allows various devices on the processor board to interrupt the processor 102. In an alternative embodiment, the accelerated graphics port ("AGP") may also be utilized to signal the processor 102.

With a DIMM MAP® element 112 associated with what might be an

25 entire DIMM slot 114, the PC will allocate a large block of addresses, typically on the order of 1 GB, for use by the DIMM MAP® element 112. While some of these can be decoded as commands, many can still be used as storage. By having at least as many address locations as the normal input/output ("I/O") block size used to transfer data from peripherals, the conventional

30 Intel™ chip sets used in most PCs (including controller 104) will allow direct I/O transfers into the DIMM MAP® element 112. This then allows data to arrive from, for example, a disk and to pass directly into a DIMM MAP®

element 112. It then may be altered in any fashion desired, packetized and transmitted to a remote node over connection 120. Because both the disk's PCI bus 110 and the DIMM MAP® element 112 and DIMM slots 114 are controlled by the PC memory controller 104, no processor bus bandwidth is 5 consumed by this transfer.

It should also be noted that in certain PCs, several DIMMs within the DIMM slots 114 may be interleaved to provide wider memory access capability in order to increase memory bandwidth. In these systems, the 10 previously described technique may also be utilized concurrently in several DIMM slots 114. Nevertheless, regardless of the particular implementation chosen, the end result is a DIMM-based MAP® element 112 having one or more connections to the PCI bus 110 and an external switch or network over connection 120 which results in many times the performance of a PCI-based connection alone as well as the ability to process data as it passes through 15 the interconnect fabric.

With reference additionally now to Fig. 2A, a functional block diagram of an exemplary embodiment of a switch/network adapter port 200A in accordance with the present invention is shown. Like structure and functionality to that disclosed with respect to the foregoing figure is like 20 numbered and the foregoing description thereof shall suffice herefor. The switch/network adapter port 200A is shown in a by-two configuration of interleaved DIMM slot form factor SNAP elements 204A and 204B coupled to a common control element 202. In this embodiment, the controller 104 is an interleaved memory controller bi-directionally coupled to the DIMM slots 114 25 and SNAP elements 204A, 204B by means of a Channel A 216A and a Channel B 216B.

With reference additionally now to Fig. 2B, a functional block diagram of another exemplary embodiment of a switch/network adapter port 200B in accordance with the present invention is shown. Again, like structure and 30 functionality to that disclosed with respect to the preceding figures is like numbered and the foregoing description thereof shall suffice herefor. The switch/network adapter port 200B is shown in a a by-four configuration of

interleaved DIMM slot form factor SNAP elements 204A through 204D coupled to a common control element 202. In this embodiment, the controller 104 is again an interleaved memory controller bi-directionally coupled to the DIMM slots 114 and SNAP elements 204A, 204B, 204C and 204D by means 5 of respective Channel A 216A, Channel B 216B, Channel C 216C and Channel D 216D.

With reference additionally now to Fig. 3, a high level functional block diagram of an SRC-6E computer system 300 available from SRC Computers, Inc. is shown. The computer system 300 may incorporate the switch/network adapter port elements 204 of Figs. 2A and 2B and comprises a number of 10 MAP® elements 302₀ through 302_N, each coupled to a corresponding microprocessor board 306₀ through 306_N by means of a respective switch/network adapter port 304₀ through 304_N. A number of chain ports 308 with a bandwidth of 1200 MB/Sec. or more, couple the various MAP® 15 elements 302₀ through 302_N to others of the MAP® elements 302₀ through 302_N. In operation, the computer system 300 provides equal MAP® element 302 to processor utilization with low data sharing, high MAP® element 302 input/output (“I/O”) utilization and very low inter-processor communication.

With reference additionally now to Fig. 4, a more detailed functional 20 block diagram of a portion of the SRC-6E computer system 300 of the preceding figure is shown as a computing system 400 and illustrating one-half of a processor board 306 in conjunction with one half of a MAP® board 302. As illustrated, the processor board 306 comprises a pair of microprocessors 402₀ and 402₁, each coupled to an associated level 2 (“L2”) cache 404 and a 25 memory and I/O controller (e.g. “North Bridge”) 406. The controller 406 provides access to, for example, one or more peripheral component interconnect (“PCI”) slots 408 and a private memory 410. The private memory 410 may comprise, for example, interleaved DIMM or RIMM slots having an associated switch/network adapter port 304 as will be more fully 30 described hereinafter.

The MAP® element 302 includes a control block, or control chip 412 which is coupled to and controls the switch/network adapter port 304 in the

processor board 306. The control chip 412, which may comprise a field programmable gate array ("FPGA"), is coupled to an on-board memory 414 which may comprise, for example, dual-ported static random access memory ("SRAM") or other types of memory devices. The on-board memory 414 is 5 also coupled to one or more user chips (which may also be provided as FPGAs) 416₀ and 416₁.

With reference additionally now to Fig. 5, a further, more detailed functional block diagram of the MAP[®] element 302 of the preceding figure is shown. In this view, a functional block 500 is illustrated showing in greater 10 detail the control FPGA 412, on-board memory 414 and user FPGAs 416₀ and 416₁ and the bandwidths of the various connections therebetween. A microcode read only memory ("ROM") 502 and configuration ROM 504 for specifying the operation of the control FPGA 412 are also illustrated as well as a number of Joint Test Action Group ("JTAG") Institute of Electrical and 15 Electronic Engineering (IEEE 1149.1) standard boundary scan connections. A number of chain ports 308 are shown for directly coupling various MAP[®] elements 302 in an overall computing system.

In a preferred embodiment, the control FPGA 412 may comprise a Virtex II XC2V6000 device available from Xilinx, Inc. containing six million 20 gates implemented with a direct memory access ("DMA") function to control the user FPGAs 416. Similarly, the user FPGAs 416 may also comprise XC2V6000 devices implemented with a 100 MHz clock phase locked loop ("PLL") for multiple frequencies. The on-board memory 414 may comprise six dual-ported memory banks of 24 Mbytes capable of operating at 100 MHz 25 with 64 bit data paths.

With reference additionally now to Fig. 6 a further, more detailed functional block diagram of the control block 412 of the MAP[®] element 302 illustrated in Figs. 4 and 5 is shown as a functional block 600. As depicted, the control block 412 is coupled to the switch/network adapter port 304 of the 30 processor board 306 as well as to the on-board memory 414 and user chips 416.

The control block 412 includes a command processor 602 for processing a command list ("ComList") and issuing DMA requests. A DMA engine for on-board memory 604 is coupled to the command processor 602 and is used to address the six memory banks of the on-board memory 414 in linear strides and provides sixteen entry irregular access thereto. Another DMA engine for external accesses 608 is also coupled to the command processor 602 and performs DMA operations to memories outside of the MAP[®] element 302. It too has strided capability and maintains tracks of out-of-order responses.

10 The DMA engine 608 is coupled to a packetizer 606 which forms thirty two byte packets for transmission and generates error correction codes ("ECC") to ensure the validity of the data transmitted and received. A depacketizer 610 is also associated with the DMA engine 608 and functions to check data packet validity as well as to check and correct ECC data. A 15 data register 612 and flag register 614 are coupled between the command processor 602 and the user chips 416. The former holds sixteen sets of data, with each set comprising thirty two 64 bit values while the latter holds thirty two 1 bit values.

20 The control block 412 also includes a number of storage locations which, in the exemplary embodiment illustrated, comprises storage 0 block 616 and storage 1 block 618 coupled between the command processor 602 and the DMA engine 608. The storage 0 block 616 is utilized to hold the first ComList as well as the first portion of the user configuration data for the 25 control block 412. The storage 1 block 618 holds the second ComList as well as the second portion of the user configuration data. A user chip configurator 620 is coupled to the command processor 602 and the user chips 416 to control the user logic bitsream and to configure each of the user chips 416₀ and 416₁.

30 With reference additionally now to Fig. 7, an additional, more detailed functional block diagram of a representative switch/network adapter port 304 is shown as a functional block 700 for possible use with the computer system illustrated in Figs. 3 and 4. The switch/network adapter port 304 may

conveniently be provided in a form factor to plug into a DIMM or RIMM slot of a microprocessor board 306 (Fig. 3) and provides for the interconnection of a MAP® element 302 to one or more microprocessors 402, e.g. 5 microprocessors 102 of Figs. 1, 2A and 2B. In this regard, the switch/network adapter port 304 may be coupled to exchange data and address information with, for example, a double data rate ("DDR") memory interface 702 to the one or more microprocessors 402.

The switch/network adapter port 304 comprises a packetizer 704 which is coupled to the MAP® element 302 to provide DMA header packet reply and 10 retry functionality as well as ECC checkbyte generation. A depacketizer 706 is also coupled to the MAP® element 302 and performs a DMA header packet sort, data request packet sort and ECC check and correct function on data received by the switch/network adapter port 304. A read buffer 708 is coupled to the depacketizer 706 and is provided with, in the exemplary 15 embodiment illustrated, 32 Kbytes or more capacity for holding 1024 cache lines of "read" data. A corresponding write buffer 710 is coupled to the packetizer 704 and may also have a capacity of 32 Kbytes or more for holding 1024 cache lines of "write" data.

A memory bus data multiplexer 712 is coupled between the data bus of 20 the DDR memory interface 702 and the read and write buffers 708, 710. The memory bus data multiplexer 712 functions as a data multiplexer between these elements while an address and command decode block 714 is coupled between the address bus of the DDR memory interface 702 and the multiplexer 712 to control data read/write functions of the memory bus data 25 multiplexer 712, the information to be maintained in status and configuration registers 718 and to send direct commands to the MAP® element 302. In this regard, a serial direct command formatter 716 is coupled to the address and command decode block 714 to provide MAP® element 302 direct command support.

30 With reference additionally now to Fig. 8, an illustration showing the interaction between the MAP® control block 412, switch/network adapter port 304 and one or more microprocessors 402 (e.g., of Fig. 4) is provided as an

example of the operation of a computer system utilizing a switch/network adapter port in accordance with the present invention.

In this representative “read” operation, the MAP® control block 412 first processes the ComList and issues a “read” request which includes the address and length of the data requested. This request is forwarded to the associated switch/network adapter port 304 where it is processed. The switch/network adapter port 304 updates its control registers and sends the address and length of the data requested on to a microprocessor 402. The microprocessor 402 has a DMA thread waiting for either a “read” or “write” request. It then processes the data request and transfers the requested data to the requesting switch/network adapter port 304. The switch/network adapter port 304 processes a data write of the received data and starts filling a first-in, first-out register for transfer of the data to the MAP® control block 412. The DMA engine 608 (Fig. 6) of the MAP® control block 412 receives the data and the data is written to the on-board memory 414 by the DMA engine 604. In the exemplary embodiment illustrated, the total latency for this entire operation is on the order of 1.0 nanoseconds from the time the MAP® control block 412 issued the request until the first stream of data is received.

While there have been described above the principles of the present invention in conjunction with specific system implementations and technologies, it is to be clearly understood that the foregoing description is made only by way of example and not as a limitation to the scope of the invention. Particularly, it is recognized that the teachings of the foregoing disclosure will suggest other modifications to those persons skilled in the relevant art. Such modifications may involve other features which are already known per se and which may be used instead of or in addition to features already described herein. Although claims have been formulated in this application to particular combinations of features, it should be understood that the scope of the disclosure herein also includes any novel feature or any novel combination of features disclosed either explicitly or implicitly or any generalization or modification thereof which would be apparent to persons skilled in the relevant art, whether or not such relates to the same invention

as presently claimed in any claim and whether or not it mitigates any or all of the same technical problems as confronted by the present invention. The applicants hereby reserve the right to formulate new claims to such features and/or combinations of such features during the prosecution of the present 5 application or of any further application derived therefrom.

What is claimed is:

CLAIMS

WE CLAIM:

1. A computer system comprising:
 - at least one processor;
 - 5 a controller for coupling said at least one processor to a control block and a memory bus;
 - a plurality of memory module slots coupled to said memory bus;
 - an adapter port associated with a subset of said plurality of memory module slots; and
- 10 a processor element coupled to said adapter port.
2. The computer system of claim 1 wherein said controller comprises an interleaved memory controller.
3. The computer system of claim 1 wherein said plurality of memory module slots comprise DIMM memory module slots.
- 15 4. The computer system of claim 3 wherein said adapter port comprises a DIMM physical format for retention within one of said DIMM memory module slots.
5. The computer system of claim 1 wherein said plurality of memory module slots comprise RIMM memory module slots.
- 20 6. The computer system of claim 5 wherein said adapter port comprises a RIMM physical format for retention within one of said RIMM memory module slots.
7. The computer system of claim 1 wherein said control block provides control information to said adapter port.
- 25 8. The computer system of claim 1 wherein said control block provides control information to said processor element.

9. The computer system of claim 1 wherein said control block comprises a peripheral bus control block.

10. The computer system of claim 9 wherein said peripheral bus control block provides control information to said adapter port.

5 11. The computer system of claim 9 wherein said peripheral control block provides control information to said processor element.

12. The computer system of claim 1 wherein said control block comprises a graphics control block.

13. The computer system of claim 12 wherein said graphics control
10 block provides control information to said adapter port.

14. The computer system of claim 12 wherein said graphics control block provides control information to said processor element.

15. The computer system of claim 1 wherein said control block comprises a systems maintenance control block.

15 16. The computer system of claim 15 wherein said systems maintenance control block provides control information to said adapter port.

17. The computer system of claim 15 wherein said systems maintenance control block provides control information to said processor element.

20 18. The computer system of claim 1 wherein said adapter port comprises:

a memory bus data multiplexer adapted to receive data information from said memory bus.

25 19. The computer system of claim 18 wherein said adapter port further comprises:

an address and command decoder adapted to receive address information from said memory bus.

20. The computer system of claim 18 wherein said adapter port further comprises:

a write buffer for coupling said data information received from said memory bus to said processor element; and

5 a read buffer for coupling data provided by said processor element to said memory bus.

21. The computer system of claim 20 further comprising:

a packetizer for coupling said write buffer to said processor element; and

10 a depacketizer for coupling said processor element to said read buffer.

22. The computer system of claim 18 further comprising:

a plurality of registers coupled to said memory bus data multiplexer.

23. The computer system of claim 22 wherein said plurality of registers comprise configuration registers.

15 24. The computer system of claim 22 wherein said plurality of registers comprise status registers.

25. The computer system of claim 19 further comprising:

a command formatter coupling said address and command decoder to said processor element.

20 26. The computer system of claim 1 wherein said processor element comprises a reconfigurable processor element.

27. The computer system of claim 1 wherein said processor element is operative to alter data received from said controller on said memory bus.

25 28. The computer system of claim 1 wherein said processor element is operative to alter data received from an external source prior to placing altered data on said memory bus.

29. The computer system of claim 1 wherein said processor element comprises:

a control block coupled to said adapter port.

30. The computer system of claim 29 wherein said processor element further comprises:

at least one field programmable gate array configurable to perform an identified algorithm on and operand provided thereto by said adapter port.

31. The computer system of claim 30 further comprising:
a dual-ported memory block coupling a control block coupled to said adapter port to said at least one field programmable gate array.

32. The computer system of claim 1 wherein said processor element comprises:
a chain port for coupling said processor element to another processor element.

33. The computer system of claim 29 wherein said processor element further comprises:

a read only memory associated with said control block for providing configuration information thereto.

34. A computer system comprising:
20 at least one processor;
an interleaved controller for coupling said at least one processor to a control block and a memory bus;
a plurality of memory slots coupled to said memory bus;
an adapter port associated with at least two of said plurality of memory 25 slots; and
a processor element coupled to each of said adapter ports.

35. The computer system of claim 34 wherein said plurality of memory slots comprise DIMM memory module slots.

36. The computer system of claim 35 wherein said adapter port comprises a DIMM physical format for retention within one of said DIMM memory module slots.

37. The computer system of claim 34 wherein said plurality of 5 memory slots comprise RIMM memory module slots.

38. The computer system of claim 37 wherein said adapter port comprises a RIMM physical format for retention within one of said RIMM memory module slots.

39. The computer system of claim 34 wherein said control block 10 provides control information to said adapter port.

40. The computer system of claim 34 wherein said control block provides control information to said processor element.

41. The computer system of claim 34 wherein said control block comprises a peripheral bus control block.

15 42. The computer system of claim 41 wherein said peripheral bus control block provides control information to said adapter port.

43. The computer system of claim 41 wherein said peripheral control block provides control information to said processor element.

44. The computer system of claim 34 wherein said control block 20 comprises a graphics control block.

45. The computer system of claim 44 wherein said graphics control block provides control information to said adapter port.

46. The computer system of claim 44 wherein said graphics control block provides control information to said processor element.

25 47. The computer system of claim 34 wherein said control block comprises a systems maintenance control block.

48. The computer system of claim 47 wherein said systems maintenance control block provides control information to said adapter port.

49. The computer system of claim 47 wherein said systems maintenance control block provides control information to said processor 5 element.

50. The computer system of claim 34 wherein said adapter port comprises:

a memory bus data multiplexer adapted to receive data information from said memory bus.

10 51. The computer system of claim 50 wherein said adapter port further comprises:

an address and command decoder adapted to receive address information from said memory bus.

15 52. The computer system of claim 50 wherein said adapter port further comprises:

a write buffer for coupling said data information received from said memory bus to said processor element; and

a read buffer for coupling data provided by said processor element to said memory bus.

20 53. The computer system of claim 52 further comprising:

a packetizer for coupling said write buffer to said processor element; and

a depacketizer for coupling said processor element to said read buffer.

25 54. The computer system of claim 50 further comprising:

a plurality of registers coupled to said memory bus data multiplexer.

55. The computer system of claim 54 wherein said plurality of registers comprise configuration registers.

56. The computer system of claim 54 wherein said plurality of registers comprise status registers.

57. The computer system of claim 51 further comprising:
a command formatter coupling said address and command decoder to
5 said processor element.

58. The computer system of claim 34 wherein said processor element comprises a reconfigurable processor element.

59. The computer system of claim 34 wherein said processor element is operative to alter data received from said controller on said
10 memory bus.

60. The computer system of claim 34 wherein said processor element is operative to alter data received from an external source prior to placing altered data on said memory bus.

61. The computer system of claim 34 wherein said processor
15 element comprises:
a control block coupled to said adapter port.

62. The computer system of claim 61 wherein said processor element further comprises:
at least one field programmable gate array configurable to perform an
20 identified algorithm on and operand provided thereto by said adapter port.


63. The computer system of claim 62 further comprising:
a dual-ported memory block coupling a control block coupled to said adapter port to said at least one field programmable gate array.

64. The computer system of claim 34 wherein said processor
25 element comprises:
a chain port for coupling said processor element to another processor element.


65. The computer system of claim 61 wherein said processor element further comprises:

a read only memory associated with said control block for providing configuration information thereto.


5

**Fig. 1**

PCI, AGP or SM Bus Connection

**Fig. 2A**

3/8

**Fig. 2B**

4/8






Fig. 3

Fig. 4

5/8

**Fig. 5**

6/8

**Fig. 6**




Fig. 7

8/8




Fig. 8